

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](http://www.nctatechnicalpapers.com).

ISBN Number: 0-940272-58-X  
©2019, NCTA – The Internet and Television Association.  
All rights reserved.

Jason Cole - Cox Communications

***Segment Routing and Enterprise: What It Is and Why It Matters..... 1***

Chris Busch; Dave Baran - CommScope

Jeff Finkelstein - COX Communications

***Creating The Intelligent Edge: Increasing DAA Velocity Using Service Orchestration ..... 30***

Colin Howlett; Douglas Johnson; Kai Meisen - Vecima Networks

***Delivering QAM Video in Distributed Access Architectures ..... 51***

Loay Kreishan - Charter Communications

***Offloading Data Using Unlicensed LTE (CBRS)..... 78***

Ron Hranac - Cisco Systems

James Medlock - Akleza, Inc.

Bruce Currivan - JJP Development

Roger Fish; Tom Kolze - Broadcom

Jason Rupe; Tom Williams - CableLabs®

Larry Wolcott - Comcast

***Characterizing Network Problems Using DOCSIS® 3.1 OFDM RxMER Per Subcarrier Data ..... 89***

Matthew Olfert - Shaw Communications

***Maximizing the Capacity and Reliability of a DOCSIS Network ..... 142***

Noé Morales; Asten Fenby - Shaw Communications Inc.

***The Role of Lean in Shaw..... 164***

Jay Liew; Mark Teflian; Bruce Bacon; Jay

Brophy; Randy Pettus - Charter Communications

***New Generation Data Governance for Charter Network:1 ..... 187***

Tom Holloran - Charter Communications

***Using MILP (Mixed Integer Linear Programming) for RF Bandwidth Optimization ..... 203***

Kashif Shakil - Ericsson

***Can Wireless Compete With Wired Access To The Home: A Review of Fixed Wireless Access Technology And Economics ..... 225***

Srilal M Weerasinghe; Robbie Mills III; Vipul Patel; Basil Badawiyeh; Mike Terada - Charter Communications

***Machine Learning Applications in Cable TV Advertising – Usage and Challenges.....248***

Patricio Sebastian Latini - CASA Systems

***An Analysis of How to Deploy Low Power WAN IoT Using HFC and Fiber Network Infrastructure.....260***

Dr. Claudio Righetti; Emilia Gibellini; Carlos

Germán Carreño Romano; Gabriel Carro - Telecom Argentina S.A.

***Can Future Networks Survive Without Artificial Intelligence?.....288***

Roger Brooks; Pankaj Kumar; Mudit Jain; Megha Vij; Nandit Jain; Andrew Colby - Guavus

***Customer First: CX-Driven Augmented Operations.....318***

Elaine Yeo - Charter Communications

***Segment Routing Proof of Concept for Business Services .....330***

Joe Rodolico - Comcast

***IoT Device Energy Harvesting Technologies and Implementations .....369***

Joe Rodolico - Comcast

***Methods to Maximize IoT Battery Life .....379***

Mohamed Daoud; Matthew Hubbard; Rajeev

Aggarwal; Hossam Hmimy - Charter

Communications

***On The Performance Of CBRS Fixed Wireless Access: Coverage And Capacity Field Study .....386***

Elliott Hoole; Joshua Sanders - Charter

Communications

***Building a Technology Platform for Smart Agriculture Deployments Using C-Band and Unlicensed Technologies.....427***

Ralph Bachofen - Triveni Digital

***ATSC 3.0: A Look at the Infrastructure and Possible Impact that Next-Gen TV Will Have on Cable Operations .....439***

Stuart Kurkowski, PhD - Comcast Incorporated

***Broadcast and Digital Evolution: The Evolution of Delivering to Any Screen.....446***



John Ulm; Tom Cloonan - CommScope  
***The Broadband Network Evolution Continues  
– How Do We Get To Cable 10g?..... 455***

Jason Rupe, Ph.D. - CableLabs®  
***A General-Purpose Operations Cost Model to  
Support Proactive Network Maintenance and  
More..... 486***

Joe Keller; Sam Plant - Cox Communications  
***Rethinking Customer Support - Proactive  
Customer Engagement: Experimentation in  
Real-Time Data ..... 519***

Patrick Goemaere - Technicolor  
***The Evolution of Network Virtualization In  
The Home: Improving User Experience And  
Manageability..... 532***

Tao Wan - CableLabs  
Mansour Ganji - Rogers Communications  
***Security Analysis Of 5G Mobile Networks  
..... 576***

Tao Wan; Max Pala - CableLabs  
Yildirim Sahin - Charter Communications  
***Authentication In 5G Wireline And Wireless  
Convergence..... 586***

Charuhas Ghatge - Nuage Networks, a Nokia  
Company  
***SD-WAN 2.0: A Platform for Multi-Cloud,  
Security and Value Added Services ..... 596***

Vipul Patel - Charter Communications  
Xavier Denis - CommScope  
***Scaling IP Advertising Using Manifest  
Manipulation..... 604***

Ian Wheelock; Charles Cheevers - CommScope  
***2019 Virtualized CPE Services Have Finally  
Arrived Via Service Delivery Platforms..... 630***

Matthew Tooley; Thomas Belford - NCTA – The  
Internet & Television Association  
***Detecting Video Piracy with Machine Learning  
..... 664***

David John Urban - Comcast  
***The Importance Of Wi-Fi 6 Technology For  
Delivery Of gbps Internet Service..... 689***

Arash Pendari; Giles Wilson - VionLabs AB  
Michael Eagles - Liberty Global  
***Next Player Video Service: The Case For  
Bringing Playlists to TV ..... 722***

Drew Davis - Cox Communications  
Anish Kelkar - Bell Labs Consulting  
***The Imperative of MSO Future Wireless... 740***

Mark Vogel - CommScope  
***HFC Spectrum Expansion: Design and  
Component Impacts..... 751***

Shaul Shulman - Intel Corporation  
***Operating Legacy Cable Modems in an FDX  
Environment ..... 765***

Zhensheng (Steve) Jia, Ph.D.; L. Alberto Campos,  
Ph.D.; Mu Xu, Ph.D; Haipeng Zhang, Ph.D.; Junwen  
Zhang, Ph.D; Chris Stengrim; Curtis Knittle, Ph.D. -  
CableLabs  
***Ultra Low-Cost Injection-locked FP Laser  
Source for Coherent Access Networks ..... 786***

Edouard Karam; Greg Spear - Accedian  
***Cost-Effective, Scalable Quality of Experience  
(QoE) Monitoring for SD-WAN Networks  
..... 806***

Venk Mutalik; Bob Gaydos; Dan Rice; Doug Combs  
- Comcast  
***Fifty Shades of Grey Optics: A Roadmap for  
Next Generation Access Networks ..... 824***

Michael Ting Wang, P. Eng. - Shaw  
Communications Inc.  
***Disaggregated, Coherent DWDM Solution at  
Shaw's Newest Cloud Datacentre Interconnect  
..... 856***

Derek Strauss - Shaw Communications Inc.  
***Operational Transformation: Modernizing  
Field Operations ..... 873***

Doug Jones - CableLabs  
***DOCSIS® 4.0 Technology Realizing  
Multigigabit Symmetric Services..... 887***

Kyle Haefner - Cable Television Laboratories Inc.  
***Predicting the Evolution of Distributed Denial  
of Service Attacks on Carrier Networks..... 905***

Fady Masoud, M. Eng. - Infinera <b><i>Preparing the Metro Core Network for Disruptive Technologies Like DAA and 5G</i></b> ..... 922	Andrew Bender - VMware <b><i>A Roadmap for Virtualization in HFC Networks: Use Cases and Considerations</i></b> .....1112
Alan Evans - EDGE GRAVITY by Ericsson <b><i>Why Gaming Needs An Edge</i></b> ..... 932	J.R. Flesch; Charles Cheevers; Kurt Lumbatis - Commscope <b><i>The Promise of WiFi in the 6 GHz Band</i></b> .....1127
John Douglas - Ericsson North America <b><i>Operational Impacts of Network Slicing, Leveraging Network Slicing Technologies to Offer Innovative Business Services</i></b> ..... 945	Bill Beesley - Fujitsu Network Communications <b><i>DAA, GAP, and Cloud Computethe Network of the Future</i></b> ..... 1154
Dr. Bill Wall; Michael Cooper; David Job - Cox Communications <b><i>Practical Considerations For Full Duplex Deployments In N+x Environments</i></b> ..... 954	Karthik Krishna - Nokia <b><i>How to deliver QAM video in a DAA world</i></b> .....1161
Javier Ger; Esteban Poggio; Miguel Masache Ojeda - Telecom Argentina Furquan Ansari; Ben Tang - Bell Labs Consulting <b><i>Telecom Argentina: Transport Network Evolution For Future Services</i></b> ..... 970	Marcin Godlewski - Technicolor <b><i>RDK All Access (Networks): DOCSIS, DSL, PON and Beyond</i></b> ..... 1176
Brian A. Scriber - CableLabs <b><i>Changing the World: IoT Chaos as a Ladder to Improving Security</i></b> ..... 992	Bryan Kelly - Comcast Cable <b><i>What Can Your CPE Tell You?</i></b> ..... 1187
Ryan Michael Cunningham - Comcast Corporation <b><i>The Pivotal Role of Cable Gateways in the Internet of Things</i></b> ..... 1003	John Ulm; Zoran Maricevic - CommScope <b><i>Cable 10G vs. Wireless 5G – Foe or Friend? A Survey of Next Gen Network Directions</i></b> ..1203
Asaf Matatyau - Harmonic, Inc. Brian Bendt - Comporium <b><i>Practical Lessons of a DAA Deployment with a Virtualized CMTS</i></b> ..... 1022	Nader Foroughi - Shaw Communications <b><i>Upgrading the Plant to Satisfy Traffic Demands: The One Touch Approach</i></b> ..... 1240
Roger G Stafford - Charter Communications, Inc. <b><i>The Generic Access Platform: What's in it for me?</i></b> ..... 1038	Henk Heijnen; Philippe Gilberton; Jean-Ronan Vigouroux - Technicolor France <b><i>Smart Home &amp; Smart Notifications</i></b> ..... 1278
Arun Ravisankar - Comcast Corporation <b><i>New Sensing Techniques For Advanced IoT Applications</i></b> ..... 1065	Maher Harb; Jude Ferreira; Dan Rice; Bryan Santangelo; Rick Spanbauer - Comcast <b><i>A Machine Learning Pipeline for D3.1 Profile Management</i></b> ..... 1311
John Gammons - Cox Communications <b><i>Defining the Premise and the Edge: A Managed Wi-Fi Service Application</i></b> ..... 1081	Karthik Sundaresan; Jay Zhu; Mayank Mishra - CableLabs James Lin - Kyrio/CableLabs <b><i>Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA)</i></b> .....1354
Cliff Lane; Kishan Ramaswamy - VMware <b><i>How to Leverage SD-WAN to Accelerate Time to Market and Revenue</i></b> ..... 1100	Avinash Raghavendra - Ericsson North America <b><i>Containers – To Use It or Not to Use It...</i></b> 1394
	Hani Beshara - Ericsson Inc. <b><i>The Evolution Of Cellular IoT</i></b> .....1403

Steve Condra; Kari Maki; Arttu Purmonen - Teleste Intercept LLC

***Extended Spectrum DOCSIS®: A Pragmatic Approach..... 1416***

Thomas Priore; Kevin Alcox - iEldra

***Connected Independence..... 1431***

Brian A. Scriber - CableLabs

***Opting-In: Designing Privacy Tracking for Consumer Confidentiality & Cryptographic Assurance for Enterprises..... 1449***

Rajiv Asati; Alon Bernstein – Cisco Systems

***Cable Edge Compute: Transforming Cable Hubs into Application-Centric Cloud..... 1462***

Alon Bernstein; Rajiv Asati; Sangeeta Ramakrishnan - Cisco Systems

***Winning the Gaming War: Play for Cable Operator..... 1496***

J. Clarke Stevens - Shaw Communications

***Building a Cable-Friendly Internet of Things ..... 1526***

Ron Wolfe - Charter Communications, Inc.

***Layer 1 Considerations for Extended Spectrum Utilization in Hybrid Fiber Coax & Distributed Access Architecture Networks ..... 1547***

Joe Mocerino - Fujitsu Network Communications

***5G Backhaul/Fronthaul Opportunities and Challenges..... 1559***

John T Chapman; Hang Jin - Cisco Systems

Thushara Hewavithana - Intel Corporation

Rainer Hillermeier - Qorvo

***Blueprint for 3 GHz, 25 Gbps DOCSIS® ..... 1567***

Tong Liu, PhD; John T Chapman - Cisco Systems Inc

***R-PHY with Remote Upstream Scheduler ..... 1618***

Garey Hassler - Comcast

***Delivering the Highest IP Video Quality Efficiently While Improving Customer Experience..... 1643***

Pawel Sowinski; Andy Smith; Tong Liu - Cisco Systems Inc.

***Remote PHY 2.0 - The Next Steps For Remote PHY Technology..... 1654***

Kjell Johansson - Ericsson Inc

***Capitalizing On The Evolved Communications Experience..... 1674***

Derek Rieckmann - Midco

***Operational Transformation Using GIS.. 1699***

Matt Carothers; Damien Whaley - Cox Communications

***Customer Safety Initiative (CSI) ..... 1706***

Craig Schwechel; Marth Wilson; Lauren Buhl; Tomislav Marcinko - inCode, a division of Ericsson

***Mid-band Spectrum Opportunities And Challenges..... 1722***

Greg White; Karthik Sundaresan; Bob Briscoe - CableLabs

***Low Latency DOCSIS: Overview And Performance Characteristics..... 1742***

Yair Neugeboren; Greg Cyr; Chris Zettinger - CommScope

***Experiment Results for Supporting LTE-FDD, LTE TDD, and 5G Timing Synchronization Over DOCSIS CAA and DAA..... 1769***

Tom Cloonan; Ayham Al-Banna; Frank O’Keeffe;

John Ulm - CommScope

Ruth Cloonan - BlueOpus

***Capacity Planning, Traffic Engineering, and HFC Plant Evolution for the Next 25 Years ..... 1798***

Andrew Joseph Milley - Cox Communications Inc.

***Proactive Customer Maintenance ..... 1855***

Jason Rupe, Ph.D.; Jingjie Zhu - CableLabs®

***Kickstarting Proactive Network Maintenance with the Proactive Operations Platform and Example Application ..... 1870***

# **Segment Routing and Enterprise: What It Is and Why It Matters**

## **Replacing LDP in the Metropolitan Network**

A Technical Paper prepared for SCTE•ISBE by

**Jason Cole**  
Senior IP Engineer  
Cox Communications  
Atlanta, GA USA  
404-971-9307  
jason.cole@cox.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
1. Introduction .....	4
1.1. Purpose .....	4
1.2. Scope .....	4
1.3. Prerequisites .....	4
2. Segment Routing .....	5
2.1. History .....	5
2.2. Evaluation .....	6
2.2.1. RSVP .....	7
2.2.2. LDP .....	7
2.3. Value .....	9
3. Hardware and Software .....	9
3.1. Platforms .....	9
3.2. Label Depth .....	10
3.2.1. TI-LFA .....	10
4. Global Block .....	12
4.1. Overview .....	12
4.2. Planning .....	13
4.2.2. Reuse .....	14
4.2.3. Consistency .....	16
4.2.4. Size .....	16
4.2.5. Derivation .....	18
4.2.6. Partitioning .....	20
5. Interworking .....	20
5.1. Overview .....	20
5.2. Mapping Server .....	22
6. Implementation .....	23
6.1. Prerequisites .....	23
6.2. Current .....	24
6.3. Dual Stack .....	25
6.4. Mapping Server .....	25
6.5. Protocol Preference .....	26
6.6. Interworking .....	27
Abbreviations .....	28

## List of Figures

Title	Page Number
Figure 1 - The four domains and their protocols .....	5
Figure 2 - Current residential topology and metrics used by Phoenix .....	11
Figure 3 - Inter-domain services using BGP-LU and Inter-AS Option A .....	12
Figure 4 - SRGB and index example .....	13
Figure 5 - Reuse within an autonomous-system using a SR-PCE and explicit hops .....	16
Figure 6 - Deriving indexes via IP addresses with an SRGB size of 256K .....	19
Figure 7 - Ships in the night .....	21

Figure 8 - Interworking terminology in the SR IGP area.....	22
Figure 9 - Label swapping.....	22
Figure 10 - Sample topology - current state.....	24
Figure 11 - Sample topology - Dual Stack.....	25
Figure 12 - Topology - mapping server .....	26
Figure 13 – SR preference should not impact transit LSPs.....	27
Figure 14 – Same topology – Protocol Preference .....	27
Figure 15 - Topology - Interworking.....	28

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Cox's domain protocols.....	6
Table 2 - Summary of technology comparison .....	8
Table 3 - SRGB parameters summarization .....	17
Table 4 - Node requirements. Largest IS-IS L1 areas.....	17
Table 5 - 64K indexes.....	18
Table 6 - 32K indexes.....	18
Table 7 - 8K indexes.....	18
Table 8 - MPLS Nodes and Index Requirements by Domain .....	20

# 1. Introduction

## 1.1. Purpose

This document details the deployment of Segment Routing in the residential domains. It lays the foundation for the sustaining engineering team to begin low-level planning and testing. The earliest sections of the document explain the value of Segment Routing; primarily in the residential domain. This is done by comparing SR's technical features against LDP and RSVP-TE. Most of this document describes the operational challenges of implementation. The challenges are arranged into three categories:

1. Hardware and Software
2. Global Block (SRGB) Standardization
3. Interworking

Matrices describe vendor software support, as well as hardware and software readiness of each market. The same section details why Segment Routing will not be deployed on all hardware, even though supported in code. Recommendations for code versions are provided for initial lab testing. The document also describes each platform's label depth, factors impacting label depth and why, for the foreseeable future, label depth will not be a concern in the residential network. Index assignments, which are derived from Segment Routing Global Blocks (SRGB), require planning. This document describes factors impacting planning and will recommend the assignment method. Last of the challenges is interworking. This will be illustrated in various drawings and described using CLI configurations. The concluding section will describe the implementation process and a FOA recommendation. It will detail the steps required to get to a "SR Core" in a residential market, using illustrations and configurations.

## 1.2. Scope

All residential domains will be impacted by these changes. The business and data center domains are not specifically covered by this document. In the initial deployment, only MX and NCS chassis will be impacted. At the time of this writing, there is a plan to integrate a Nokia SR-1 platform as the replacement for the ASR9K COI router. Depending on when the Segment Routing project is implemented, the scope will include that platform. The same could be said about a replacement chassis for the Cisco Nexus router. The Cisco Nexus is currently IP-only (no MPLS) but will likely be replaced by a MPLS enabled device.

## 1.3. Prerequisites

Readers of this document require a solid understanding of SR-MPLS. This document was not created to explain the technology. Rather, it describes the high-level application and deployment process. This document has not validated code and only serves to provide a code *recommendation*.

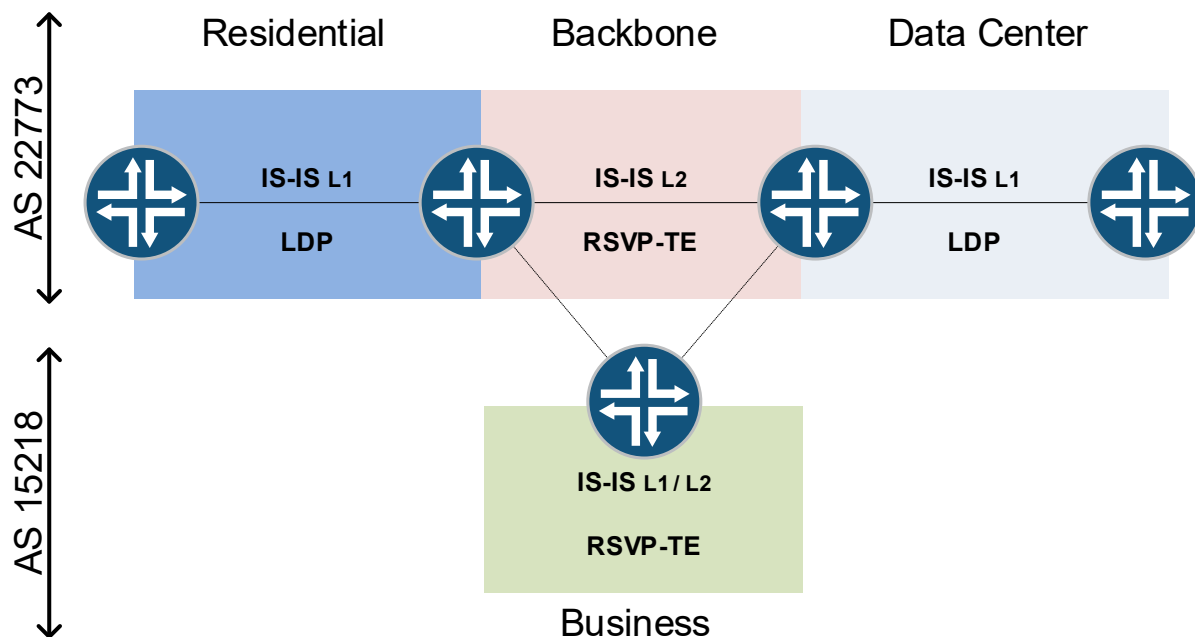
**Disclaimer:** Code testing will be performed by the sustaining engineering team, so no specific code mentioned in this document is guaranteed to be used for deployment.

## 2. Segment Routing

### 2.1. History

Because SR (RFC 8402) was designed to replace both RSVP-TE (RFC 3209) and LDP (RFC 3036) for most service provider deployments, RSVP-TE and LDP are briefly described.

RSVP-TE is a feature-rich MPLS protocol. Its most commonly deployed solely for Fast-Reroute (FRR). In fact, a survey<sup>1</sup> indicates that 90% of RSVP-TE deployments are solely for FRR<sup>2</sup>. It's less commonly deployed to leverage advanced features such as bandwidth reservation and auto-bandwidth. Cox deploys RSVP-TE for both reasons described previously. The business domain requires FRR for Metro-E services and the backbone domain leverages FRR, bandwidth reservation and auto-bandwidth for lowering cost in the long-haul fiber network. LDP is, for the most part, a simple protocol. It was the first plug-and-play, non-proprietary MPLS protocol. It provides the immediate benefits of enabling MPLS in a service provider network. These have traditionally been “BGP-free core” and transport for MPLS-based services (L2VPN/EVPN/L3VPN). Features have continued to be developed even after LDP's inception, but most development has slowed or stopped. Modern LDP code supports IPv6 (RFC 7552), FRR (RFC 8102) and multicast (RFC 6826). These features will be discussed in a later sub-section.



**Figure 1 - The four domains and their protocols**

Segment Routing is emerging as an elegant technology. It was designed to operate with the simplicity of LDP – being plug-and-play. It also provides *native* support for features like IPv6 and FRR. SR was also designed to support advanced features, such as On-Demand Next-hop (ODN), Flexible-Algorithm and multicast, by leveraging a SR-Path Computational Element (SR-PCE). By decoupling some control-plane functions, the protocol becomes even more flexible. It was designed for this at its inception. Some of the

<sup>1</sup> Insert here



advanced features for SR are still in development. Table 1 below describes each domain within the Cox network. The “next-generation protocol” is the strategic vision of Cox’s architecture.

**Table 1 - Cox's domain protocols**

Domain	Existing Protocol	Next-Gen Protocol	SR-MPLS Comments
Backbone	IS-IS L2 / RSVP-TE (FRR / Bandwidth Reservation / Auto-BW)	RSVP-TE	Requires SR-PCE Requires multicast Requires bandwidth reservation Requires high visibility (statistics)
Residential	IS-IS L1 / LDP	SR-MPLS (IPv4/IPv6)	Superior to LDP Requires SRMS
Data Center	IS-IS L1 / LDP	SR-MPLS (IPv4/IPv6)	Superior to LDP
Business	IS-IS L1, L2 / RSVP-TE (FRR)	SR-MPLS (IPv4/IPv6)	Superior to LDP More scalable SR-PCE use-case (inter-area LSPs)

## 2.2. Evaluation

This sub-section does not compare advanced features of Segment Routing like ODN and Flexible-Algorithm. These features are not applicable to the residential domain, and it is undetermined whether they’ll be leveraged in the backbone and business domains, as of the time of this writing. A brief description of these two technologies is provided to reduce obscurity. ODN leverages a SR-PCE to provide a label stack to an ingress LSR, when the LSR does not have reachability information about the egress LSR. This scenario occurs when both LSRs are in different domains (IGP areas or autonomous-systems) and an end-to-end LSP is required. The most likely application of this is in the business domain. Flexible-Algorithm uses interface monitoring probes to derive delay. These can then be translated and used to create a dual plane, IGP topology which may route based on delay. To achieve this, delay-based metrics are created and advertised by the IGP. Then, the nodes create an additional topology map which is independent from the standard, static-metric topology. The ingress LSR will associate a MPLS service to the label stack derived from either topology. The most likely application of Flex-Algorithm are the backbone and business domains. These are advanced features which distinguish Segment Routing from LDP and RSVP-TE, but they’re not applicable to the residential domain.

### **2.2.1. RSVP**

To understand SR and its holistic deployment in the Cox network, a few points will address the practical differences between RSVP-TE and SR. When RSVP-TE is used solely for FRR, as is the case in Cox's business domain, SR is the preferred protocol. SR is plug-and-play whereas RSVP-TE requires tunnel configurations. The elimination of MPLS tunnel configuration simplifies automation. RSVP-TE maintains lots of network state. When a network element fails, it creates churn and overhead. Depending on the number of LSPs, this can be quite significant. Segment Routing reduces state, configuration, complexity and convergence time, especially when replacing RSVP-TE. There are many other caveats to RSVP-TE when used solely for FRR. The points above are a shortlist which are most relevant. The backbone domain's deployment of RSVP-TE will not be directly compared against Segment Routing, because a one-for-one swap is not technically feasible at this time, but it is something that will be investigated in the near-term future.

### **2.2.2. LDP**

The residential and data center domains in the Cox network run LDP. LDP adds an additional layer of complexity. It is often referred to as a "piggyback protocol", because LDP relies on IGP entries in the RIB to perform label allocations. Thus, it can be assumed that the IGP has already converged before the LDP process begins. The process begins with each node allocating labels for the prefixes in its RIB. The label information is stored and maintained in the LDP binding table. With LDP, the number of labels is proportional to the number of nodes and links. However, this can be adjusted so that only host routes are allocated labels. Label information is exchanged via LDP adjacencies on a hop-by-hop basis like distance vector routing protocols which "route-by-rumor". Each node in the domain performs the same task until the network has converged. Whenever a network element fails, LDP relies on the piggyback process to converge. This increases the overall convergence time.

Segment Routing eliminates a layer of protocol machinery required for label distribution – LDP is no longer required for label allocation, storage and exchange. SR is unique from LDP in that it leverages the IGP to propagate and store label information. This makes the network more efficient, thus reducing convergence time. In Segment Routing, the number of labels is proportional to the number of global SIDs. Leveraging the IGP for label distribution and storage reduces the overall cost on the protocol machinery. SR inherently eliminates LDP-IGP synchronization issues.

There are three other key features which are supported by both LDP and Segment Routing that may be immediately deployable or are of some interest to Cox in the future. These three features are: MPLS for IPv6, Fast Re-Route (FRR) and MPLS multicast. LDP has been extended to support LDPv6, R-LFA and mLDP. These enhancements are for the most part successful and effective. However, some LDP extensions are not as efficient as Segment Routing's built-in features. Segment Routing's equivalent features are: SR-MPLSv6 (not "SRv6"), Topology Independent-Loop Free Alternate (TI-LFA) and Tree-SID. The three features require a brief comparison.

#### **2.2.2.1. IPv6**

IPv6 support in LDP (LDPv6) creates additional state in the network; both LDPv4 and LDPv6 processes run concurrently. This means that additional adjacencies and tables are required for label exchange and storage. SR uses a small TLV for the IPv6 node-SID, which is for the most part identical to the IPv4 node-SID. It is transported in the same IGP advertisement. It requires less memory for state and less information is exchanged than LDPv6. Thus, SR-MPLSv6 is more efficient than LDPv6.

### 2.2.2.2. Multicast

Cox's current multicast deployment in the residential and data centers domains is referred to as "IP multicast" and not "MPLS multicast"; multicast forwarding is done via PIM and multicast packets are not label switched. Multicast LDP (mLDP) is an effective and simple protocol. However, this add-on, like LDPv6, adds additional state to the LDP binding database that is not required using SR. SR's Tree-SIDs avoid additional overhead on nodes by using a SR-PCE. Thus, mLDP adds more cost to the routers than Tree-SIDs. However, Tree-SIDs move cost from the routers to a SR-PCE. An alternative to Tree-SIDs is Bit Index Explicit Replication (BIER). Prototypes have shown that it may be the next-generation MPLS multicast protocol of choice, due to its efficient addressing schema. BIER is recommended to be run in tandem with SR for maximum efficiency. As of the time of this writing, MPLS multicast is not planned for implementation in the residential and data center domains. This comparison is provided for a comprehensive analysis.

### 2.2.2.3. Fast-Reroute

TI-LFA provides 100% link and node protection and micro-loop avoidance in all topologies without the need for T-LDP, unlike R-LFA. Another plus for TI-LFA is that it will always route traffic on the post-convergence path, where R-LFA may have to move the protected traffic off the repair path and then onto the post-convergence path. It is worth noting that SR supports SRLG and TI-LFA can account for SRLGs, unlike R-LFA. As a side note, both TI-LFA and R-LFA support link and node protection. Contrary to the previous points describing TI-LFA's benefits, Cox's residential and data center domains use hub-and-spoke topologies. Such a topology should be fully covered by R-LFA, because there should only ever be two paths out of a hub-site. For instance, regarding TI-LFA's post-convergence optimization, the repair path should also be the post-convergence in link or node failure scenarios using R-LFA. Thus, the primary benefit of TI-LFA in the residential and business domains is micro-loop avoidance and not post-convergence optimization.

**Table 2 - Summary of technology comparison**

Feature	SR-MPLS	LDP
Configuration	IS-IS	LDP & IS-IS
Synchronization	Not required	IGP/LDP
State Information	No state	Minimal state
Protection	Link/Node (TI-LFA), SRLG	Link/Node (R-LFA) using T-LDP
IPv6 Support	IPv6 over SR	LDPv6
Multicast/P2MP	Not supported as of 2018	Supported
Anycast	Supported	Not supported
Flex Algorithm	Supported	Not supported
Label Learning	Leverages IGP	Piggybacking protocol
No. of Labels	SIDs proportional to nodes	Proportional to nodes or nodes & links
Label Scope	Local or Globally Unique	Locally unique
Egress Peering Engineering	Supported	Not Supported

## 2.3. Value

Traditional MPLS networks lack flexibility and programmability. Segment Routing provides this, at scale, with advanced capabilities such as On-Demand Next-Hop, Service Disjointness, Flexible-Algorithm and Service-Chaining. These advanced features, along with others, can be leveraged using a SR-PCE deployed in a centralized or hybrid deployment model. New capabilities supported in Segment Routing, which are being driven by new customer services, are causing network architects to prepare their infrastructure for the future. The new infrastructure should be able to provide end-to-end paths, between multiple domains, that meet higher-bandwidth and lower-latency requirements than before.

Although Segment Routing provides a new architecture for service providers, it will be deployed in its most basic form in the residential domain. This will introduce support for IPv6 and FRR. These features should be standard in today's networks. It will eliminate LDP-IGP synchronization outages. It is plug-and-play and will require minimal configuration<sup>3</sup>. This will remove one layer of protocols and simplify the network. However, the transitional state, which will run both LDP and Segment Routing, will introduce complexity. This complexity can be overcome, if new software/tools and proper training is provided. Given the Segment Routing architecture, it makes sense for Cox to begin migrating its network domains to Segment Routing to prepare for the future.

## 3. Hardware and Software

### 3.1. Platforms

There are five platforms in the residential domain which are Segment Routing capable. The ASR9K is SR capable, however, it has been decided to not enable SR on it; this is explained in a later sub-section. Also, the QFX<sup>4</sup> and SR-1 have not yet been deployed in the Cox network, so these are not an immediate concern. The five platforms are listed below:

1. Vendor A / platform A
2. Vendor A / platform B
3. Vendor B / platform A
4. Vendor B / platform B
5. Vendor C / platform A

Each platform's support for SR is detailed in the table below. "SR Support" does not include sub-features required to support Cox's deployment of SR; it only indicates minimum support. Also detailed in the table below is Cox's current and future codes. There is a subset of SR features that each platform's code should support. Cox should plan to deploy Segment Routing based on the dates of these feature's availability. There are some features which are mandatory and others which are not. Mandatory features should determine target dates. Optional features are flexible when determining target dates.

The Juniper QFX will be deployed for two different use-cases: Remote-Phy (R-PHY) and service aggregation<sup>5</sup>. When deployed for R-PHY, the QFX does not require MPLS. Thus, Segment Routing is not applicable in this scenario. Using the QFX for service aggregation is still in the high-level design phase. However, it will be deployed with MPLS, unlike the R-PHY deployment. Thus, Segment Routing should

---

<sup>3</sup> A Segment Routing Mapping Server will be required to interwork LDP/SR. This requires using a static one-to-one or dynamic block mapping between a label(s) and prefix(s)

<sup>4</sup> Deploying the QFX as a service layer router (SLR) as a Nexus replacement is a current discussion.

<sup>5</sup>Currently, this has not been finalized.

be enabled on the QFX in this instance. The two deployments are not discussed any further in this document, and the specific design documents should be referenced for further details

The Nokia SR-1 is the replacement for the Cisco ASR used for COI services. Currently, the ASR used for COI may also terminate Metro-E services. This will not be the case with the SR-1. The current design, which is near finalization, will deploy it identically to the Cisco ASR. Thus, it will be in the residential domain and MPLS enabled. Because this device is required to only support Internet traffic, it does not technically require MPLS now.

## **3.2. Label Depth**

A platform's supported label depth is a concern with SR-MPLS. SR will increase the size of the label stack in the following cases: TI-LFA, SR-TE and traffic accounting. In Cox's residential deployment, TI-LFA is the concerning factor. Label depth is limited by both software and hardware. Typically, large differences in label depth limits are seen in platforms leveraging merchant and custom silicon. "Silicon", in the context of this sub-section, refers to chipsets, or components, within a platform that impose forwarding limitations like label depth. The time it takes for a vendor to adjust forwarding capabilities, like label depth, is often determined by whether the platform is merchant or custom silicon based.

Merchant silicon is designed and manufactured by a 3<sup>rd</sup> party – not by a vendor. The 3<sup>rd</sup> parties produces platform components, like ASICs, that adhere to industry standards. A vendor will then implement it as a component within their platform. Leveraging a 3<sup>rd</sup> party reduces cost. However, 3<sup>rd</sup> party manufacturers tend to be slower to adapt to industry changes. Examples of platforms using merchant silicon are the Juniper PTX 10K, Juniper QFX 5K and Cisco NCS 5K.

Custom silicon is completely designed and produced in-house by a vendor. This often makes a product costlier than if it were to leverage merchant silicon. However, it allows the vendor full control and flexibility to modify their platform to meet new requirements as the industry changes. These platforms tend to be richer in features. They also provide more accountability, because the vendor owns the entire product. The Juniper MX, Juniper QFX 10K and Cisco ASR are examples of custom silicon.

### **3.2.1. TI-LFA**

Metro Ethernet over HFC is a commercial VPN service that exists in the residential domain. Metro-E over HFC, along with voice services, are drivers for implementing TI-LFA. TI-LFA inherently protects label switched traffic. Because video services are IP multicast in the residential domain, they will not be protected. Services may be disrupted by fiber cuts or node failures, when a fast re-route mechanism is not in place. TI-LFA can operate in either link protection or node protection modes.

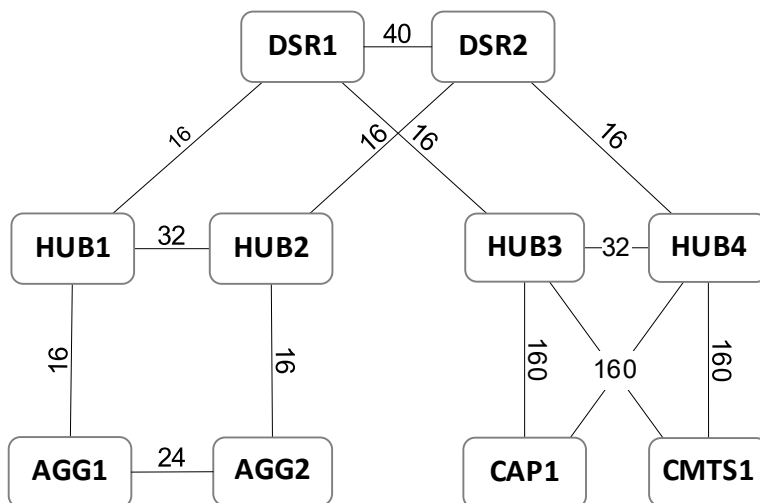
Both link and node protection leverage additional SIDs, ultimately more MPLS labels, to route around link or node failures in under 50ms. The question becomes, "how many labels are required for link protection and node protection?" The case study below provides insight to this answer. Clearly, link and node protection have different label requirements. Also, the number of labels required by either protection mechanism varies. The additional labels for FRR capabilities adds cost and complexity *to the data-plane*. This is different from traditional FRR techniques provided in RSVP-TE for example, which add additional network state *to the control-plane*.

The topology and metrics of a network determine the number of SIDs required in either mode. The residential and business domains are "hub-and-spoke" topologies with static metrics. The number of labels required for each mode was found by modeling the deployment of both using the Phoenix market's IS-IS

database, because it contains the most network elements. The results are below and the model has been attached.

Node Protection SIDs: 4

Link Protection SIDs<sup>6</sup>: 2



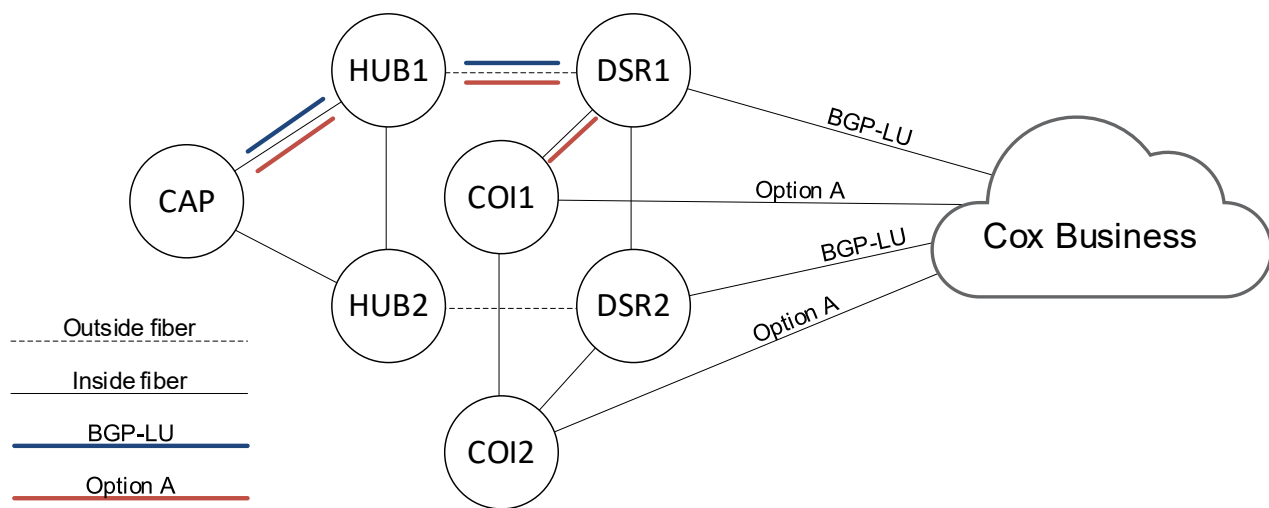
**Figure 2 - Current residential topology and metrics used by Phoenix**

Once the number of labels required for both protection modes was found, the cost of additional labels, per protection mode, had to be weighed against the likelihood of network element failures within the path of a Metro-E service. As this time, the majority of Metro-E over HFC services are type E-LINE and E-TREE. The remote ends of the service most commonly terminate in the business domain. Cox has two different methods for inter-domain services: BGP-LU and L2VPN Inter-AS Option A. These methodologies and designs are beyond the scope of this document. The key relating to TI-LFA is that the LSPs for the service terminate between the CAP/CMTS and DSR/COI router. The network elements in this path carry the most weight in the decision.

The illustration below describes two LSPs for both inter-domain methods. A service can use either LSP, depending on how it is built. CAP is a headend for both LSPs. The blue LSP's tailend is the DSR. The red LSP's tailend is COI. Both LSPs traverse the link which is most susceptible to fiber cuts (HUB-DSR). None of the nodes associated to these two LSPs are deployed with a non-redundant control-plane (RE/RSP/supervisor)<sup>7</sup>. Thus, node protection provides little value. It is worth noting that only "P" nodes without redundant control-planes are relevant to node protection. For instance, if the ingress or egress nodes have a non-redundant control-plane, that failure would cause the service to go down entirely; TI-LFA provides no protection in that case.

<sup>6</sup> Technically, the model revealed that 3 labels were required, but this was due to an incorrect metric on a link. So, 2 is the true number based on a standard topology and metrics.

<sup>7</sup> The business team plans to deploy a Nokia SR-1 as the new COI platform. This platform does not provide control-plane redundancy.



**Figure 3 - Inter-domain services using BGP-LU and Inter-AS Option A**

## 4. Global Block

### 4.1. Overview

MPLS nodes use the Label Switching Database (LSD) to store labels. The total number of labels supported by the LSD is defined by the operating system, it is normally fixed and only adjustable in software upgrades. The LSD is partitioned such that “label spaces”, which are unique ranges of labels, are defined for different MPLS applications. Examples of MPLS applications include: RSVP-TE, SR, LDP, BGP-LU, L3VPN, L2VPN, etc. Normally, applications share label spaces, and they’re most commonly group by their application type: “service” or “transport”. When an application requires a label, it will make a request to the label manager, the process responsible for allocating labels to the applications, it will receive one from its designated label space. Segment Routing is unique from traditional MPLS applications, because a unique label space must be defined by the network operator that only SR-MPLS will use. This is known the Segment Routing Global Block (SRGB).

In SR-MPLS, there are two general types of label spaces. The SRGB is used for *global* labels like node-SIDs. Node-SIDs are globally unique identifiers which require static assignments. Thus, planning is required. The Segment Routing *Local* Block (SRLB) is used for *local* labels like Adjacency-SIDs. Adjacency-SIDs are dynamically assigned by the router without any type of operator configuration. The SRLB is the term used to describe the existing, “dynamic” label space. This label space is shared with other MPLS applications; only the SRGB label space is dedicated to SR-MPLS. SID type details can be found in RFC8402.

Each node in an IGP area, which is enabled with SR-MPLS, requires static label assignments from the SRGB. These label assignments are referred to as “indexes” and are relative to their position within the SRGB range. The example below provides an example of a label assignment using a SRGB and index.

If SRGB is [16,000 (min) - 23,999 (max)]:  
Index Range = 1 – 7,999

If index is 1:  
Node Label = 16000 + 1

#### **Figure 4 - SRGB and index example**

Each vendor defines their platform's label spaces differently. There are several attributes that may differ between vendors. A few significant examples: the number of unique label spaces, the size of the labels spaces, association of MPLS applications to label spaces and the default label spaces.

## **4.2. Planning**

This sub-section explains the planning of the SRGB and index assignments for all domains at a high-level and provides recommended SRGB and index assignment methods. To do this properly, the network should be approached holistically like IP addressing. The approach assumes Segment Routing in *all* network domains – backbone, data center, residential, business. The proceeding sub-sections will describe:

1. Features requiring assignments
2. Reuse
3. Using the same SRGB
4. Selecting the size of the SRGB
5. Deriving index values
6. Partitioning the indexes/SRGB

### **4.2.1.1. Mapping Server (SRMS)**

SRMS is currently in draft<sup>8</sup> but is supported by almost all vendors already. The Segment Routing Mapping Server feature allows for LDP/SR interoperability. Using IS-IS TLVs for Segment Routing, each SR-MPLS node advertises its SRGB. Each node also advertises the loopback prefix and node-SID binding within the IS-IS TLV. Thus, any node not advertising this is not SR-MPLS enabled. Nodes which do not support Segment Routing are allocated a node-SID by the mapping server and this is advertised using another IS-IS TLV. These allocations can be done in two ways: using 1-to-1 (node-SID to node) or using a pool of node-SIDs. Either way, a block of node-SIDs must be reserved for the SRMS(s). It is required that Cox plan for SRMS index allocation in the residential domain.

### **4.2.1.2. Flexible Algorithm (Flex-Algo)**

Flexible Algorithm is currently in draft<sup>9</sup>. Flex-Algo can be used to create a dual-plane topology. Traditionally, there has been a “single-plane”, or “single topology”. A topology consists of nodes, links and metrics. Metrics often reflect the bandwidth of a link. In this case, the preferred path is the one with the most bandwidth. Cox does not use bandwidth-based metrics in the residential domain. Rather, static metrics are used to try to load-balancing or avoid certain links. In either case, a “single-plane” is constructed and used to route labeled and unlabeled traffic.

---

<sup>8</sup> <https://datatracker.ietf.org/doc/draft-ietf-spring-segment-routing-ldp-interop/>

<sup>9</sup> <https://datatracker.ietf.org/doc/draft-ietf-lsr-flex-algo/>



“Dual-plane” introduces an additional topology. A second topology can be constructed using the same nodes and links. The metric, or “constraint”, can be unique to that topology. For instance, the metric for the second topology could be based on delay. Thus, two topologies exist. One topology based on bandwidth. Another based on delay. The shortest path between two nodes could be different whether you route based on bandwidth or delay. A service can then be associated to either topology – service to LSP mapping.

#### 4.2.2. Reuse

The planning of SRGBs and SIDs is like IP addressing. Like IP addressing, SIDs *can* be reused. However, there are restrictions, similarly to when duplicate IP addresses are assigned. For instance, duplicate IP addresses can be used when the nodes do not require direct communication. Issues may arise down the line when requirements change. This may cause a massive undertaking to re-IP the nodes<sup>10</sup>. In an ideal SR-MPLS deployment, every node in the autonomous-system is assigned a unique SID and a consistent SRGB. This allows end-to-end, inter-domain LSPs via IGP route leaking or BGP-SR. However, when reuse is leveraged, inter-domain LSPs via these two methods is not available and only a SR-PCE can accomplish this. Currently, Segment Routing is in its infancy and side-effects from reuse are not fully realized. Thus, it is a best practice to avoid reuse when possible.

Reuse allows SID assignments to scale “infinitely”. It is only the largest carriers, having a single domain requiring more SIDs than the largest available SRGB, that may *intentionally and initially* design for reuse.

A completely different carrier may *unintentionally* reuse SIDs when the number of required SIDs exceeds the number of available SIDs than the *existing* SRGB can accommodate. Growth in SID requirements may be caused by drastic, unanticipated changes in the network. For instance, applications<sup>11</sup>, planning<sup>12</sup>, mergers, acquisitions, collapsing<sup>13</sup>, etc. It is challenging to plan for changes that are unlikely to occur. By doing so, the designer may over plan and create an entirely different set of challenges. Rather than completely avoiding planning for unanticipated, drastic, network changes, it is recommended that Cox plans for reuse but only “tactically”<sup>14</sup>. In this case, reuse is only to be leveraged when an emergency arises.

For example, suppose the number of LSRs (MPLS enabled routers) in ASN22773 is 2,000. A SRGB supporting 16,000 SID assignments accommodates this requirement, as well as provides room-to-grow and flexibility in planning. Suppose the business domain collapses into the residential domain and 1,000 new LSRs are installed in each market in ASN22773. If Cox has 20 markets, 20,000 additional, unique SIDs are required. Thus, a total of 22,000 assignments are needed. The current SRGB of 16,000 configured on LSRs cannot accommodate this. It is a large undertaking to change the SRGB on the 2,000 existing LSRs<sup>15</sup>. If tactical reuse is initially planned for, it accommodates this requirement. Tactical reuse is described in further detail in this section.

---

<sup>10</sup> At one point, Cox used duplicate IP addresses in markets, when it was expected that the nodes would not require direct communication. Since that time, extensive work has been done to undo the planning error.

<sup>11</sup> SR-MPLSv6, Flex-Algo, etc.

<sup>12</sup> Partitioning node SIDs via geographical location, owner, function, etc.

<sup>13</sup> Collapsing the business and residential networks

<sup>14</sup> This is not a technology specific term

<sup>15</sup> Each LSR would require a reboot

If just 1,000 SIDs are set aside for reuse at the initial stage of planning, all assignments can be accommodated<sup>16</sup>. No configuration changes are required on the existing LSRs; they would continue to have globally unique SIDs; and there would still be ~15,000 unused SIDs for growth. The new, business LSRs in market #1 can be assigned SIDs from the reserved pool. The new LSRs in market #2 will be assigned the *same* SIDs from the reserved pool as market #1. If end-to-end inter-domain LSPs are required between the new devices, a SR-PCE is required if SR is to be leveraged. This is a “tactical” approach for reuse. It provides an additional degree of flexibility without over accommodating for large, unlikely changes.

**Because of this, domain owners in the same autonomous-system can use the same SRGB and scale for unexpected changes. Thus, it is recommended to use the same SRGB in the backbone, data center and residential domains. Like all addressing deployments, the selected SRGB must be partitioned in a way that scales for the domains, markets, applications, projected nodal growth and unexpected tactical reuse. This allows for maximum scale, simplifies planning and eases troubleshooting concerns.**

#### NON-TACTICAL SCENARIO

Rather than explaining tactical reuse first, this section describes the Cox network if it were designed with reuse for all nodes at initial deployment. **This is not recommended and is only used to illustrate the concept of reuse.** When planning index assignments for a single autonomous system,

:

1. Leaking from L1 to L2 should not be performed
2. Indexes on L1 routers must be unique within the area but can be re-used in other L1 areas
3. Indexes on L1/L2 and L2-only routers cannot be re-used

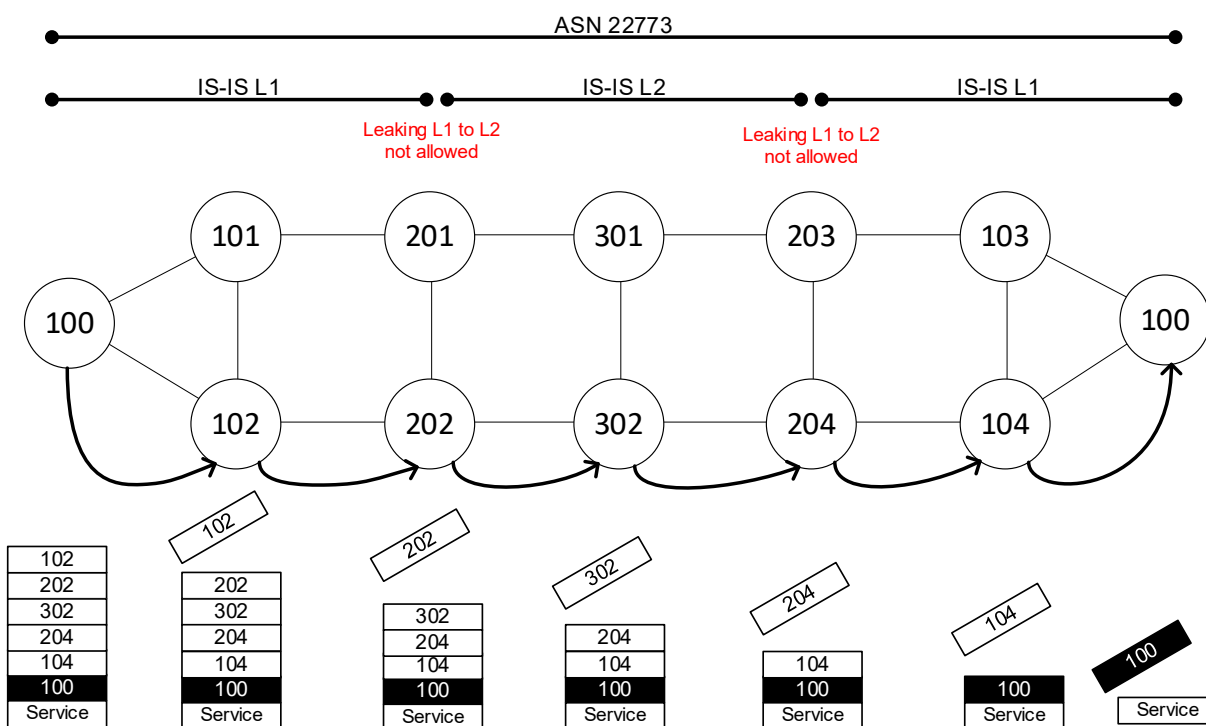
The figure below illustrates index assignment rules and tactical index reuse using ASN22773. This case requires a SR-PCE to facilitate the creation of an end-to-end LSP. The LSP is explicitly routed at each hop in the path. Index assignments are unique on all nodes in the topology except for the edge nodes (100,100) in each L1 area. There are two “core” nodes (101,102 & 103,104) in each L1 area. There are four L1/L2 nodes<sup>17</sup> (201,202,203,204). There are two L2-only nodes<sup>18</sup> (301,302). L1/L2 and L2-only indexes must be unique in this design. Also, this design cannot leak IS-IS prefixes between L1 and L2 areas. The diagram describes the stack, which can be thought of as indexes or MPLS labels, that is provided to the ingress LSR by the SR-PCE. The label of interest is label 100 which is being reused tactically. Although the ingress LSR has been assigned index 100 to itself, it is also able to use a label stack which contains its own label. This will be properly routed to the remote edge node (100).

---

<sup>16</sup> This assumes that route leaking between the markets does not occur and when/if an end-to-end LSP between the markets is required, a SR-PCE is used to build the SID list.

<sup>17</sup> DSRs/HDRs are L1/L2 routers.

<sup>18</sup> BBRs are L2 routers



**Figure 5 - Reuse within an autonomous-system using a SR-PCE and explicit hops**

### 4.2.3. Consistency

As described in the previous sub-section, index assignments can overlap in different domains; exceptions to this were also described. Examples used a consistent SRGB; the SRGB was the same on all nodes in all domains. Vendors recommend this method. If domain owners select different SRGBs upon deployment of SR-MPLS, inconsistency occurs. Although this is technically feasible, it complicates troubleshooting and verification, particularly with SR-TE and a controller. Another reason for inconsistent SRGBs is when a domain grows beyond the total number supported indexes within the SRGB. If this occurs the SRGB must be adjusted on nodes in the domain<sup>19</sup>. **It is recommended that Cox select a consistent SRGB for all domains that scales.**

### 4.2.4. Size

Below is a table summarizing SRGB parameters of platform A, platform B, platform C, platform D. These values are used to determine available SRGBs. Default values are not defined in JunOS and SR-OS, because they do not perform SRGB label preservation like Cisco software. Label preservation was described in an earlier sub-section. From the table below, various SRGBs can be defined. Although the CBR-8 does not support Segment Routing currently, the latest IOS-XE version supports a SRGB size up to 64K indexes. This may change when the CBR-8 platform supports SR-MPLS, but it will be used as the lowest common denominator of all operating systems.

<sup>19</sup> Although it is technically possible to use different SRGBs within a domain/IGP area, this is not recommended. Thus, all node's SRGB would have to be modified.

**Table 3 - SRGB parameters summarization**

Parameter	Platform A	Platform B	Platform C	Platform D	Platform E
Min (default)	none	16000	16000	16000	none
Max (default)	none	23999	23999	23999	none
Min	16	16000	16000	16000	18432
Max	1048575	1048575	1048575	1048575	1048575
Total Limit	1000000	65536	262144	65536	131072

The table below summarizes node counts in all domains. Domain owners provided estimates of future node counts. The values will be used to determine the requirements for the SRGB size based on node counts only. The table does not consider features described in a previous sub-section that require additional indexes from the SRGB.

**Table 4 - Node requirements. Largest IS-IS L1 areas.**

Domain	LSR Count (2018) <sup>20</sup>	LSR Count (2028)	Major Impacts
Backbone	< 110	< 120	None
Residential <sup>21</sup>	< 650	< 1,000	vCCAP buildout Data Center buildout
Data Center <sup>22</sup>	< 40	< 1,000	More virtual routers
Business Core	< 520	< 1,100	Additional S-PEs in each market
Business Access <sup>23</sup>	< 250	< 500	Possible standard change

Based on the estimated node counts, it is apparent that a small SRGB suffices. However, there are always challenges with forecasting node counts due to new products and services. For instance, the access engineering team has not decided on the ratio of virtual CAPs to RPDs. Because the CAP will be virtualized as a container or virtual machine, and automation and orchestration will be used to simplify deployments, it is possible that the ratio could be 1:1 to simplify the automation architecture. If this is the case, the number of SR-MPLS enabled nodes would increase dramatically in the largest market – estimate at around 50,000 in 10 years. Although an option, it is unlikely to be adopted since it would introduce FIB scaling and address allocation problems. Nonetheless, it emphasizes the difficulties of forecasting unknowns. The table below illustrates the forecasting of RPDs in the Phoenix market.

The table below describes a SRGB with 64K indexes. 64K is the recommended SRGB size, and it will be used in the examples to describe partitioning. 20000-83999 is the recommended SRGB range, given the available minimum and maximum limits on the different platforms. Label preservation is lost, but the solution scales. When the SRGB is configured, the operating system adjusts all label blocks. All four network operating systems support 1 million labels. Thus, less than 7% of all total label will be reserved for the SRGB. Vendors have noted that reserving 64K indexes should not cause problems, especially when SRGB size limitations are imposed<sup>24</sup>. However, label usage should always be tested.

<sup>20</sup> All data is collected

<sup>21</sup> Nodes in the markets owned by the data center team are classified as residential.

<sup>22</sup> Duke and Deer Valley

<sup>23</sup> Current standard as of 2018 is less than 250 nodes in an IS-IS L1 area or OSPF non-area zero

<sup>24</sup> JunOS does not limit the SRGB size.

**Table 5 - 64K indexes**

<b>SRGB Size</b>	<b>+ / -</b>	<b>Description</b>
64K	+	Size currently supported on all platforms
	+	Administered by each autonomous-system owner
	+	Features, immediate and future, supported
	+	Node growth supported
	+	Flexible index partitioning – features, reserved, etc.
	+/-	Utilizes <7% of total labels in the LSD
	-	Label preservation

The table below describes a SRGB with 32K indexes. This solution also sacrifices label preservation scale. Less than 4% of all total label will be reserved for the SRGB. This is less of a concern than 64K labels but should still be tested.

**Table 6 - 32K indexes**

<b>SRGB Size</b>	<b>+ / -</b>	<b>Description</b>
32K	+	Size currently supported on all platforms
	+	Administered by each autonomous-system owner
	+	Features, immediate and future, supported
	+	Node growth supported
	+	Flexible index partitioning – features, reserved, etc.
	+	Utilizes <4% of total labels in the LSD
	-	Label preservation

As described earlier, 8K index trades simplicity for scaling. It is ideal for the smallest deployments of SR-MPLS on Cisco software. The initial deployment of SR-MPLS in the residential domain is on less than 20 nodes in the largest market, so the initial deployment scope is relatively small. Thus, this SRGB can be eliminated.

**Table 7 - 8K indexes**

<b>SRGB Size</b>	<b>+ / -</b>	<b>Description</b>
8K (eliminated)	+	Size currently supported on all platforms
	+	Administered by each autonomous-system owner
	+	Features, immediate and future, supported
	-	Node growth supported
	-	Flexible index partitioning – features, reserved, etc.
	+	Utilizes <1% of total labels in the LSD
	+	Label preservation

#### **4.2.5. Derivation**

Index values must be assigned to each SR-MPLS node for them to generate node-SIDs. The method for deriving index values per node must be defined. There were two high-level options for this:

1. Based on static information residing on the node e.g. IP address

## 2. Partitioning indexes and tracking assignments using a database

The first option *appears* to be the most ideal. There is no point in maintaining an additional database if the index value can be derived from some value that is unique to the node. Using the IP address, an index value can be constructed so long as the index value can be large enough to support represented all or at least some of the octets. As stated in an earlier sub-section, lowest common denominator of currently available SRGB sizes is 64K. A future version of IOS-XE, may allow for 256K. With 256,000 index values, the last two octets of an IP address could be represented. This is illustrated in the figures below.

<div>Loopback Example: 172.[Market].[0-255].[0-255]  Loopback Range – Low: 172.29.0.1  Loopback Range – High: 172.29.255.255  Index Values Required: 1 - 255,255</div>	
<div>Example #1  SRGB Base: 16,000 – 271,255  Loopback: 172.29.193.119  Index: 193,119  Label: 209,119</div>	<div>Example #2  SRGB Base: 16,000 – 271,255  Loopback: 172.29.210.24  Index: 210,024  Label: 226,024</div>

**Figure 6 - Deriving indexes via IP addresses with an SRGB size of 256K**

There are many challenges associated with this method. For instance, markets may not always use the same first two octets on all loopbacks. Phoenix uses 10.119.x.x, 10.120.x.x, 10.122.x.x, 172.29.x.x, 172.16.x.x, etc. addresses. Thus, it is possible that two nodes have the same last two octets, and this method begins to breakdown<sup>25</sup>. Another challenge is wasted indexes. No market should ever need 256K indexes, however, that is the minimum number required to derive two octets. Another downside is that there is no room for partitioning. Global indexes are used for various features – tactical reuse, IPv4, IPv6, SRMS, Anycast, Flex-Algo, etc. This method does not allow for partitioning by ownership either.

On the other hand, leveraging a database so that indexes can be partitioned providing structure based on function and group is most ideal. As mentioned above, it allows for clear delineation of ownership. For example, the data center and residential teams can each be allocated their own index blocks. Each team can further partition it, if they choose. For instance, indexes 1-99 can be allocated for nodes at the top tier of the hierarchy, and indexes 200-299 can be allocated for the next tier, etc. <sup>26</sup> As a recommendation, domain owners reserve a sub-block for any future technologies they may be considering that were not

<sup>25</sup> Using this method, no market node could share the same last two octets as a backbone node either, because backbone node indexes must be unique within the AS.

<sup>26</sup> A partitioning structure recommendation is provided in the next section.

addressed in this document. This method of index assignment also uses a smaller SRGB, because the indexes are used efficiently. This reduces the likelihood of label depletion. Database tracking, which is admittedly cumbersome in some ways, is required. This creates additional management overhead. However, it is still the recommended method to assign indexes.

#### 4.2.6. Partitioning

This document has probably blurred the lines of domain ownership. It has, at a high-level, assessed all domains in the Cox network at one point or another, but it has tried to avoid specific statements about the business and data center domains. It has assessed features – whether needs are immediate, future or uncertain – for all domains e.g. Anycast, IPv6, Multicast, Flex-Algo, end-to-end LSPs, etc. It has addressed platform support, label depth, obsolesce, forecasting, etc. These factors were used to determine the SRGB size. Once that has been determined, along with the method to assign indexes, domain owners can plan their blocks. The tables in this section are the recommendations for label partitioning.

The first step is to get holistic view of the MPLS enabled routers in ASN22773 and to begin the initial phase of partitioning the SRGB indexes. The first phase partitions the block by domain ownership. The two tables below describes the total number of MPLS enabled nodes in each domain, as well as the percentage of nodes per domain. A domain is then allocated X indexes which is relative to the percentage of nodes in the domain. An additional column adds growth to the relative percentages. The actual index block assignments are also described. These numbers are based on a SRGB of 64,000. It is recommended that the data center and residential domain owners carves their index blocks by geolocation, mapping server and reuse.

**Table 8 - MPLS Nodes and Index Requirements by Domain**

Devices	MPLS Node Count (Q1 '19)	Percentage of Total MPLS Nodes	# of Indexes Relative to %	# of Indexes Relative to % + Growth
Backbone	106	5%	3,200	5,000
Residential	2000	93%	59,520	54,000
Data Center	42	2%	1,280	5,000

## 5. Interworking

### 5.1. Overview

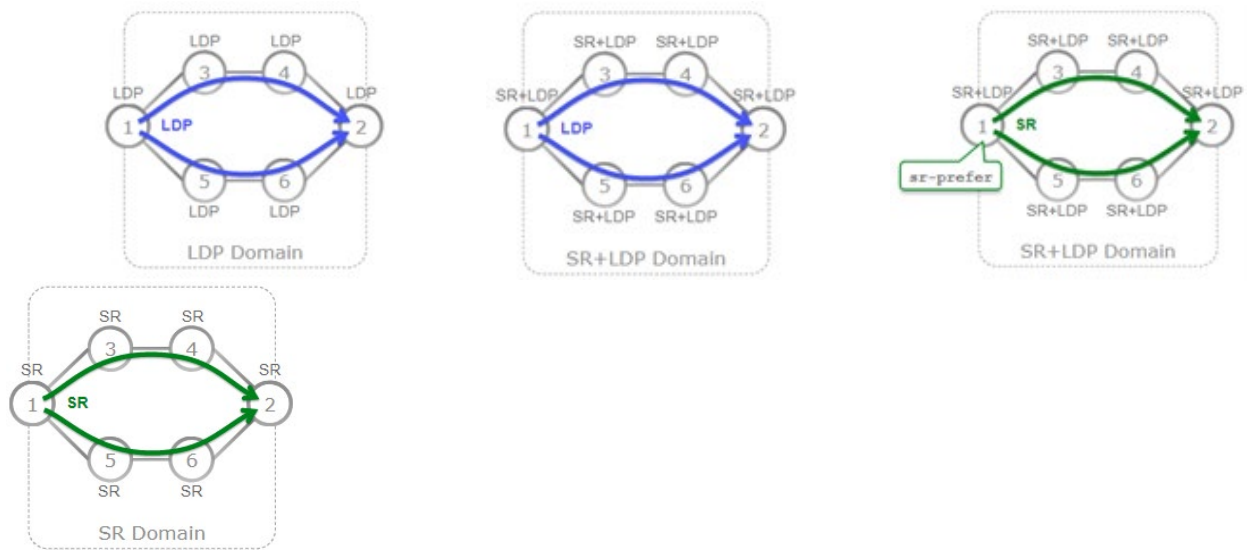
Segment Routing provides two deployment models for integrating SR into a LDP-based network. One of the models, “ships in the night”, describes dual-stacking LDP and SR. This is the simplest implementation but is only leveraged when all nodes support SR. The other model, “interworking”, allows LSPs between LDP-only and SR-only nodes (and vice versa); Cox requires interworking. In this model, functions are required at both control-plane and data-plane layers. The mapping server (SRMS) facilitates control-plane learning. Certain nodes are required for perform SR-to-LDP and LDP-to-SR label swaps in the data-plane. The following sub-sections describe interworking as well as the challenges within Cox.

“Ships in the night”, or dual-stacking, is the simpler deployment model and will be briefly described here for completeness. There are four stages when transitioning.

1. LDP only
2. Dual stack
3. Prefer SR

#### 4. SR only

The starting state has only LDP configured on all nodes. The ending state has only SR configured on all nodes. The second state dual stacks SR on LDP. Thus, the control-plane for SR is synchronized via the IGP. However, the FIB, or data-plane, is not programmed to push, pop or swap SR labels<sup>27</sup> until the third stage. In the third stage, SR is configured as the preferred protocol. On Juniper platforms, this is done by manipulating the route-preference. On Cisco platforms, “sr-prefer” is a macro which makes SR preferred over LDP. When a change is made on a node and because LSPs are unidirectional, only the LSP(s) for which that node is the ingress LSR will be impacted. This must be done on a node-by-node basis, until all nodes have transitioned to prefer SR. In the final stage, LDP is removed and the domain becomes SR only. The stages are illustrated below.



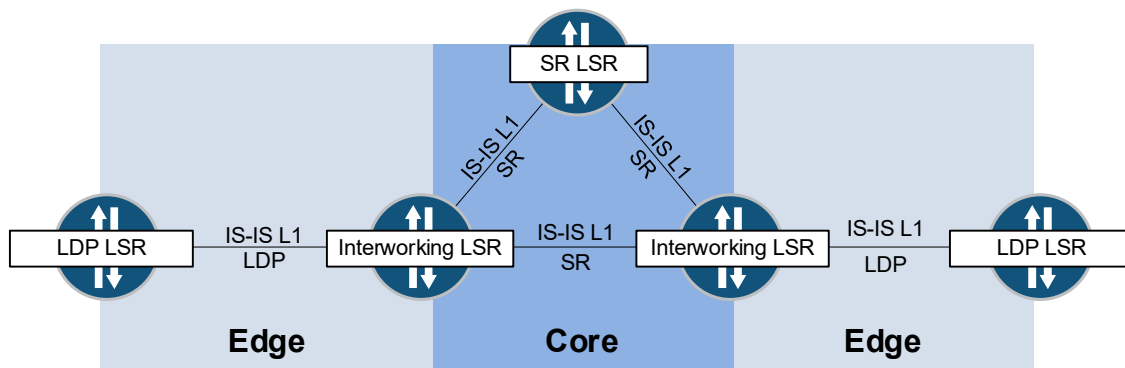
**Figure 7 - Ships in the night**

Interworking expands on ships in the night. The difference between ships in the night and interworking is that interworking accounts for LDP remaining in the network. Terminology has been created to define new nodal functionalities introduced in this deployment model. In the end state, the SR deployment can be classified into two sections: “the core” and “edges”<sup>28</sup>. Unique nodal functions occur in both sections. “SR LSRs” perform label swaps using SR labels. “Interworking LSRs” swap SR-to-LDP and LDP-to-SR labels. SR LSRs run SR only, and an interworking LSRs run both SR and LDP. However, the interworking LSR does not have dual stacked interfaces. The “edge” is where “LDP LSRs” resides. These nodes only run LDP. This is illustrated in the figure below.

<sup>27</sup> These are traditional MPLS terms, and they have been redefined in the SR framework.

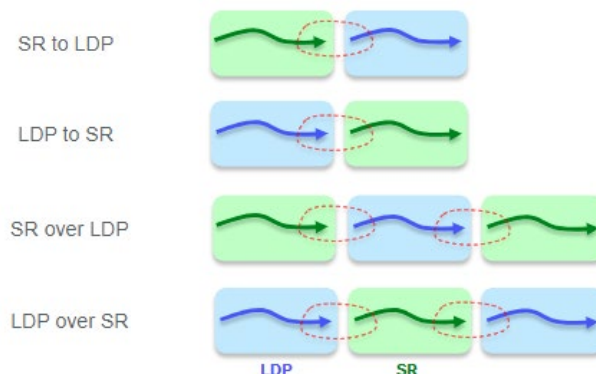
<sup>28</sup> This is not an industry or technology term. It has not been defined and is made up here.





**Figure 8 - Interworking terminology in the SR IGP area**

It's important to keep in mind traffic flows and LSP termination points i.e. the ingress and egress LSRs of an LSP. Traffic flows and LSPs determine which node types the LSP must traverse. An LSP could terminate on any of the nodal types defined in the interworking model. For instance, a service is defined on two LDP LSRs. Thus, an LSP must exist. If the LDP LSRs are in different edges, the LSP must be programming in at the LDP LSRs and Interworking LSRs FIBs, using the diagram above. This requires swapping LDP-to-SR and SR-to-LDP at the interworking LSRs. However, if the service is defined on two LDP LSRs in the *same* edge, LDP only swaps occur. These scenarios are illustrated in the figure below.



**Figure 9 - Label swapping**

To understand the type of label operation being performed, one must identify the LSP's termination points. Knowing the termination points also allows determining whether a mapping server is required for the LSP. A matrix has been created to reflect the residential domain and identify when the mapping server is required.

## 5.2. Mapping Server

A mapping server is configured with individual or blocks of IP addresses and indexes to be assigned to LDP LSRs. The mapping server will identify LDP LSRs via its IGP database. It then allocates and advertises, via its own LSP, an index on the nodes' behalf using a new TLV. From the perspective of a SR LSR, an index assignment exists for every node, including LDP LSRs. Thus, all nodes appear SR enabled.

This allows SR LSRs to push SR labels for services residing LDP LSRs. Once the mapping server allocates indexes on behalf of LDP nodes, the interworking LSR can program the FIB for label swapping between protocols. . This process is out of the scope of this document and should be understood via RFCs, technical documentation, training and testing.

At least two mapping servers should be defined within an IS-IS area. It functions like a BGP route-reflector; it is not required to be in the data-plane and is purely a control-plane function. The mapping server can reside on any node within the IGP area. Keep in mind that it will serve all nodes in the IGP area; this includes nodes owned by the data center team. Thus, the location of the mapping server should account for index block planning, because it may serve data center nodes<sup>29</sup>. The selection of mapping server placement is highly dependent upon operational perspectives no recommendation will be made on mapping server placement.

The mapping server is responsible for advertising prefix and index associations of LDP LSRs. It is recommended to dedicate a range for mapping server allocations. A recommended SRGB plan was defined in the Global Block Planning section. Currently, there are two methods to perform label allocation:

1. **1-to-1** using individual prefixes and individual indexes
2. **Many-to-many** using prefix and indexes ranges

The mapping server allocation method impacts operations and provisioning teams. Typically, new LSR deployments in the residential domain are irregular, compared to the business domain. In the business domain, or a domain where new LSR deployments are more frequent, this brings more value than in the residential domain. In today's state, simplifying provisioning does not outweigh the increase in operational complexity. Also, most mapping server entry challenges will be on "day one" deployments – not "day two" provisioning challenges. It is recommended that 1-to-1 be deployed to reduce the complexity of the interworking deployment. Each time a LDP LSR is added to the network, the mapping server must be updated. This assumes that the 1-to-1 entry type is being used. Operations and provisioning teams should be aware of this new, additional step in the deployment process.

With the mapping server, an error can occur when two mapping servers advertise different index bindings. This is likely an operational issue, rather than by design. In this case, a preference selection can occur, so that all nodes within the IGP area have selected the same mapping server for bindings. However, this was not the case in the initial phase of testing in the lab. In fact, this function is being defined within the IETF and is still in the "draft" state<sup>30</sup>. However, vendors often implement features before the IETF standardizes it. The current mapping entry preference is summarized below. However, this is subject to change.

## 6. Implementation

### 6.1. Prerequisites

After the base configuration is applied to all Vendor A and Vendor B platforms, it is recommended a tool be used to aide in verification as well as to generate and validate mapping server entries. This tool should be provided to the field along with the LLD. It should leverage NETCONF transport and YANG data models to collect IS-IS and LDP database information. Collection should occur via two nodes for

---

<sup>29</sup> If the data center receives their own index block from the residential SRGB, that team may not want the allocations performed on the NCS.

<sup>30</sup> <https://datatracker.ietf.org/doc/draft-ietf-spring-conflict-resolution/>


redundancy. By cross referencing both IS-IS and LDP data collection, a node can quickly be classified as one of the following:

1. **SR LSR** – index in IS-IS database
2. **LDP LSR** – entry in LDP database
3. **Interworking LSR** – index and entry in both databases
4. **IP-only** – no index or entry in either databases

All nodes generate and advertise an IS-IS LSP. These contain hostnames, prefixes and metrics. Nodes enabled with SR also advertise their index and SRGB assignments. Mapping servers also advertise the entries used for interworking. All information should be made available to an operator via a GUI. This allows them to quickly parse the SR information of any node which is not an easy task otherwise. This also allows an operator to easily identify LDP LSRs and generate mapping server configurations. Using this same information, two automation tasks can be performed. These tasks require cross examining LDP LSRs against the mapping server's advertisement. By cross examining these, it can be determined which nodes require an entry. Thus, mapping server configurations can automatically be generated. This is ideal for “day one” deployments where markets have 400+ LDP LSRs. The tool can also routinely validate the mapping server entries and generate reports when there is a LDP LSR without a mapping server entry. This is ideal for “day two”. If a tool is constructed with these features, it will greatly reduce the workload of operators.

## 6.2. Current

This section addresses the high-level implementation phases. It also defines specific challenges and concerns. The current network state is illustrated below. All nodes are LDP LSRs. Only MPLS-enabled nodes are shown; IP nodes are not impacted by SR-MPLS. Topologies and hostnames vary market to market.



**Figure 10 - Sample topology - current state**

### 6.3. Dual Stack

After upgrading both MX and NCS platforms to the preferred software versions for SR-MPLS, the base SR configuration can be applied. This should be non-service impacting<sup>31</sup>. Like administrative distance (Cisco) and route preference (Juniper), MPLS transport protocol preferences exist. This allows an operator to toggle between different MPLS transport protocols, if label switched paths exist for both. By default, LDP is preferred over SR-MPLS in both platforms<sup>32</sup>.

SR-MPLS and LDP are dual stacked with LDP remaining the preferred MPLS transport protocol. SR LSRs will program an additional LFIB entry for other SR LSRs. Thus, two entries for an SR LSR will exist; an entry to reach the SR LSR via LDP and another entry via SR. In addition to this, link protection also requires additional hardware resources to program backup routes for FRR capabilities. It is recommended to test the change in FIB/LFIB utilization caused by dual stacking and TI-LFA.




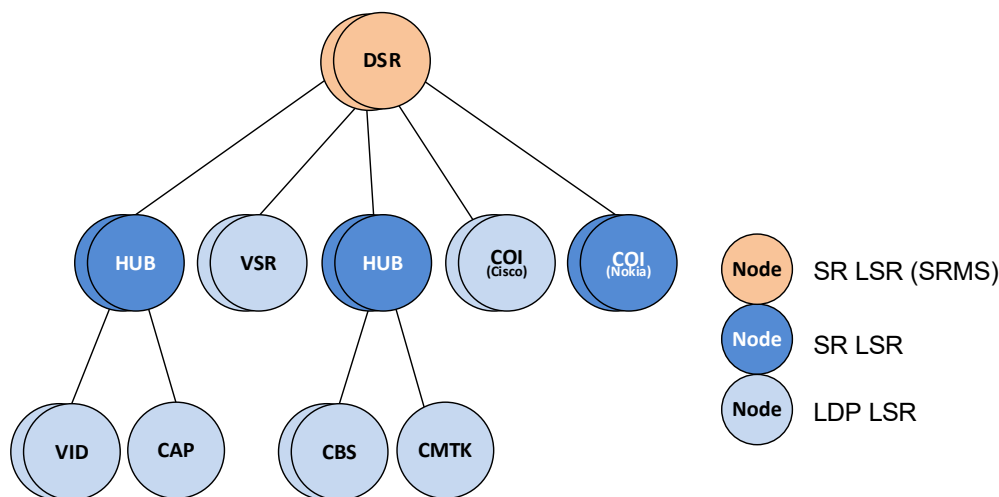
Figure 11 - Sample topology - Dual Stack

### 6.4. Mapping Server

This task, and the following, will leverage the tool described in the earlier sub-section. This should be non-service impacting. Since the mapping servers' location is currently undecided, the illustration uses the DSRs for explanatory purposes. As noted earlier, the mapping server is control-plane only. Thus, the functionality is the same no matter where implemented in the IGP area. After implementing the mapping server, every SR LSR should have an index for all MPLS nodes. SR LSRs will program an additional LFIB entry for LDP LSRs. Thus, the LFIB utilization on the NCS and MX will increase. The workload in this phase is greatly reduced if the SR tool mentioned in the previous section is available. It can be used for an immediate validation of mapping server entries and determining if entries are missing.

<sup>31</sup> Only nodes dual stacking will be impacted – MX/NCS. This most likely reason for service impacts, if any would be scale. Dual stacking on the MX and NCS should be scale tested prior to deployment

<sup>32</sup> This should be revalidated with the selected SR-MPLS codes.



**Figure 12 - Topology - mapping server**

## 6.5. Protocol Preference

When an operator changes the label preference on the node, the LSR will reprogram its FIB to “push” SR labels instead of pushing LDP labels. Thus, LSPs for which the LSR is the headend are impacted. When a LSR terminates a service and has the preference changed, a service interruption can occur. This is scenario is applicable to the MX. This may apply to the Nokia COI router, depending on its deployment. Since LSPs are unidirectional, the remote LSR’s LSP is not impacted by the change. The operation at the headend LSR is summarized in the following points:

- Push LDP, by default
- Push SR, if configured<sup>33</sup>

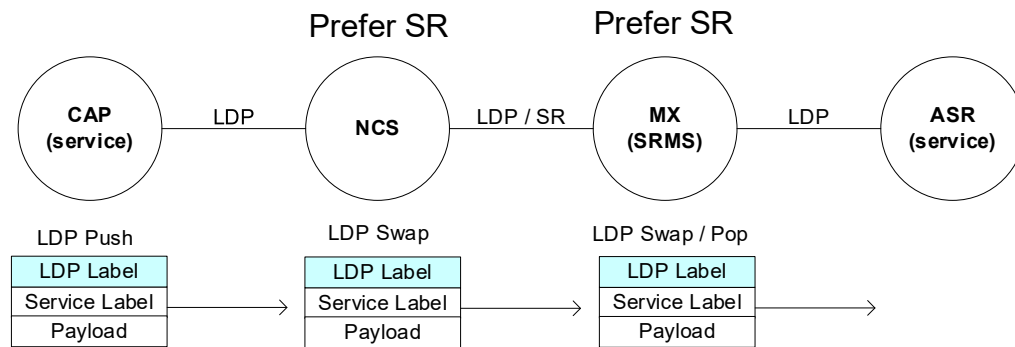
When the preference is modified on a the NCS, a “P” LSR, transit LSPs should not be impacted. Changing the preference on a P LSR does not cause it to perform the interworking functionality. In other words, it will not change its swap operation from a LDP-to-LDP swap to a LDP-to-SR swap. It is recommended to validate that the preference does not delegate interworking for transit LSPs using the software versions selected for deployment. This being the case, modifying the preference is less likely to impact services than the next stage. The operation at the transit LSR is summarized in the following points:

- Preference does not influence the swap operation
- Swap to LDP, when a SR LSP exists but leveraging a mapping server<sup>34</sup>
- Swap to SR, when a SR LSP exists without leveraging a mapping server

An illustration of how the protocol preference at the transit LSR and SRMS does not impact the LSP. The MX and NCS are dual stacked with LDP and SR. The DSR is performing the mapping server function. Thus, both the MX and NCS having SR labels for both CAP and COI. A service terminates on the CAP and COI. The preference is modified on the MX and NCS. However, if the transport layer does not change, the service is not impacted. The NCS continues to swap LDP labels because of the rules stated above.

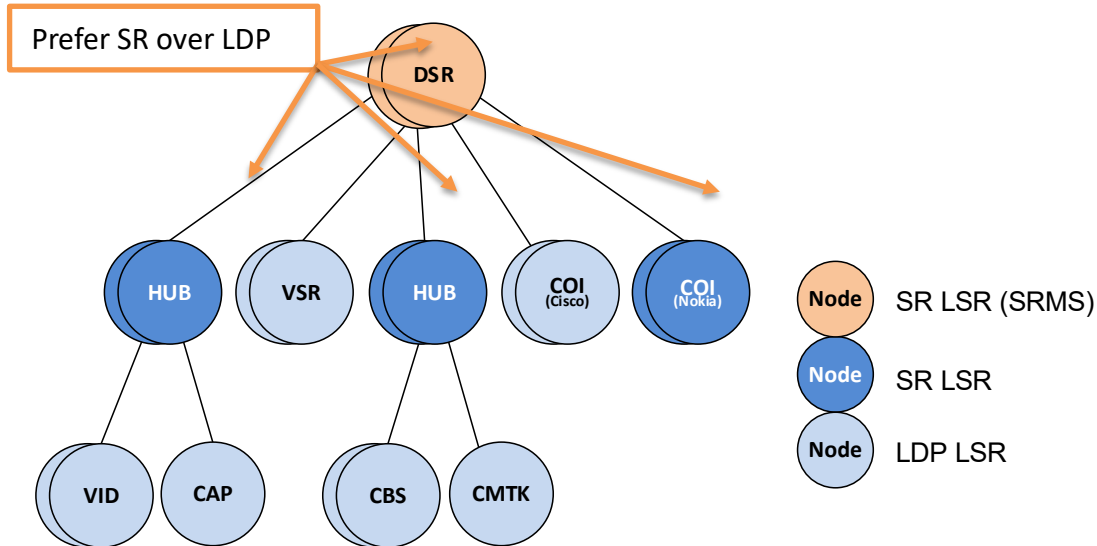
<sup>33</sup> Whether the SR LSP is derived from mapping server or is end-to-end is irrelevant.

<sup>34</sup> The MPLS protocol preference does not impact this decision, per code testing.



**Figure 13 – SR preference should not impact transit LSPs**

The example topology that has been referenced in the previous sub-sections is supplied below. The MX and NCS are configured to prefer SR over LDP.



**Figure 14 – Same topology – Protocol Preference**

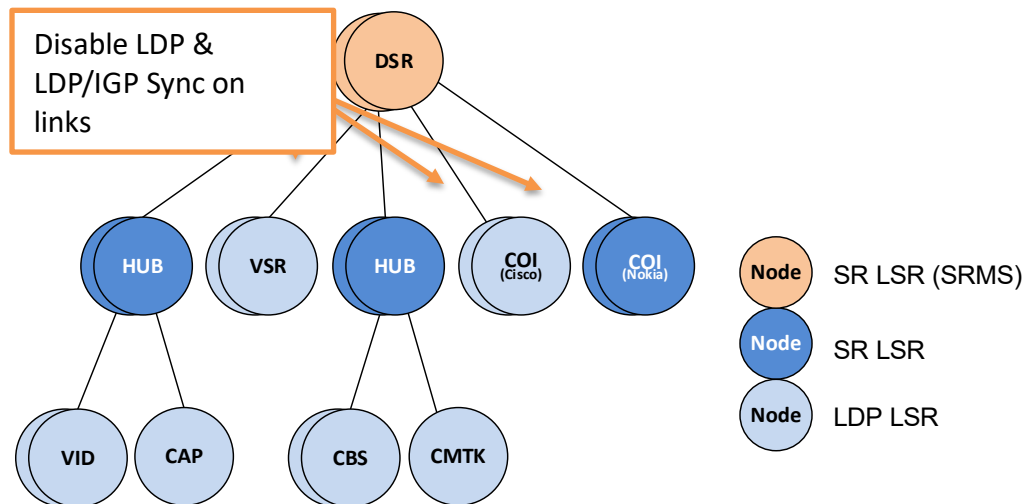
## 6.6. Interworking

At this point, all MX and NCS nodes should prefer SR as the MPLS transport protocol. Every MPLS node requires an index. SR LSRs have the index defined in their configuration. LDP LSRs are assigned indexes via the mapping servers. The final stage of implementation is disabling LDP and relying on interworking. In other words, LDP-to-SR and SR-to-LDP swaps.

When a site is selected to transition to Segment Routing, LDP should be removed on the NCS uplinks and interconnects. LDP-IGP synchronization should also be removed from those links or the IGP will advertise them with “max-metrics” as designed. LDP should remain on links to CBR-8, uBR, ASR, etc. platforms. Once the configurations have been committed, the NCS will perform the interworking function and forward the label switched traffic using the new label operation.

It is less likely that Internet and video services being impacted, because they are, or can be, transport via native IP. Thus, L2VPN and L3VPN customers, whether internal and external, should be aware of the potential for service impact. The most likely reason for impact to services is incorrect or incomplete

mapping server configurations. After decommissioning LDP on the defined links at all hub-sites, the market will have a “SR core” and this project is complete.



**Figure 15 - Topology - Interworking**

## Abbreviations

Architectures	
SR-MPLS	A deployment of SR where MPLS is the data-plane and IPv4 is the control-plane
SR-MPLSv6	A deployment of SR where MPLS is the data-plane and IPv6 is the control-plane
SRv6	A deployment of SR where IPv6 is the data-plane and IPv6 is the control-plane
Unified SR	An architecture using BGP-LU to unify multiple SR domains
Unified MPLS	An architecture using BGP-LU to unify multiple MPLS domains
Distributed TE / SR-TE	A TE deployment not utilizing a SR-PCE, thus having policies defined on the node
Centralized TE / SR-TE	A TE deployment utilizing a PCE, thus having policies defined on the PCE
Terminology	
Backbone Domain	ASN22773: L1/L2 or L2-only nodes
Data Center Domain	ASN22773: L1 nodes – Duke and Deer Valley
Residential Domain	ASN22773: L1 nodes – Phoenix, Las Vegas, Hampton Roads, Rhode Island, etc.
Business Domain	ASN15218: All nodes
Business Core	ASN15218: L1/L2 or L2-only nodes

Business Access	ASN15218: L1 nodes
Segment Routing (SR)	
Source Packet Routing in Networking (SPRING)	Juniper's name for Segment Routing
Segment ID (SID)	A MPLS label identifying a node, IGP adjacency, prefix, LSP, etc.
Segment List	The MPLS label stack derived via Segment Routing
Global Block (SRGB)	A label range for global SID assignments. E.G. Node-SID
Lobal Block (SRLB)	A label range for local SID assignments. E.G. Adjacency-SID
Mapping Server (SRMS)	Devices which assign SR labels for non-SR devices via IGP advertisements
Mapping Client (SRMC)	Devices which receive SR labels from the SRMS
Path Computation Element (PCE)	A controller for RSVP-TE, SR or SR's ODN
Path Computation Client (PCC)	A client of the controller
MPLS-TE	MPLS protocols which provide traffic engineering capabilities such as SR and RSVP-TE
TE / SR-TE	Non-default routing – may be defined on either a PCE or directly on a node via CLI
On-Demand Next-Hop (ODN)	A model used to compute end-to-end LSPs across multiple domains via a PCE



# **Creating The Intelligent Edge**

## **Increasing DAA Velocity Using Service Orchestration**

A Technical Paper prepared for SCTE•ISBE by

**Chris Busch**

Fellow – CTO Network Solutions  
CommScope  
2400 Ogden Avenue, Lisle Illinois  
630-281-3150  
chris.busch@commscope.com

**Dave Baran**

Fellow – CTO Network Solutions  
CommScope  
101 Tournament Dr. Horsham Pennsylvania  
215-323-1973  
dave.baran@commscope.com

**Jeff Finkelstein**

Executive Director Network Architecture  
COX Communications  
6205 Peachtree Dunwoody Rd Atlanta  
jeff.finkelstein@cox.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	4
1. Goals of The Intelligent Edge .....	4
2. Orchestrating the Intelligent Edge.....	5
2.1. Creating Service Groups.....	7
3. DAA Inventory.....	8
3.1. Static-Inventory .....	8
3.2. Dynamic-Inventory.....	9
3.3. Using a Mobile App .....	10
4. DAA Inventory Provisioning.....	11
5. CIN Toplogy.....	16
6. Dynamic Edge for DAA .....	17
6.1. LLDP Stage.....	17
6.2. 802.1x Stage.....	18
6.3. DHCP Stage.....	19
Conclusion .....	20
Abbreviations.....	20

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Product Catalog Creation.....	5
Figure 2 - A Single Service Group.....	6
Figure 3 - Service Group to Named Service.....	7
Figure 4 - Service Creation with Active Resources .....	8
Figure 5 - Warehouse QR Scan.....	10
Figure 6 - Mobile App QR Scan Service Group to Inventory During Field Installation.....	11
Figure 7 - Initial Switch Config .....	12
Figure 8 - OSS Consumption of LLDP .....	13
Figure 9 - Intelligent Edge with Dynamic Inventory .....	14
Figure 10 - CIN Port Dynamic Profile Provisioning .....	15
Figure 11 - Any DAA Service Any Port.....	16
Figure 12 - EVPN VxLAN CIN.....	16
Figure 13 - Intelligent Edge Discovery and s-Leaf Provisioning Sequence.....	18
Figure 14 - AAA Sequence .....	18
Figure 15 - DHCP Sequence .....	19

# Introduction

Distributed Access Architectures (DAA) open many new business opportunities at the edge of the network for the cable operator. Increasing service options in bandwidth terms, overall service quality, reduced cost of optics, and lower latency, are key values of the architecture. The infrastructure introduces all digital wavelength division multiplexed 10-100 Gigabit Ethernet from Hub to pole with highly precise timing support. This further creates a simpler path to 5G roll out for the MSO, in addition to residential and business services over any form of wireline access.

While Remote PHY is the most common distributed access technology today, there are multiple forms of DAA including Remote PHY, Remote MACPHY, and OLT-PON, in addition to emerging 5G backhaul networks.

Serving each of these platforms is an Ethernet system known as the Converged Interconnect Network (CIN). This digital infrastructure enables multiple services over each of these remote functions. These same CIN systems enable 5G backhaul opportunities that require access network sharing, sometimes termed network slicing functionality. Lastly, DAA brings with it support for the era of cable Network Function Virtualization (NFV). In fact, DAA is not any one element, it is all these platforms working together.

However, one fundamental role that the CCAP platform had must be introduced into the Distributed Access Architecture. The automatic discovery, provisioning, and telemetry for Physical layer and MAC layer functions in the topology must exist for DAA as they do today within any Integrated-CCAP chassis. Without this functionality in place, the increased management complexity due to the larger number of DAA devices can negatively impact operations and expected cost savings.

The authors seek to address the opportunity to add intelligence at the DAA edge using a select set of back office systems creating an SDN-based ‘Intelligent Edge’ for the DAA. Through this Intelligent Edge, operators may now see these distributed and often decoupled systems realized not at a network or element level, instead being brought together to deliver services in an end-to-end fashion, with dynamic provisioning effectively returning the same Zero-Touch onboarding and deployment a CCAP-based chassis provides for HFC-based PHY and MAC functions.

In this paper, we will review options for top down orchestration based on industry practice. Additionally, an approach to bottom up network discovery and considerations for automation and device provisioning from the edge of the network into the Hub and Headend are detailed to provide guidance on Zero-Touch provisioning for DAA.

# Content

## 1. Goals of The Intelligent Edge

For the cable operator, introducing Distributed Access Architecture involves many stakeholders across the business. Engineering needs to understand the network bandwidth and interconnect needs for each DAA remote device type. Operations needs to understand how performance and troubleshooting data will be made available to preserve service assurance systems. Warehouse and Inside Plant Engineering teams have needs where logistics of inventory and assignment to outside and inside physical plants are concerned. Field technicians are concerned with time spent during maintenance windows and ability to close work orders for each installation. Service assurance as well has additional resources available for monitoring and managing the performance of the holistic HFC network.

Goals for most operators deploying DAA include:

- Simple inventory association
- Automated provisioning for any remote device
- Minimize remote device installation time
- Service assurance of remote device operations

From a high level, many if not all of these goals and stakeholders can be aided through an effective onboarding process for DAA remote systems.

Onboarding is the function of network element discovery and provisioning based on role and required service attributes.

To enable onboarding, multiple systems in the current back office, specifically the Operational Support System (OSS) evolve to support effective DAA deployment.

This evolution influences and creates the Intelligent Edge through:

- Inventory Process
- Device and Topology Discovery
- Service Activation
- Device Authentication
- Dynamic Device Addressing
- Device Provisioning
- Network Provisioning
- Field Installation Application

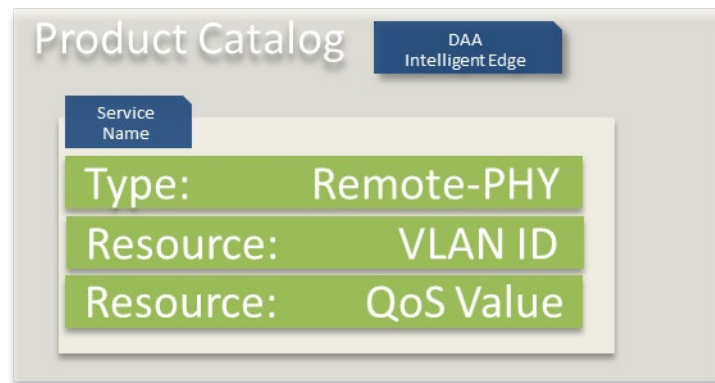
Perhaps more accurately, the Intelligent Edge is initially created through integrations among each of these systems to enable the goals operators seek in deploying DAA technologies as a solution.

## 2. Orchestrating the Intelligent Edge

The Intelligent Edge has knowledge of the resources in the DAA network and their relationships. Resources can be thought of generally as the various actors in the network involved with any DAA service and the values these actors must have configured for each DAA service and function to operate.

We can build up an effective relationship design for the Intelligent Edge by thinking of any DAA as a service that consumes available inventory in part or in whole. Services are defined at a top level for the operator and become associated to operating regions and entities across the network.

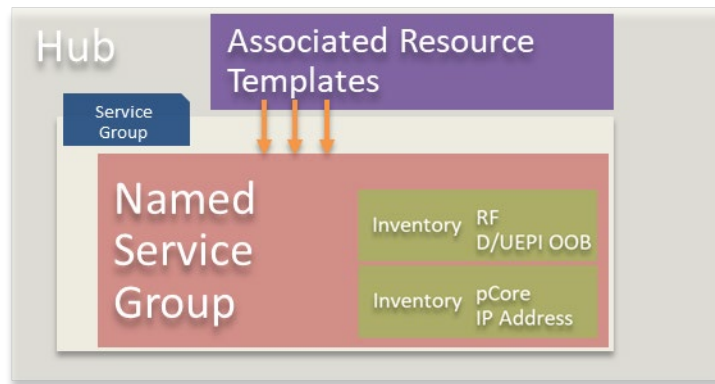
The approach presented here is a product catalog or so called ‘model driven’ approach to service creation whereby association to common elements minimizes manual entry and promotes high reusability of configuration data such that the least amount of configuration is required when a device is attached to the system. The idea of the catalog is to define standard products that are available for sale with standardized attributes that ensure that the provisioning/service delivery phase has all of the necessary information collected to ensure proper implementation.



**Figure 1 - Product Catalog Creation**

Following that service model from a catalog approach, a Remote-PHY service type can be defined in the catalog to cascade down global configuration to members of the service. In this example, assume that the Remote PHY service includes a VLAN identifier and a QoS value for traffic separation on the network link. From this stage onward, other associations with increasing granularity may be defined and associated in a chain like manner.

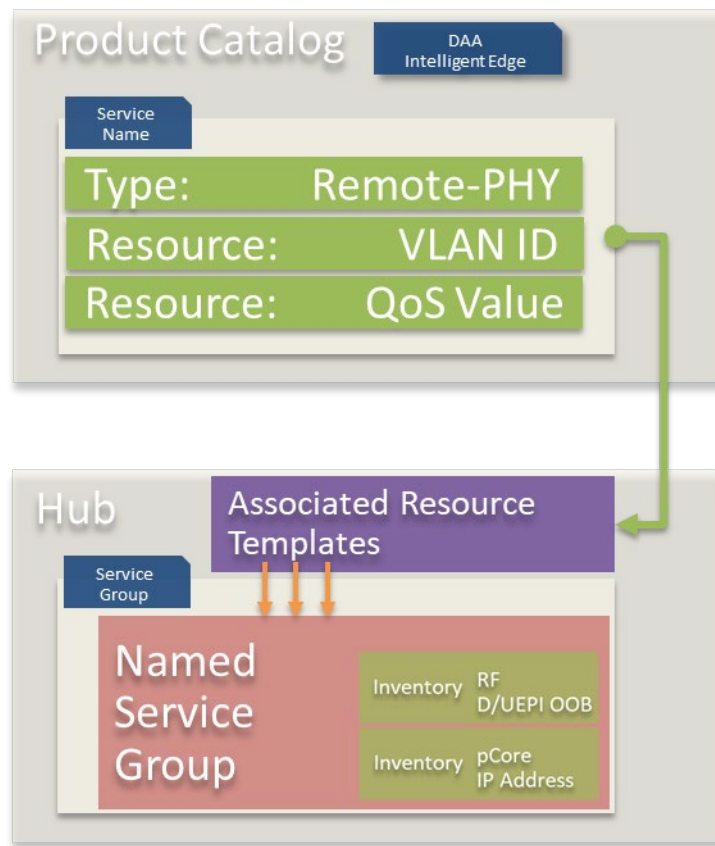
For example, ‘RPD\_Service\_A’ is added to an entity such as ‘North-East’. Within North-East, there may be multiple Headends. ‘Headend\_A’ includes resources such as PNF (Physical Network Function) and VNF (Virtual Network Function) DOCSIS cores for Remote-PHY association. Headend\_A also includes defined RF resources and Headend\_A may have a Hub level resource below it that inherits from the entire chain described. The RF resources are a many to one asset, meaning they are defined as a named template and associated to an asset in the Headend or Hub, specifically the named service group RPD\_Service\_A which becomes a member of the Headend or Hub inheriting access to resources at that level.



**Figure 2 - A Single Service Group**

When the product catalog for the RPD Service is associated to the Named Service Group, additional inheritance of functionality occurs. The product catalog is less specific in terms of Headend or Hub resources, such as a RF forward and return or video out-of-band configuration values. Certainly, if the operator had the same DOCSIS RF resource plan and Video QAM RF resource plan across many Hubs then these could be configured up in the catalog. The relationship is purely one of re-use and ownership to facilitate orchestrating a service.

In the example demonstrated, all Remote-PHY service groups have a Remote-PHY VLAN ID value and QoS value the example here chooses to apply to all Hubs that have joined the service. By having this as part of the product catalog, the necessary information is obtained up front to ensure proper configuration of both the Remote-PHY device and the Secure Leaf (s-Leaf).




**Figure 3 - Service Group to Named Service**

## 2.1. Creating Service Groups

To create the service-based view of DAA systems, we have borrowed the concept of a 'Service Group' and use this as a simple entity to hold members of configuration objects both provisioned and learned. This 'Named Service Group', we could say 'Service Group 405', will have associated Resource Templates and be a member of various inventory objects in the network. These are the provisioned objects.

Resource Templates are values for configuration such as Video Out-of-Band, Downstream and Upstream channels, VLAN value, QoS value that service 'Service Group 405' may inherit. As inventory is associated, such as a Remote PHY DOCSIS device, the resource templates are applied by network function and executed in the service specific manner updating the device record held in our DAA Inventory function as indicated by the 'Service Group 405' record.



**Figure 4 - Service Creation with Active Resources**

Secure Leaf switch is also known as an s-Leaf exists at the edge of the network to accept a variety of distributed access device connections. A Secure Leaf is so termed given its role in supporting IEEE 802.1AE MACsec, and encryption standard for Ethernet physical layer.

When a physical RPD is attached to the network and joined to the service, its network address and s-Leaf attachment information becomes dynamically known to the Named Service Group, for example Service Group 405.

Service Group 405 may have simply held the RPD MAC address and the membership to the Hub. The Hub may have the pCore (Principal CCAP Core) for Remote PHY and therefore Service Group 405 inherits this pCore IP address.

As the system evolves, we will see how more of the service group may be constructed dynamically rather than through a buildup of static configurations pushed to any single device. Instead, we will have the ability to construct and deconstruct the needed provisioning updates at run time effectively when a device is discovered at the DAA edge using SDN control functions to execute these tasks on behalf of the service. Changing the provisioning model from "before use" to "at use" simplifies service creation and allows it to be executed without human involvement which is critical to success with a DAA architecture.

### 3. DAA Inventory

#### 3.1. Static-Inventory

A Static-Inventory based system assigns the unique identification of a remote device to a service in the network. This may occur during a work order ticket for the field team as the truck is loaded for the day or may occur from other systems. When the activation process is driven from a Static-Inventory view, the field technician must use the specific remote device for that service order.



For Remote-PHY, this may mean that a pre-staging has occurred where the warehouse team have bench tested the RPD, or other DAA remote device type, prior to deployment and assigned it to its designed ‘Service Group’ where all that remains is for the physical installation in the field.

The steps involved in a Remote-PHY Service Group from Static-Inventory include:

- Association of RPD MAC address from Inventory to Service Group
  - Service Group inherits RF configuration and Core attributes
- Association of RPD MAC address to a Core
  - At minimum a Principal Core for DOCSIS services

Most DAA remote devices use some form of Ethernet backhaul such as 10 GB, 25 GB, or 40 GB Ethernet. This backhaul or access aggregation in the outside plant typically occurs over Wave Division Multiplexing (WDM) either Dense or Coarse (DWDM/CWDM). Others may use xPON or other technology (e.g. 5G, mmWave, etc.), but for the purpose of discussion we will focus on point-to-point Ethernet.

The two ends of this backhaul are pluggable optics, where the Hub or Headend side has an s-Leaf switch populated with Small Form Pluggable+ (SFP+) or a variant based on the PHY rate that is interconnected to a fiber mux tray at the top of the rack. At the other end in the outside plant (OSP), a Fiber Optic Splice Closure (FOSC) will present one or two fiber pair(s) to the remote device. The remote uses a pluggable optic to terminate the interconnect. The pluggable optics involved in this system may be fixed wavelength or dynamically tunable.

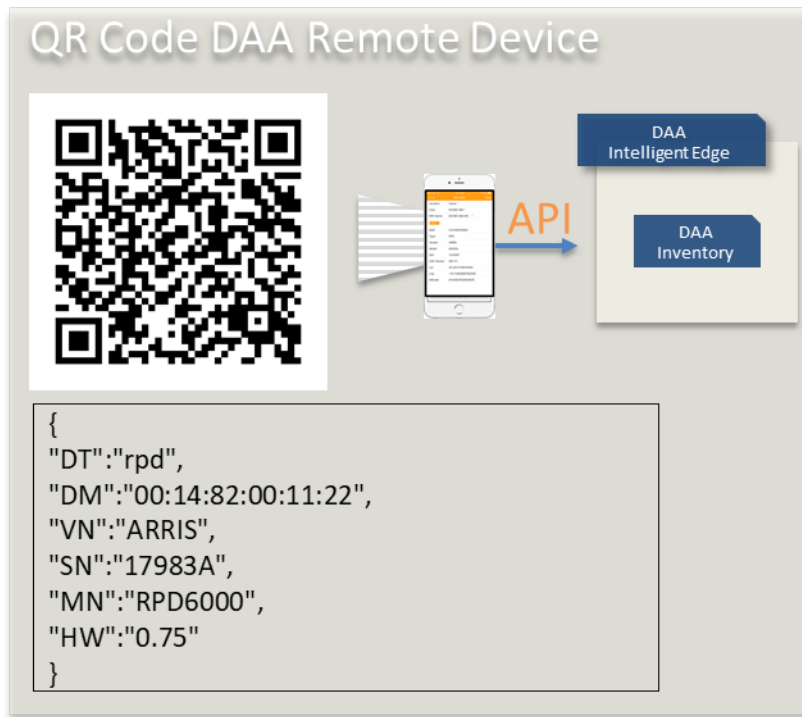
In the event the optics are on fixed wavelengths, the Static-Inventory system must include assignment of the SFP+ assets at both ends of the fiber path from s-Leaf to remote port(s).

The static model has a significant drawback as it produces distinct devices rather than fungible ones. At each stop along the installer's path, they have to get the right device and hope that it was provisioned correctly. In the event that a device doesn't work, the technician can't easily diagnose the problem by swapping another unit from the truck as that unit would be provisioned differently. The key to simplicity is to keep things interchangeable, which requires a dynamic model.

### **3.2. Dynamic-Inventory**

A Dynamic-Inventory based system assigns the unique identification of a remote device to a service in the network at the moment the technician installs the device in the OSP.

This is made possible using several systems and application integrations discussed in further detail.

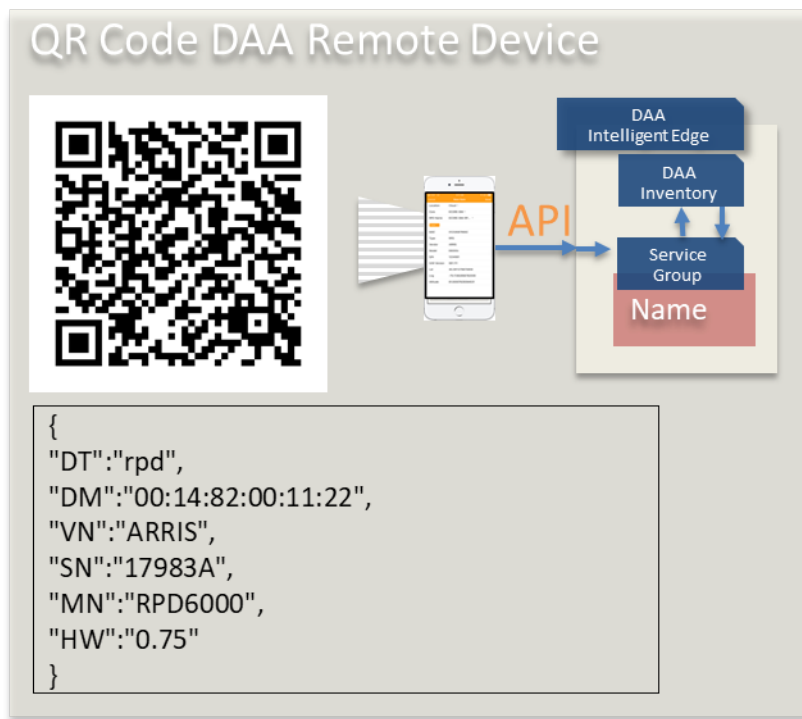


**Figure 5 - Warehouse QR Scan**

An effective system will integrate the warehouse view of physical inventory with the network view of actual in-service inventory. There can be subtle differences in how the operator may choose to assign inventory from warehouse to network. Specifically, whether a DAA remote device is assigned to a specific work order or destination service group, or whether the field technician will generate that assignment during installation. While the distinction is subtle, the implications are distinct to the back-office system.

### **3.3. Using a Mobile App**

Onboarding for DAA Remote devices can be assisted from the field using an install app from a mobile device. Where Remote PHY devices are concerned, there is a QR code available that is associated to the DAA device module. The field technician at time of installation may scan this code with the mobile app to facilitate the onboarding process.



**Figure 6 - Mobile App QR Scan Service Group to Inventory During Field Installation**

The mobile app will also send to Inventory the geo-coordinates of the scanned device. These are sent as latitude and longitude, stored with the device record in the Inventory of the Intelligent Edge.

The field technician is prompted to select the Service Group to which this device is to be associated with. The Intelligent Edge will have returned a list to choose from or optionally the service group may be searched for within the mobile application.

Once submitted, an API is updated with the same inventory information from the warehouse, including geo-coordinates and now a named service group association to the DAA Inventory device record.

This is a key stage in the onboarding process. Association of a physical device to a ‘Service’ will now have impact on the edge topology discovery and the DHCP provisioning discovery stages.


## 4. DAA Inventory Provisioning

At the heart of the Intelligent Edge is the idea of an ‘Active Inventory’. This ‘Active’ state is simply a closed loop of updates from SDN control functions to the Inventory database of the Intelligent Edge.

If we consider the Integrated-CCAP chassis model of line cards for PHY, MAC and routing, the fundamental of these interconnects was the chassis backplane. In a DAA system, the backplane is the CIN. As an important first step, enabling the CIN to itself be Zero-Touch onboarded is important.

The management network DHCP uses a lookup to the DAA Inventory to learn the configuration file value for the CIN element assuming the device is onboarded in a configuration file download manner. This approach is similar to ‘golden config’ as the returned configuration file will have globally significant configuration based on the device role either Spine or s-Leaf. In either case at a minimum two main


configurations are enabled. First the direction to SDN Controller for ongoing management and Link Layer Discovery Protocol (LLDP) communications. Second access across all s-Leaf ports to the management network for DAA remote device onboarding. This includes the VLAN - VTEP association in all s-Leaf devices and the participation of EVPN forwarding from s-Leaf to Spine and Spine-to-Spine forwarding.



**Figure 7 - Initial Switch Config**

Next is the dynamic provisioning of service ports for each s-Leaf switch in the CIN. As each port should be capable of supporting any DAA service, we need to ensure discovery and provisioning at the edge are working perfectly. To achieve this, it is necessary to build a view of the complete topology.

Construction of the topology is possible through a graph of LLDP based information provided by all CIN devices. Several methods for asynchronously obtaining LLDP to a northbound system such as an SDN Controller are possible and include OpenFlow, sFlow, and NETCONF or RESTCONF streaming where each of these provide for the inherent MAC to Port information necessary to construct a dynamic view of the topology.

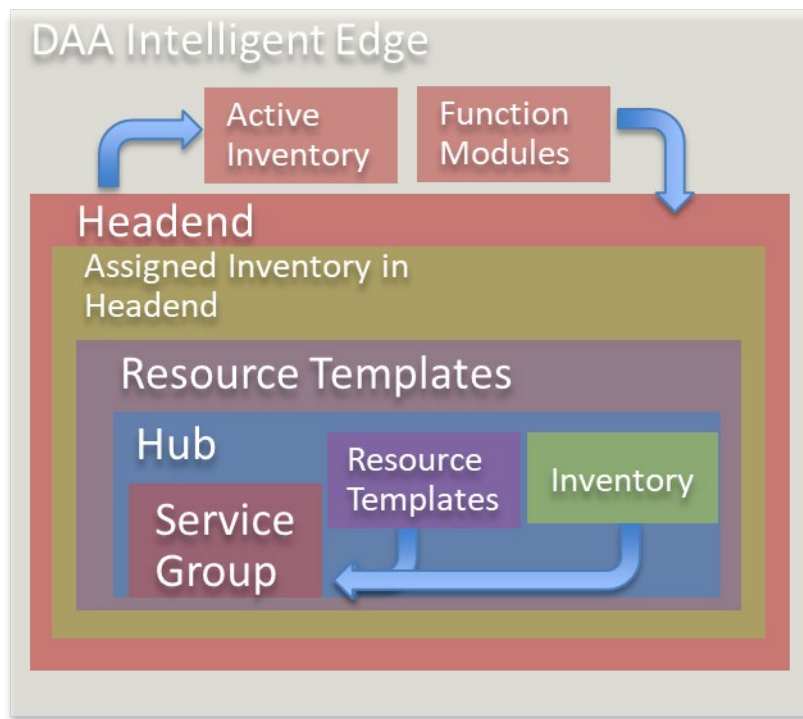


**Figure 8 - OSS Consumption of LLDP**

Additional resources in the example RPD Service Group 405 will need to associate with and involve the CIN. Specifically, a port on an s-Leaf switch.

The Intelligent Edge would also have knowledge then of the CIN elements, based on their assigned role as an ‘s-Leaf’ in the network at any location such as Headend or Hub serving Service Group 405.

At this stage, the basic system logic is in place that will support either a ‘Static-Inventory’ or a ‘Dynamic-Inventory’ based approach to deploying Remote-PHY and essentially any DAA remote technology.



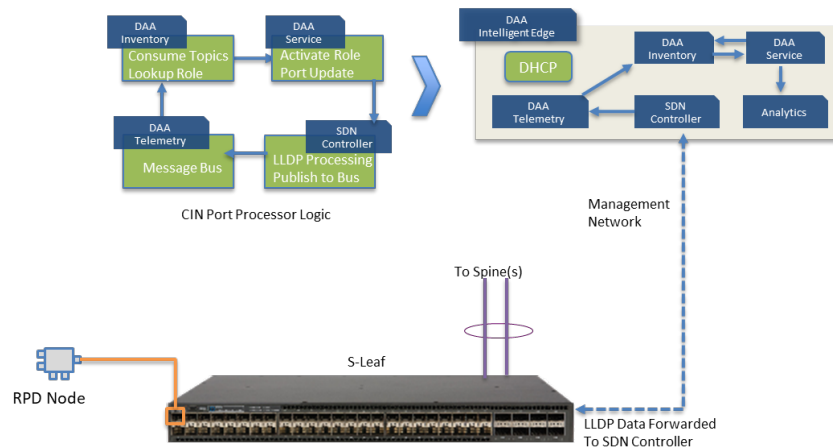
**Figure 9 - Intelligent Edge with Dynamic Inventory**

Using LLDP collection and SDN Controller to build a topology, the initial knowledge of edge port associations is made visible and accessible for automation.

The system design providing LLDP processing is quite robust. LLDP TLV data is exposed to the SDN Controller by any of the means described earlier. The SDN Controller transforms this LLDP packet data into a message bus topic and publishes this for consumption by the other functions of the Intelligent Edge.

In this system design, there is a DAA Inventory service always consuming topics from the message bus to remain in an 'active inventory' state for the entire network. As updates occur for device records, the Inventory system takes a decision if there is any change in the known data based on the information present in the message bus topic. If there is a change, the update is sent to a DAA Service Activation function which sends an API request to the SDN Controller. The SDN Controller then can modify the port profile configuration for the given service role communicated (such as RPD-Service or R-OLT Service or similar), which will inherit specific configuration values during the process.

If there is nothing to update in the network, the device record is simply updated in the DAA Inventory with any additional values such as counters or other operational values. No actions to DAA Service Activation occur in this case.



**Figure 10 - CIN Port Dynamic Profile Provisioning**

The result of this system design is all ports of the CIN are dynamically capable for any defined service known to the Intelligent Edge system as the CIN device record exists in Inventory and all ports have management network reachability. As any DAA device is attached, it is dynamically discovered before 802.1x or IP address stage such that Service Activation may occur just before the DAA device is provisioned.

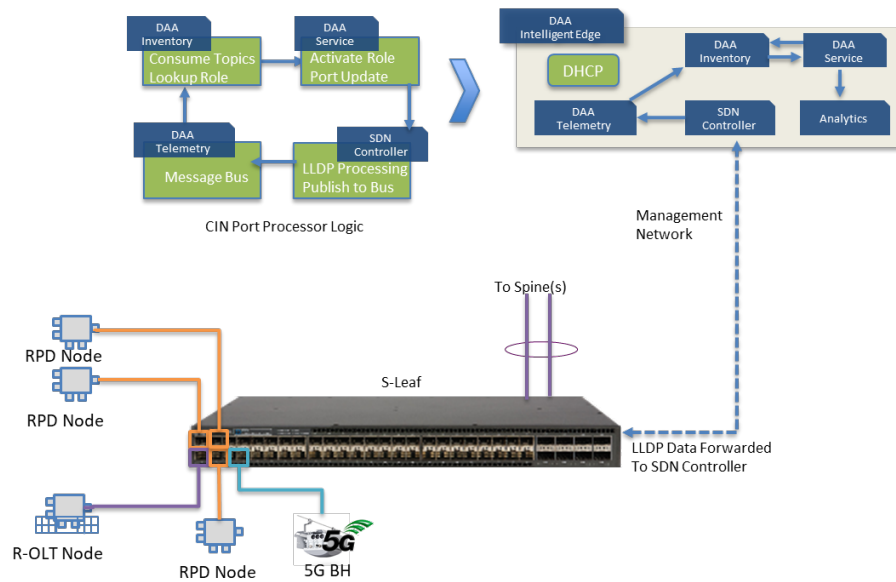
For the field technician who may be responsible for deploying Remote PHY or Remote-OLT PON, or 5G Backhaul, the simplicity at the edge could not be better. Similar to how an Integrated CCAP chassis automatically detects and onboards line cards in the chassis based on their position and role, the CIN now provides the same logic in a highly distributed and in some cases disaggregated fashion.

For the field installer, attaching a DAA remote device in the Intelligent Edge system becomes as simple as adding a line card to a CCAP chassis where the CCAP chassis has a supervision card or function on a common card watching for interrupts on the backplane that occur when a card is inserted.

When such an event occurs in a CCAP chassis, the control card will determine the type of card inserted and prepare it for operation, possibly upgrading the firmware of the card, then configuring the line card for operation.

In a distributed system, the line card, in part or in whole, is the DAA remote device, the chassis fabric is the Converged Interconnect Network and the supervisor or control card function is an OSS back office application.

For each s-Leaf port the switch sends an LLDP message with switch unique identifier, port identifier, and the attached device MAC address to the controller. The controller emits this to a common message bus where the Inventory and Service applications work together to support dynamic edge discovery and service control by consuming the new message bus topic. It may then be determined if an inventory update is required. If an update is required, it will initiate the action sending the Service Activation the matching service profile from DAA Inventory, where the s-Leaf port config is then updated based on the service profile, initiated by Service Activation, and executed by SDN Controller.




**Figure 11 - Any DAA Service Any Port**

## 5. CIN Toplogy

In our model, we have chosen to deploy a CIN based on Ethernet Virtual Private LAN (EVPN) and Virtual Extensible Local Area Network (VxLAN).

Ethernet Virtual Private LAN is a control plane for routing and forwarding Virtual Extensible LAN overlays. When these technologies are applied to the CIN, the back-office co-ordination tasks are greatly reduced as now the interconnect between s-Leaf switches over the Spine is handled within an autonomous system that acts like a forwarding fabric based on prefix identifier.



**Figure 12 - EVPN VxLAN CIN**

The EVPN with VxLAN design has the added benefit of isolating MAC address learning from the Spine. Virtual Tunnel Endpoints (VTEP) are exposed via the s-Leaf devices. Learning for MAC addresses, for example an RPD connecting to a Principal Core, is handled at each of those respective edges.



Essentially, the VTEP will encapsulate a tagged or untagged VLAN identifier from the incoming port of the s-Leaf and map it to a VxLAN identifier. The VxLAN then functions as an overlay network forwarded by prefix across the Spines for any other s-Leaf participant to join the created overlay service.

The Intelligent Edge may configure Spine devices with a uniform configuration to support all the s-Leaf devices, where the bulk of any CIN related CRUD (Create, Read, Update, Delete) actions will now occur to control VLAN to VTEP associations per device or service as required at the edge.

Multiple VLANs are possible per port, which enables the general theory of operation that management VLANs may be the default untagged VLAN for all ports of the s-Leaf in the CIN. Additional VLANs are likely to represent DAA subscriber service networks.

This management VLAN can reach the DAA Intelligent Edge OSS systems which include DHCP services for Dynamic Host Configuration Protocol IP addressing, as well as Authentication Authorization and Accounting (AAA) systems for 802.1x authentications in addition to the collection of topology related LLDP messages destined for the SDN Controller.

Additional VLANs may present per port based on the DAA remote service type, such as a remote OLT PON device with both a management interface for access to an OLT Manager in addition to a VLAN for subscriber IP internet services terminated by a Provider Edge (PE) router. In such a scenario, the PON Service would have two VLANs associated with two VTEP associations to associate forwarding across the Spine.

There are also further scenarios for remote OLT PON where multiple service providers may be served or a requirement to provision unique service separation, the same design of service profile and port configuration orchestration to VTEP associations are automated. The breakout of the PON service then occurs as a result of the PE element occupying a unique VTEP : VLAN Port appearance egressing an s-Leaf verses the management port egress to the back-office systems where the OLT Manager provides R-OLT provisioning and OAM functions with the remote device.

## **6. Dynamic Edge for DAA**

### **6.1. LLDP Stage**

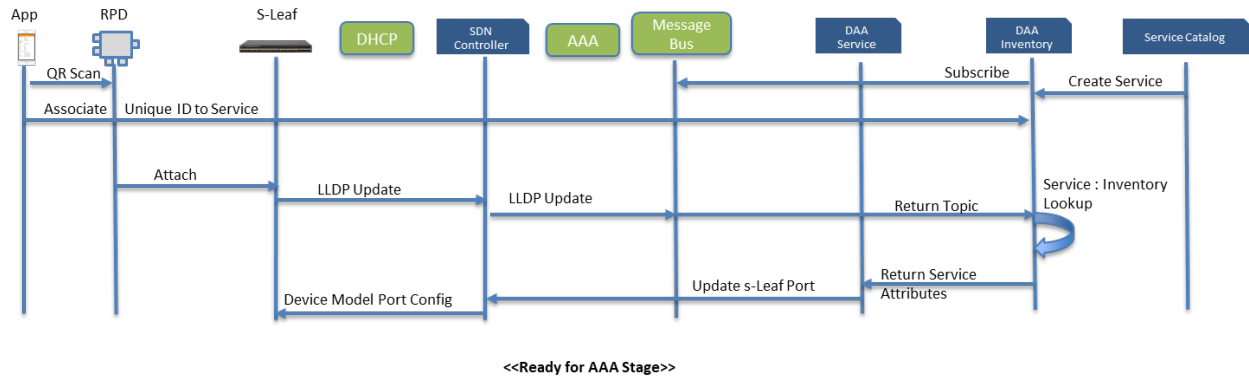
A large portion of our intelligence in the Intelligent Edge is achieved through learned changes in device operating state from the moment a device is attached to an s-Leaf port. The system is dependent on Link Layer Discovery Protocol (LLDP) packets being exposed to the DAA Intelligent Edge SDN Controller.

Each Ethernet frame contains an LLDP Data Unit (LLDPDU) which holds type length values (TLV) for several objects. These may include ChassisID, PortID, Time-To-Live (TTL), Port Description, System Name, System Description, System Capabilities, and Management Address. Additionally, vendor specific LLDPDU TLVs are possible.

Using the device MAC address with the LLDPDU contents when processed from all elements in the CIN becomes incredibly valuable to the DAA Intelligent Edge.

The use of LLDP permits the calculation of the topology graph by the controller. The topology is now available for interrogation by other systems in the Intelligent Edge back office. This is particularly useful when seeking to understand attached location or physical or logical forwarding path based on DAA remote device by a variety of OSS systems.

For our purposes, workflow is triggered from our DAA Intelligent Edge orchestration based on the identity of the attached device to an associated service.



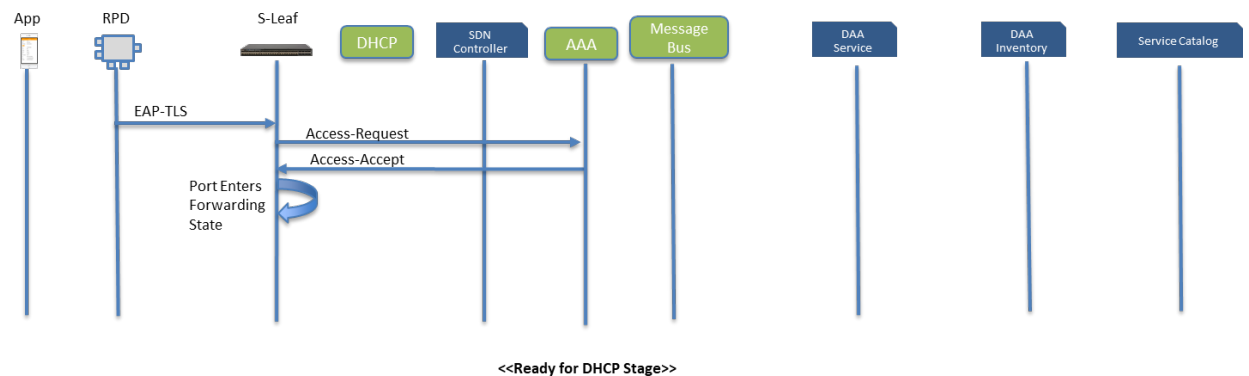
**Figure 5 - Intelligent Edge Discovery and s-Leaf Provisioning Sequence**

The Intelligent Edge Inventory holds the device records for all network elements in our DAA systems. Spine switches, s-Leaf switches, Remote PHY, Remote MACPHY, and Remote OLT PON devices all have device records in Inventory. By orchestrating the service to the device, we can now inherit dynamic provisioning of the DAA s-Leaf port.

Given the design assumption made in this paper has been a CIN based on VxLAN overlays with Spine forwarding using EVPN control, this means the association of a service based on the discovered device attachment to port permits ‘service based’ port configurations primarily consisting of VLAN and VTEP setup across the CIN.

## 6.2. 802.1x Stage

Given s-Leaf ports are configured by default to permit access to the management network, when CableLabs® standardized DAA devices attach, there is a requirement to process Extensible Authentication Process (EAP) to secure the edge network ports. EAP authentication in these terms is usually certificate driven meaning the CableLabs device will use its CableLabs X.509 based certificate or a valid vendor certificate chained from the Root-CA to complete the secure identity association.



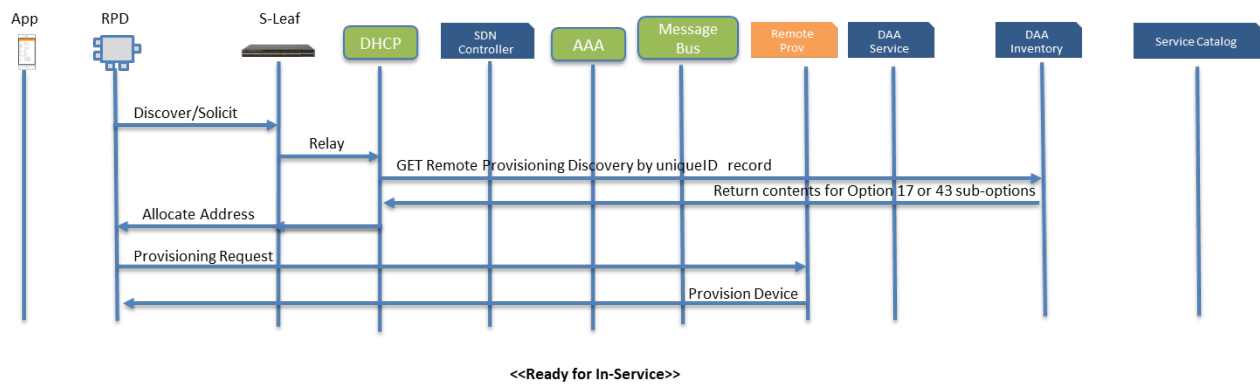
**Figure 6 - AAA Sequence**

To support this, each s-Leaf device provides an authenticator role to interwork with the AAA services in the back office. While the s-Leaf has informed the SDN Controller of the attached device by its MAC address and including the port information, Layer 2 traffic is not available until the authentication request is successful.

### 6.3. DHCP Stage

DHCP and DHCPv6 systems have formed a central role in the cable industry since the earliest dependencies on the protocol as effectively a control plane for dynamically provisioning cable modems. In DAA systems, we make use of DHCP and DHCPv6 to assist remote devices in the discovery of their management systems or some other functional actor that will complete the remote device provisioning stage.

For an RPD, this may be a Principal Core capable of configuring the RPD with DEPI, UEPI, and/or OOB RF resources over Generic Control Plane (GCP) protocol. For an R-OLT, DHCP may be used to discover the address of the OLT Manager system capable of provisioning the EPON remote OLT using some vendor defined means thereby supporting rapid deployment of CableLabs based DOCSIS Provisioning over EPON (DPoE) services.



**Figure 7 - DHCP Sequence**

DHCP systems can often enrich and expose attributes of the network edge to back office systems given they understand gateway relays that have provided edge routing of DHCP discovery and solicit messages to reach the DHCP / DHCPv6 system itself. This provides a hint about the edge of the network a device has attached to, however, it often lacks further details.

When DHCP is part of the DAA Intelligent Edge, it can be said that DHCP systems may enrich the DAA back office and the DAA back office may also enrich the DHCP systems.

# Conclusion

As exciting as the numerous new solutions are that came with these distributed system changes, DAA also brings with it a dramatic increase in managed elements in the network with relationships spanning edge into the eadend and beyond. Operators with early experience in DAA deployments have indicated challenges with Dynamic Host Configuration Protocol (DHCP) based back office integrations and a need to enhance the overall relationships of the back office to support the new distributed network in all its forms.

To achieve the business and technology goals of DAA, it will require automation and additional intelligence in the OSS back office systems that many would call Software Defined Networking. This is needed for the goal of returning to the operator the same level of autonomous operation that exists within CCAP chassis over distributed and multi-vendor network elements going forward.

In the proposed solution, the combination of orchestration concepts using a product catalog to model a DAA service were explored, application then of these DAA services based on device association were made real. Using SDN based automation and combination with DHCP systems, this specific solution enhanced the network edge for onboarding that presents both an initial step to added edge intelligence and a future option towards 'DAA as a Service' intent-based networks.

With multiple access network options, introduction of virtualization and wireless, we must co-ordinate the systems in the Hub and Headend while deploying any of these remote devices in a Zero-Touch approach whenever possible that maximizes the velocity for the operator overall and minimizes the time spent per field technician installing any of these edge solutions.

# Abbreviations

AAA	authentication authorization accounting
API	application programming interface
CIN	converged interconnect network
CCAP	converged cable access platform
DAA	distributed access architecture
DHCP	dynamic host configuration protocol
DOCSIS	data over cable service interface specification
DPoE	DOCSIS Provisioning of EPON
EAP	extensible authorization protocol
EPON	Ethernet PON
EVPN	ethernet virtual private network
FOSC	fiber optic splice closure
GCP	generic control plane
LLDP	link layer discovery protocol
OAM	operations, administration and maintenance
OLT	optical line termination
OSP	outside plant
OSS	operations support system
pCore	principal core

PE	provider edge
PNF	physical network function
PON	passive optical network
QAM	quadrature amplitude modulation
QoS	quality of service
PHY	physical interface
RF	radio frequency
RMD	remote macphy
R-OLT	remote olt
RPD	remote phy
SDN	software defined network
SFP+	small form pluggable optics module; the + indicates 10G capability
s-Leaf	secure leaf; a type of switch
TVL	type length value
vCore	virtual core
VNF	virtual network function
VLAN	virtual local area network
VxLAN	virtual extensible local area network
VTEP	virtual tunnel end point
WDM	wave division multiplexing
xPON	passive optical network

# **Delivering QAM Video in Distributed Access Architectures**

## **Architecture Options for Deployment**

A Technical Paper prepared for SCTE•ISBE by

**Colin Howlett**  
VP, Architecture  
Vecima Networks  
771 Vanalman Ave  
Victoria, BC, Canada V8Z3B8  
+1 (250) 881-1982  
colin.howlett@vecima.com

**Douglas Johnson**, Vecima Networks

**Kai Meisen**, Vecima Networks

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
QAM Video Delivery Architectures .....	4
1. Traditional Centralized QAM Video Delivery .....	4
2. Distributed QAM Video Delivery Options .....	8
2.1. Analog Overlay .....	8
2.2. Remote Edge QAM .....	10
2.3. Split Edge QAM Reference Architecture .....	11
3. Challenges in QAM Video Delivery using DAA .....	13
3.1. Content Encryption .....	13
3.2. Heterogenous Vendor Environment .....	14
3.3. Network Topology .....	14
3.4. All-IP Transition .....	14
3.5. Organizational Silos .....	14
4. Architecture Options .....	15
4.1. Integrated CCAP Core .....	15
4.2. Auxiliary Video Core .....	17
4.3. Standalone Principal Core + Video Traffic Engine .....	19
4.4. Separate Auxiliary Core and Video Traffic Engine .....	21
5. Video Traffic Engine Options .....	23
5.1. QAM RF Input .....	23
5.2. UDP IP Input .....	24
5.3. CDN Input .....	25
Conclusion .....	26
Abbreviations .....	26
Bibliography & References .....	27

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Traditional Centralized QAM Video Delivery – Standalone Edge QAM .....	5
Figure 2 – Traditional Centralized QAM Video Delivery – Integrated CCAP .....	6
Figure 3 – Traditional Centralized QAM Video Delivery – Hybrid Edge QAM + CCAP .....	7
Figure 4 – Distributed QAM Video Delivery Early Options – Analog Overlay .....	9
Figure 5 – Distributed QAM Video Delivery Early Options – Remote Edge QAM .....	11
Figure 6 – Distributed QAM Video Delivery – Split Edge QAM Reference Architecture .....	12
Figure 7 – Integrated CCAP Core Architecture .....	16
Figure 8 – Auxiliary Video Core Architecture .....	18
Figure 9 – Standalone Principal Core + Traffic Engine Architecture .....	20
Figure 10 – Separate Auxiliary Core + Traffic Engine Architecture .....	22
Figure 11 – Video Traffic Engine Options .....	23

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Integrated CCAP Core Architecture Capabilities .....	15
Table 2 – Auxiliary Video Core Architecture Capabilities .....	17
Table 1 – Standalone Principal Core + Video Traffic Engine Architecture Capabilities .....	19
Table 4 – Separate Auxiliary Core + Video Traffic Engine Architecture Capabilities.....	21
Table 5 – QAM Input Video Traffic Engine Capabilities .....	24
Table 6 – UDP IP Input Video Traffic Engine Capabilities.....	25
Table 7 – CDN Input Video Traffic Engine Capabilities.....	26



# Introduction

The shift from centralized access to distributed access architectures (DAA) represents a fundamental change in the operation of hybrid fiber-coax (HFC) networks. Operators are moving to DAA for a plethora of reasons including:

- Radio frequency (RF) signal improvements by generating signals at the node
- Evolve outside plant fiber network to an all-digital network to serve other Ethernet-based needs (wireless, Metro Ethernet)
- Hub space and power savings

While the DOCSIS portion of the HFC network is the primary focus of this DAA transition, operators cannot simply replace other key services with a full DOCSIS system. Traditional set-top boxes (STB) still require video to be delivered as MPEG Transport Streams (MPEG-TS) over J.83 QAM channels and many operators require out-of-band (OOB) signaling such as [SCTE 55-1] and [SCTE 55-2] to control those STBs, provide channel maps, program guide data, conditional access authorizations, and remote firmware upgrades.

When DAA standards were first being created back in 2014-2015, there was an implicit assumption that QAM video would be controlled and processed by a single Converged Cable Access Platform (CCAP) Core. Evolution of the overall DAA ecosystem to embrace other technologies such as virtualization and an interoperable multi-vendor environment has greatly expanded the possibilities for operators and a primary CCAP Core for QAM video delivery is just one of several architectures that may be used.

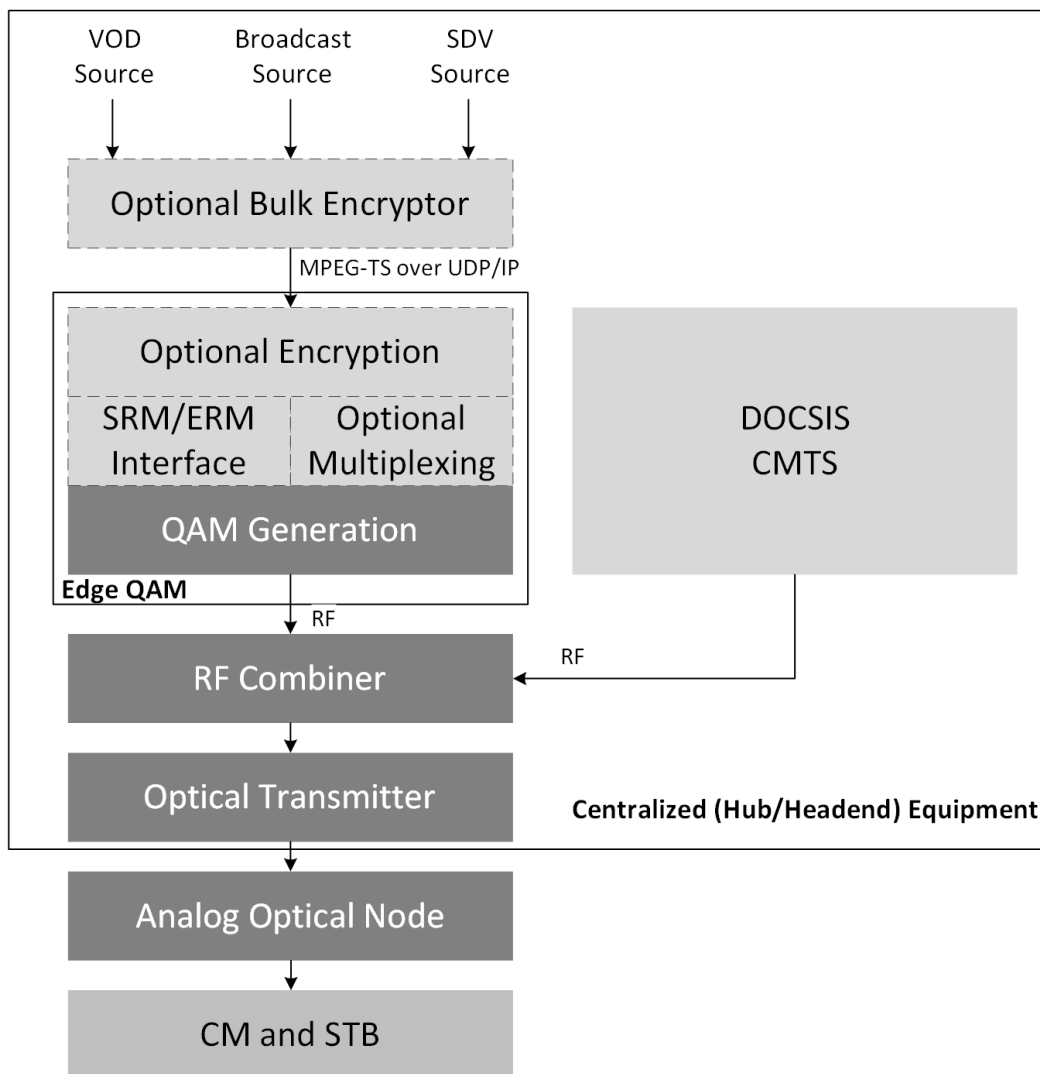
This paper outlines the challenges in delivering QAM video using DAA and compares the architectural options available to vendors and operators with specific examples of real-world operator feedback as part of early DAA deployments.

## QAM Video Delivery Architectures

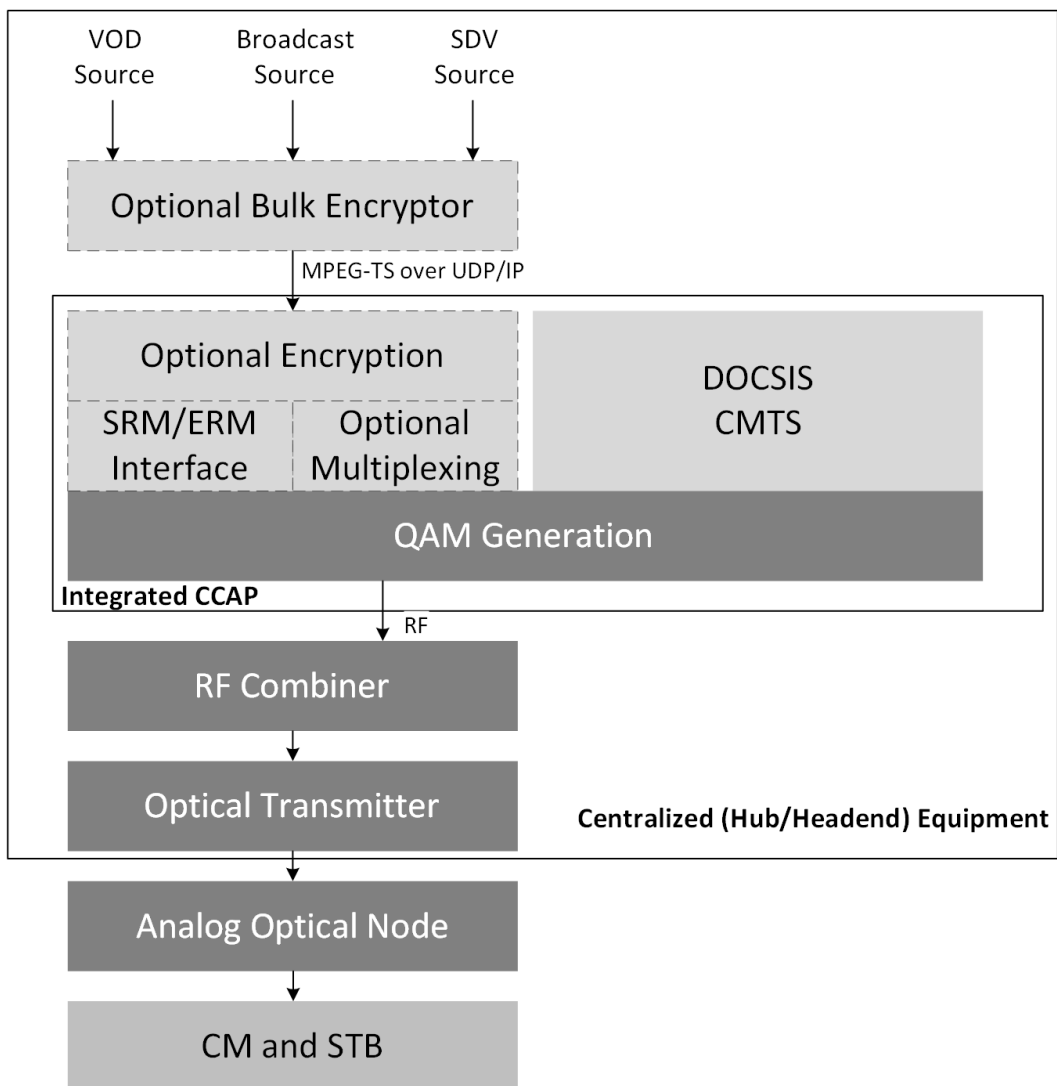
### 1. Traditional Centralized QAM Video Delivery

In order to understand the challenges and issues in QAM video delivery using DAA, we must first review the architectures used for traditional centralized QAM video delivery. The figures below show two architectures plus a third hybrid that represent the typical deployments today.

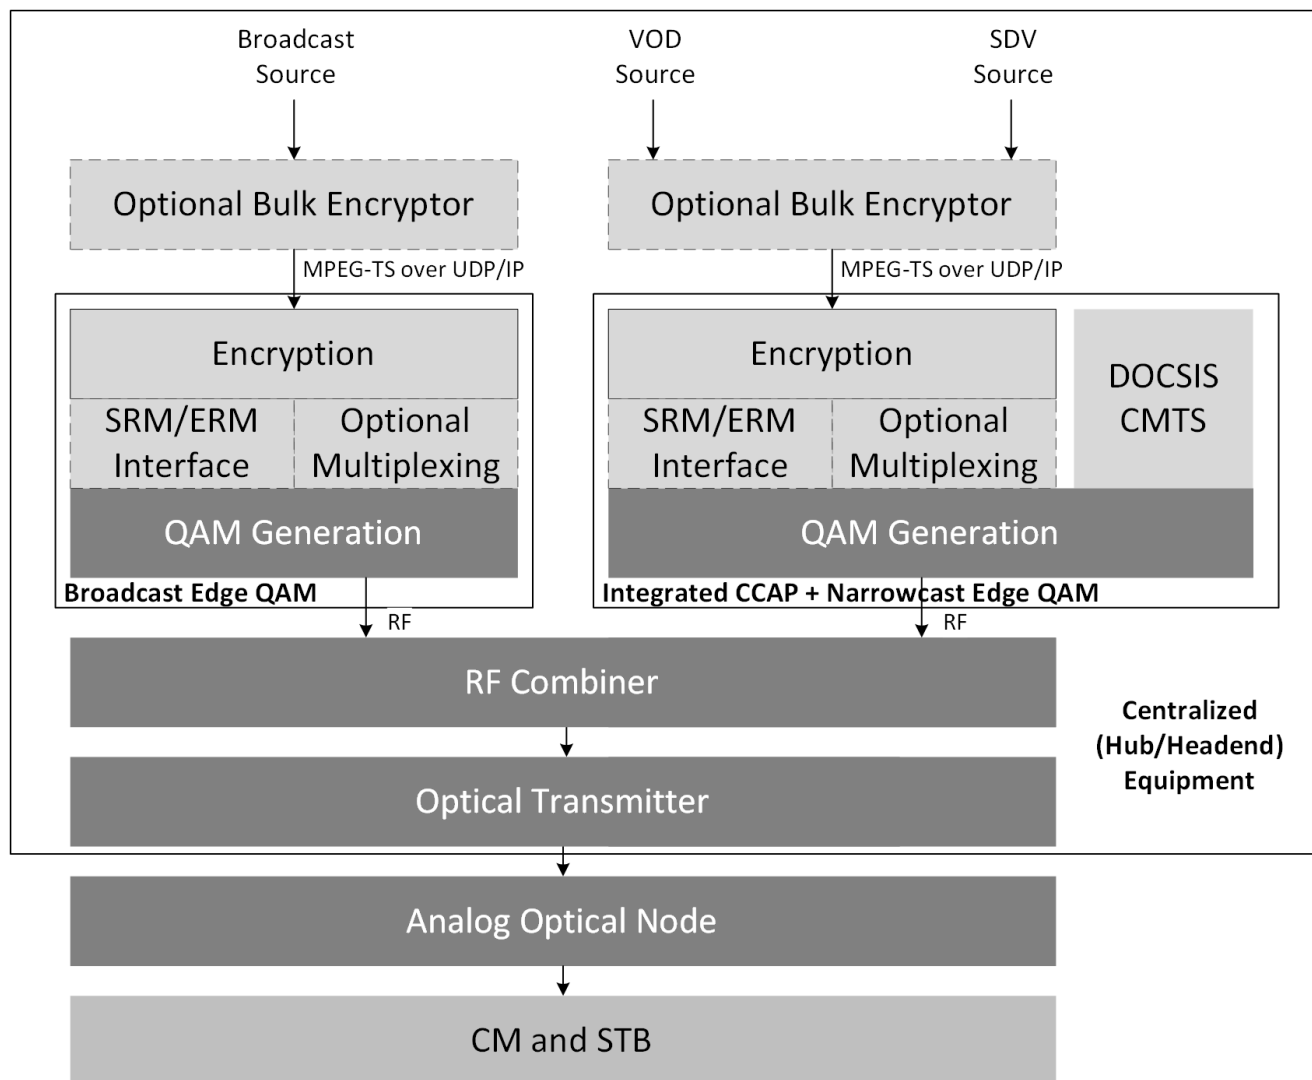
- Standalone Edge QAM (EQAM)
- Integrated Converged Cable Access Product (CCAP)
- Edge QAM + Integrated CCAP



**Figure 1 – Traditional Centralized QAM Video Delivery – Standalone Edge QAM**



**Figure 2 – Traditional Centralized QAM Video Delivery – Integrated CCAP**



**Figure 3 – Traditional Centralized QAM Video Delivery – Hybrid Edge QAM + CCAP**

In each of these architectures, the edge QAM functionality includes Session Resource Manager (SRM) and Edge Resource Manager (ERM) interfaces for narrowcast video channel allocation, content encryption, multiplexing and de-jittering of the input MPEG transport streams (MPEG-TS) and generation of the QAM signals.

In the Standalone Edge QAM architecture, DOCSIS and QAM video are handled separately and combined as RF before delivery to the subscriber via the optical transmitter and analog optical node.

In the Integrated CCAP architecture, all edge QAM functionality has been paired with the DOCSIS CMTS functionality into a single integrated CCAP device and Video-DOCSIS RF combining is removed.

A hybrid of the two is common in cases in North America where the CCAP vendor isn't the same as the conditional access (CAS) vendor. In this case, broadcast encryption cannot be supported in the CCAP device and a dedicated edge QAM is used for the broadcast channels and combined in at RF.

## **2. Distributed QAM Video Delivery Options**

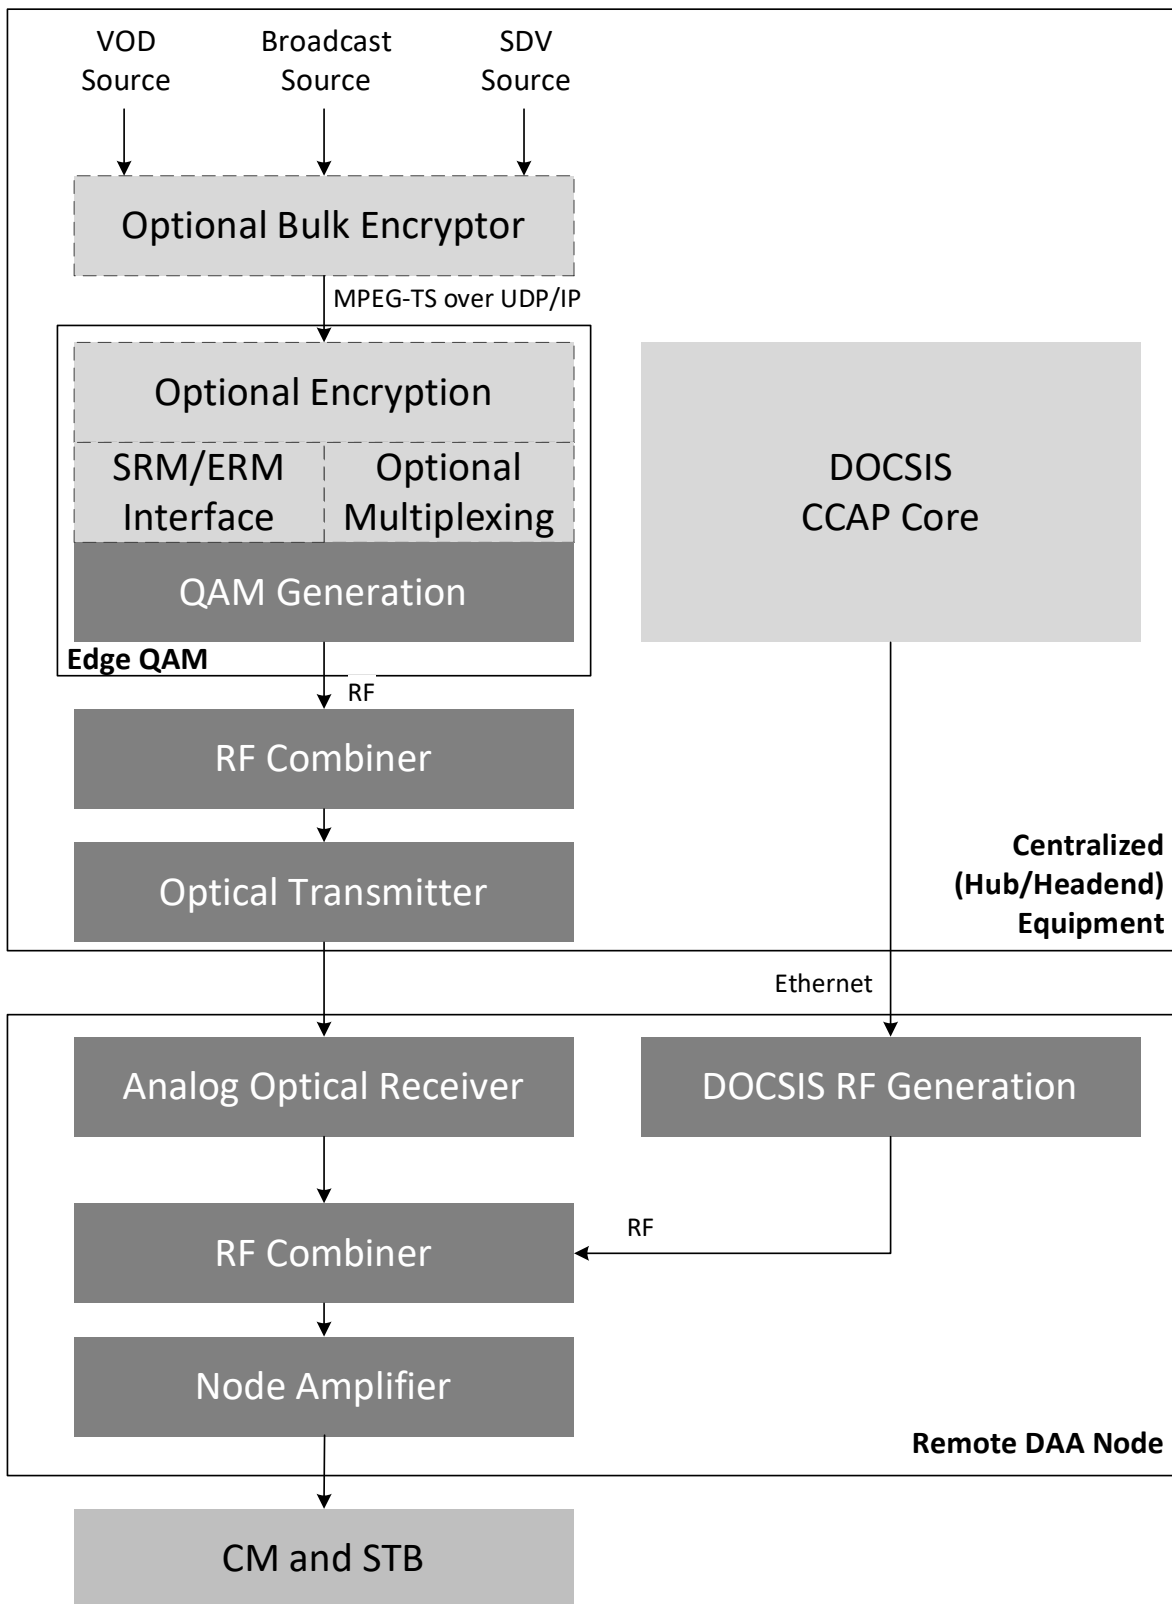
Early in DAA development, there was significant divergence in how QAM video would be delivered to remote node devices over the fiber portion of the HFC network. Three strong candidates were considered:

### **2.1. Analog Overlay**

One option, shown in Figure 4, is to keep analog fiber in place and combine this with the remote DOCSIS functionality. This analog overlay may include all the QAM channels plus the OOB or just some portion of the QAM channels (just broadcast).

Using analog overlay avoids concerns with potential complexity in QAM and overall video implementation, but has significant disadvantages as both remote nodes and the network evolves:

- 1) Analog fiber distribution still needed – signals are still dependent on analog RF distribution over fiber including distance-related SNR limitations and a reduced number of wavelengths usable for all-digital devices due to a limited amount of wavelength division multiplexing
- 2) RF combining in the node – two separate signals must be combined in the node compared to generating all the signals in alternate DAA approaches
- 3) Digital predistortion – the use of digital predistortion in remote nodes, driven by high integration of signal generation in SoC/FPGA solutions, can save 10s of watts in fiber deep scenarios. DPD requires that all signals are generated and available in the digital domain. High performance analog overlay systems generally prevent the use of digital predistortion and will limit future capability to improve overall outside plant power consumption



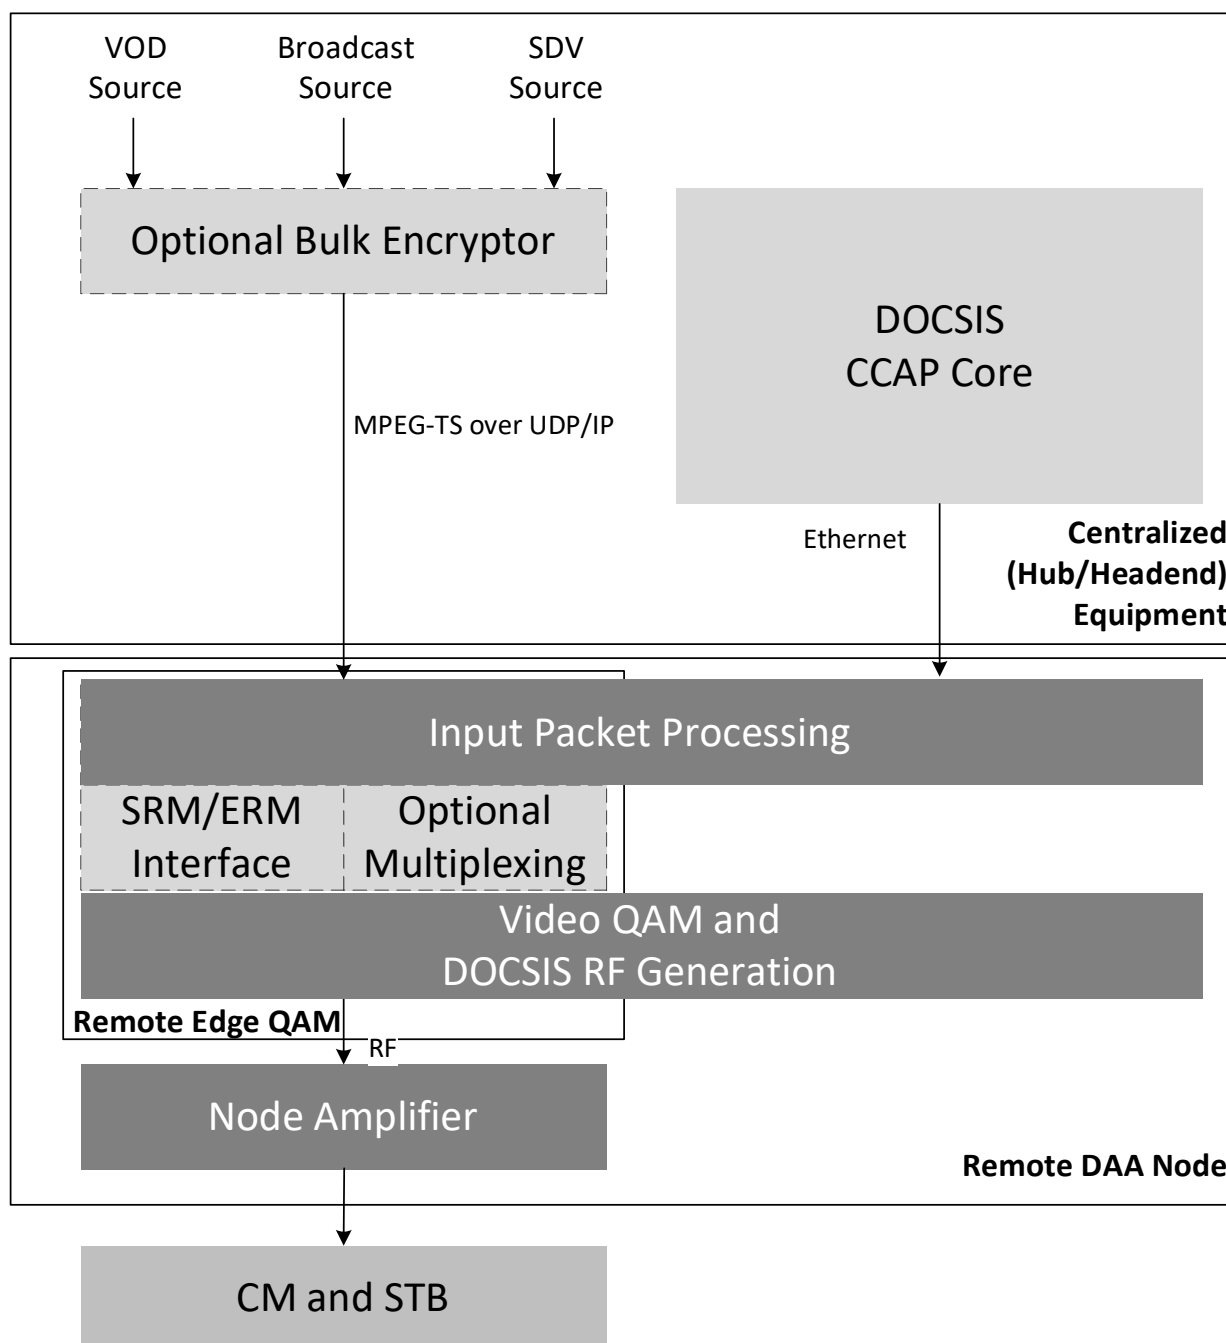
**Figure 4 – Distributed QAM Video Delivery Early Options – Analog Overlay**

## **2.2. Remote Edge QAM**

A second option, as shown in Figure 5, is to fully distribute the edge QAM functionality to the remote nodes.

While this solution may minimize space requirements within the hub/headend, there are several distinct disadvantages:

- 1) Node complexity – adding full multiplexing, encryption, and the need to provide appropriate interfaces to the ERM/SRM increases the amount and complexity of the software within the remote node.
- 2) Encryption security – encryption functions have very high levels of hardware and software applied to prevent the accidental disclosure of secrets related to conditional access operation. This places a high burden on the remote device in an untrusted domain (outside plant, basement of an apartment building) compared to a secure location within the hub/headend.
- 3) Duplication of functionality – in many deployment cases, especially for fiber deep architectures, the number of homes included in a video service group is many times the number of homes in a DOCSIS service group. Requiring full remote edge QAM functionality duplicates that power and functionality all over the outside plant.

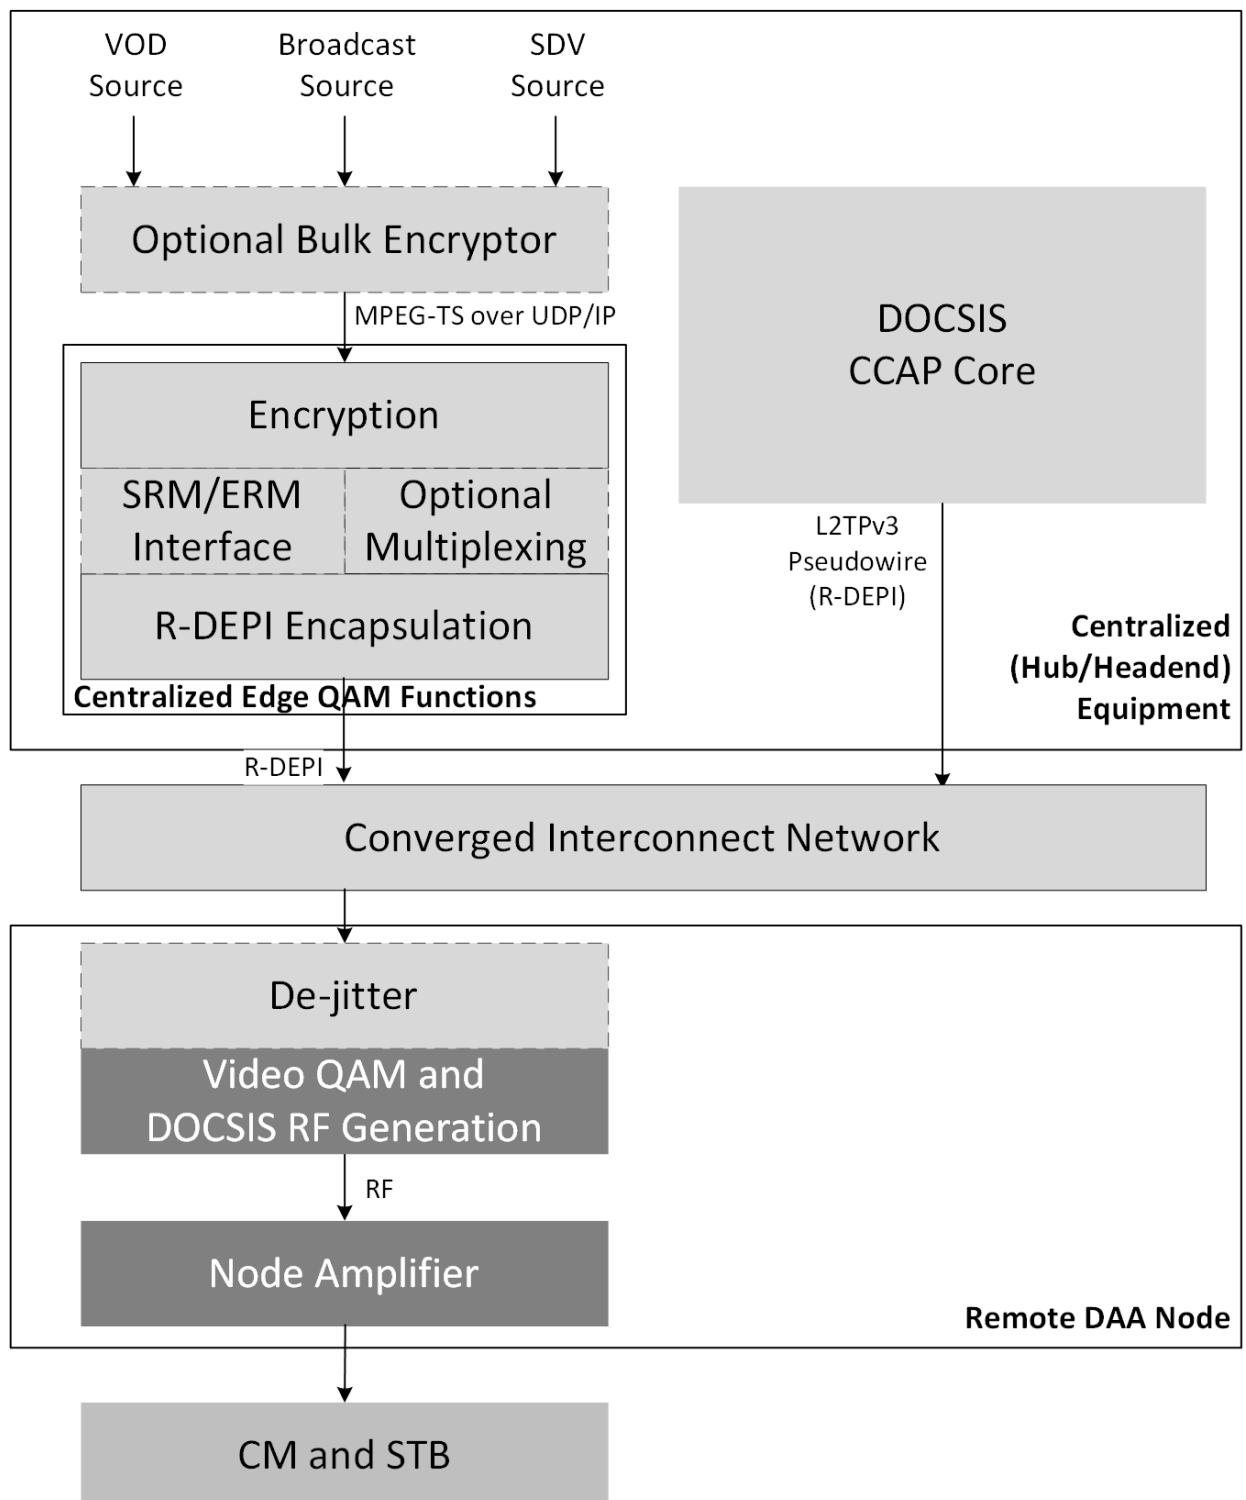


**Figure 5 – Distributed QAM Video Delivery Early Options – Remote Edge QAM**

### **2.3. Split Edge QAM Reference Architecture**

Given the concerns with analog overlay and remote edge QAM, the vendor and operator community settled on split edge QAM as the reference architecture detailed in MHAv2 for R-PHY. The high-level architecture for DAA video delivery specified in R-PHY is shown below in Figure 6. This architecture is used both for R-PHY and will be re-used as part of the next generation CableLabs Flexible MAC Architecture specification with centralized video elements able to support both types of remote devices.





**Figure 6 – Distributed QAM Video Delivery – Split Edge QAM Reference Architecture**

Vendor and operator development work on deployment architectures throughout the R-PHY specification led to several key elements also being included as part of the standardized solution:

- Separation of control plane and data plane into Cores and Traffic Engines

Cores contain control plane functionality, including either L2TPv3 control plane signaling or Generic Control Protocol (GCP) statically configured pseudowires. Cores may also contain associated data plane functionality (for example a DOCSIS Core contains both).

Traffic Engines only provide data plane functionality. In the case of QAM video delivery, a Video Traffic Engine only provides statically configured multicast pseudowire (R-DEPI) output for processing by the RPD. The RPD is configured by an associated Core to listen on the appropriate pseudowire and output to a specific QAM channel.

Further details on Cores and Engines can be seen in [Rahman].

- Simplified remote device – limited QAM functions

As shown in the diagram above, centralized edge QAM elements provide fully-formed line rate MPTS for the RPD. In the simplest case, the RPD is only responsible for de-jittering network contributions and generating the QAM signal.

Pseudowire operation may be used in either synchronous or asynchronous mode. Synchronous mode operates with both the R-DEPI Traffic Engine and the remote node synchronized using R-DTI (based on IEEE-1588 PTP). Synchronous mode is the mandatory mode of operation for remote nodes. Many remote node vendors also support the optional asynchronous mode video which allows for null stuffing/deletion and PCR restamping to avoid the need for R-DTI synchronization of the R-DEPI Traffic Engine.

The flexibility of the interfaces and interoperable standardized operation allows optimization of the QAM video delivery architecture to suit specific operator needs as discussed in section 4.

### **3. Challenges in QAM Video Delivery using DAA**

The shift to DAA requires the operator to address several challenges in providing QAM-based video to existing STB.

#### **3.1. Content Encryption**

Many operators, especially in North America, have limited options for the devices that can be used to encrypt video transport streams for conditional access and content protection. Some conditional access vendors only support broadcast channel encryption through their own edge QAM or CCAP devices, but bulk network encryptor solutions are becoming more commonplace.

The encryption system in use has a significant impact on which architecture can be deployed by the operator, especially if the CCAP Core vendor of choice does not match the conditional access vendor.

A large percentage of early deployments of R-PHY in North America have seen a mismatch between the conditional access vendor and the CCAP Core vendor resulting in the need for auxiliary core solutions to deliver broadcast video.

### **3.2. Heterogenous Vendor Environment**

Cable operators naturally operate with a mix of vendors in their networks. Operators may have a mix of DOCSIS CMTS/CCAP vendors in different parts of their network, all of which could be upgraded to support R-PHY. In some cases, the video solution is common across the network with minimal interaction between the systems since they all interface at the RF level in the combiner.

QAM video delivery architectures which focus all services (DOCSIS and video) through a single device now make this multi-vendor mix more complicated since all the video backend must support integrations with each of those vendors. Separating the DAA QAM video solution from the DOCSIS solution can better support this mixed vendor environment by minimizing the amount of video backend integration.

### **3.3. Network Topology**

Operators who serve lower density communities are turning to DAA to remove the need to deploy large CCAP platforms at each community. Centralizing the DAA components at a regional headend vs. small community hub locations often provides significant cost savings in operational expenditures (facilities consolidation) and capital expenditures (sharing a larger CCAP platform across multiple communities and getting closer to full density).

In several real deployments though, the optimum DOCSIS Core location is further from the remote nodes than the video location. This happens in cases where HITS or similar satellite-based video distribution methods are used for video. If a CCAP is used as both DOCSIS Core and Video Core, video traffic must be “hairpinned” back to the DOCSIS Core location before R-DEPI encapsulation and transport to the remote nodes. This significantly increases fiber capacity needs to those small communities so architectures which can support deeper distribution of R-DEPI encapsulation are highly desirable.

### **3.4. All-IP Transition**

Operators are now starting to embrace all-IP video delivery to take advantage of lower cost STB solutions, support non-STB mobile devices, and support more rapid advancements in video service offerings by using a platform complementary to over-the-top content providers.

This new generation of all-IP video services transition delivery of video from a UDP streaming mechanism to a content delivery network (CDN) consisting of origin servers and several levels of caches deployed throughout the operator network. Ad insertion, transcoding, event blackouts, and many other video processing elements all operate differently than in a traditional MPEG-TS environment.

There are significant operational expenditure benefits to merging the delivery to new all-IP devices and traditional QAM STB into a single unified CDN.

### **3.5. Organizational Silos**

Many MSO engineering organizations have separate video teams and access teams. Some integration of these teams has happened in the move to IP video and where operators have moved more aggressively to deliver both DOCSIS + QAM video through an integrated CCAP, but it is still common to have separate DOCSIS and video teams.

DAA QAM video architectures which consider the organizational issues and keep QAM video separate from DOCSIS may be more successful in getting to deployment earlier at lower cost. These organizational issues can overwhelm technical merit of different architectures since the burden to

implement those changes is less than deploying new equipment architectures in each silo. The key for operators is to recognize their internal capabilities and focus on architectures which can be successfully deployed.

## 4. Architecture Options

Several architecture options are identified in sections below. Each of these options is capable of multi-vendor interoperability between Core elements and remote nodes and compliant to CableLabs R-PHY and anticipated FMA standards. An operator may even choose to deploy different architectures for broadcast and narrowcast QAM video to suit the needs of their system architecture.

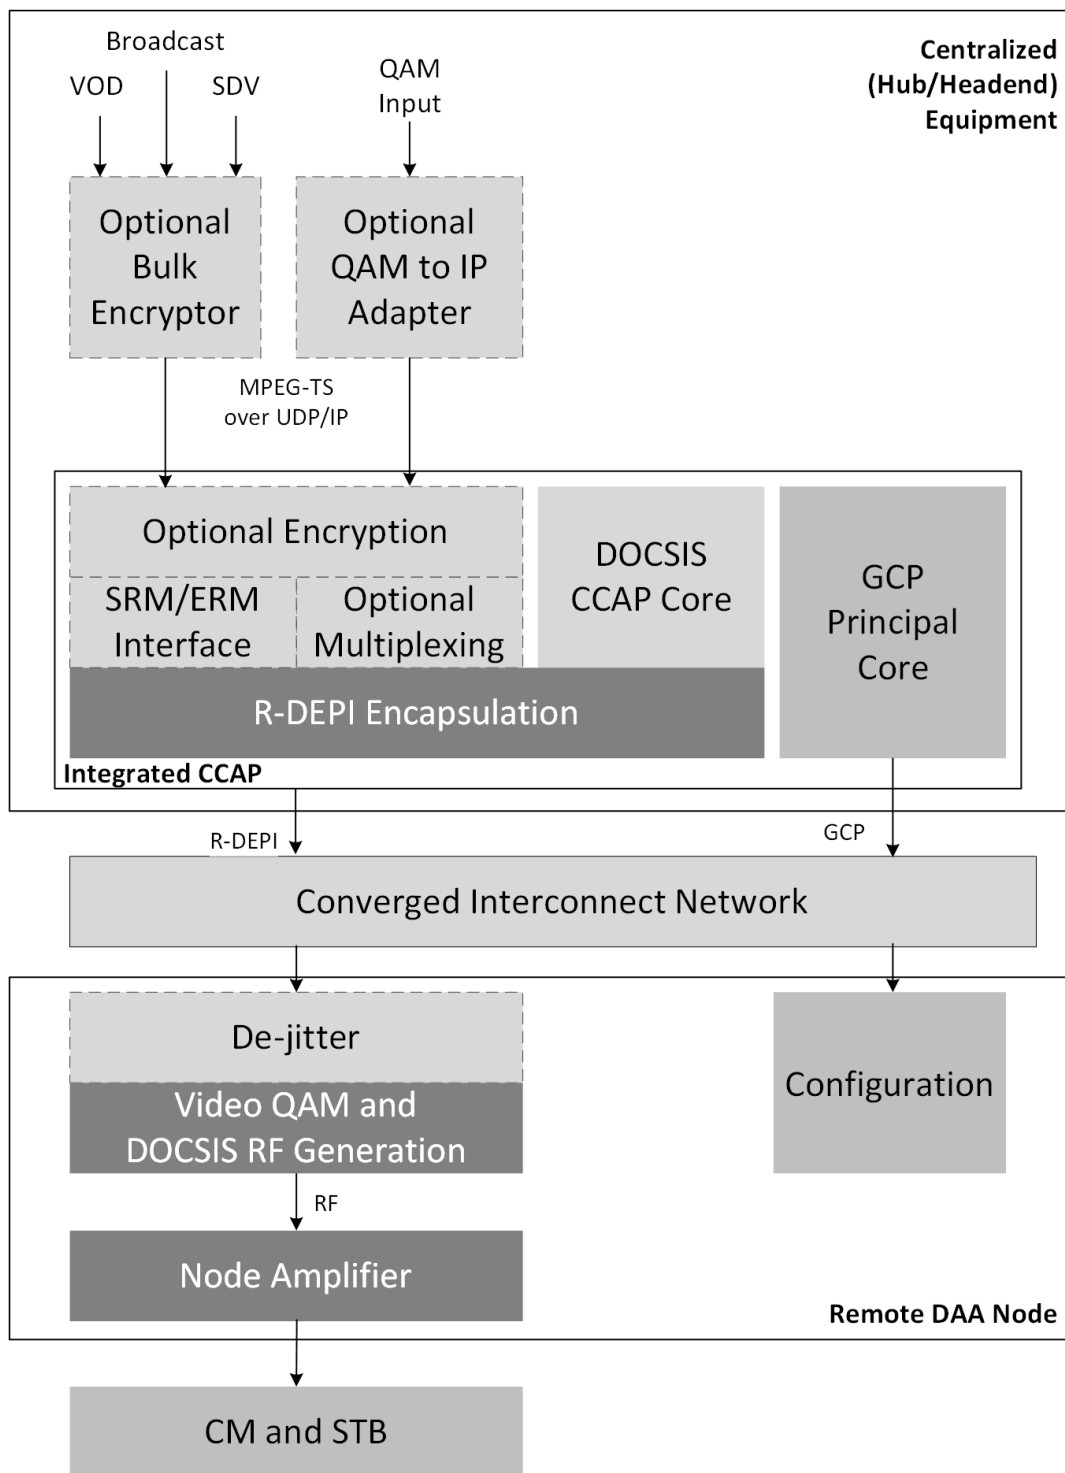
The capability of each architecture option to address the challenges identified in the previous section is listed along with information on typical usage scenarios for each architecture based on real world deployments. Details on Video Traffic Engines are included in section 5.

### 4.1. Integrated CCAP Core

The original assumption and starting point for R-PHY DAA assumed an architecture, shown in Figure 7, that used an integrated CCAP Core to fulfill both DOCSIS CMTS and video EQAM requirements. Content encryption may be provided externally through a bulk network encryptor or a QAM to IP adapter that provides pre-encrypted transport streams to the integrated CCAP Core.

**Table 1 – Integrated CCAP Core Architecture Capabilities**

Attribute	Capability
Encryption	Neutral - requires external broadcast encryption solution if Core vendor doesn't match CAS vendor
Mixed Vendor	Poor – monolithic vendor for both video and DOCSIS functions
Network Topology	Poor – integration of DOCSIS and video in the same device limits flexibility in where the two functions reside
All-IP Transition	Neutral – direct CDN input not available but MPTS passthrough allows straightforward connection to a CDN Input Video Traffic Engine
Organizational	Poor – requires coordination of video and DOCSIS teams using the exact same device
Typical Usage Scenarios	<ol style="list-style-type: none"> <li>1. MSO with pre-existing centralized integrated CCAP using DOCSIS + video</li> <li>2. MSO using QAM Input Video Traffic Engines in MPTS passthrough mode</li> </ol>



**Figure 7 – Integrated CCAP Core Architecture**

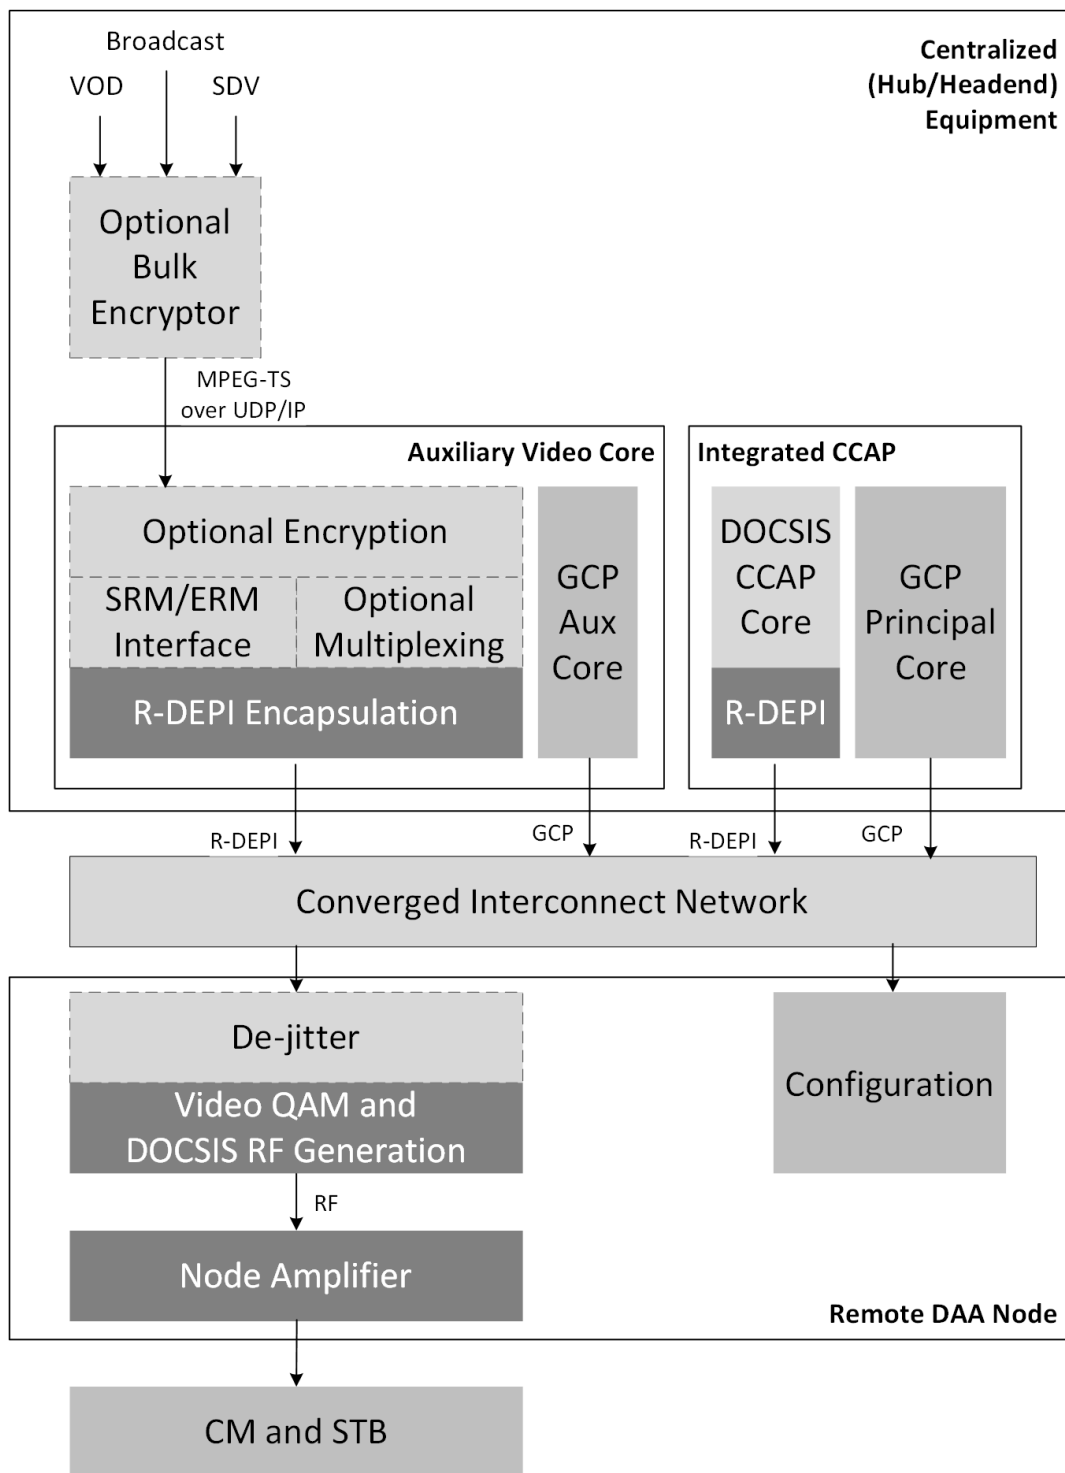
## 4.2. Auxiliary Video Core

Video QAM delivery can be implemented as a fully standalone core separate from the DOCSIS Core as shown in Figure 8. This Auxiliary Core contains both control plane (either dynamic L2TPv3 or GCP controlled static L2TPv3) and data plane R-DEPI encapsulation functions.

The Auxiliary Video Core may also integrate traditional EQAM processing functions and could be implemented as a virtual function or as part of a high density EQAM platform upgraded for DAA use.

**Table 2 – Auxiliary Video Core Architecture Capabilities**

Attribute	Capability
Encryption	Neutral – requires external broadcast encryption solution if Auxiliary Video Core vendor doesn't match CAS vendor
Mixed Vendor	Good – allows video to be completely separated from DOCSIS
Network Topology	Good - can be located separate from the DOCSIS Core wherever the video solution may be needed
All-IP Transition	Neutral – highly dependent on vendor implementation
Organizational	Good – supports separation of video and DOCSIS requirements and responsibilities
Typical Usage Scenarios	<ol style="list-style-type: none"><li>1. Reuse of existing high-density edge QAM platforms for narrowcast DAA</li><li>2. Alternate to Traffic Engines if auxiliary video core implemented by CAS vendor</li></ol>



**Figure 8 – Auxiliary Video Core Architecture**

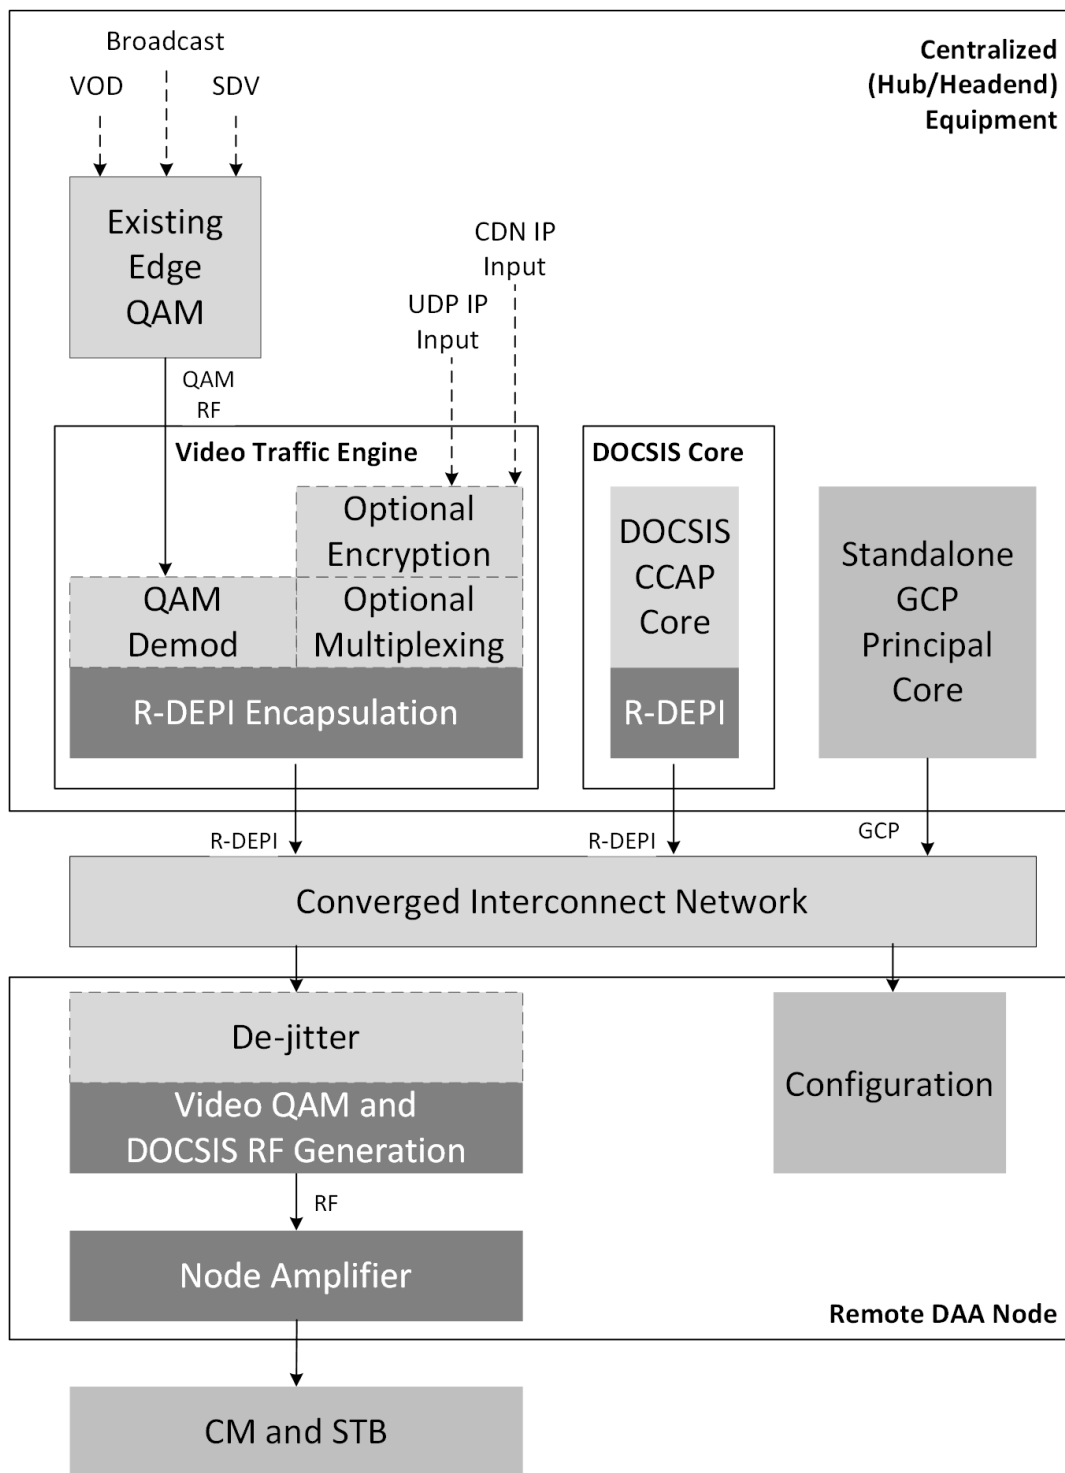
### 4.3. Standalone Principal Core + Video Traffic Engine

The R-PHY specifications support the implementation of a Principal Core as a standalone configuration-only functional entity, separate from the DOCSIS Core and the video data plane. See [Rahman] for background on the operation of this mode. In this architecture, shown in Figure 9, the separate Principal Core is responsible for overall configuration functions including the static L2TPv3 pseudowire setup for Video Traffic Engines. The Video Traffic Engine (see section 5 for input options) is responsible for R-DEPI encapsulation to the remote node.

**Table 3 – Standalone Principal Core + Video Traffic Engine Architecture Capabilities**

Attribute	Capability
Encryption	Good – encryption handled by Video Engine implementation or by existing encryptor investments (edge QAM or bulk)
Mixed Vendor	Good – maintains configuration control in a single entity for flexibility when deploying a mix of DOCSIS and video solutions/vendors
Network Topology	Good - can be located separate from the DOCSIS Core wherever video solution may be needed
All-IP Transition	Good – can support a CDN input Traffic Engine
Organizational	Neutral – separates video from DOCSIS but requires cross-coordination amongst teams on the joint Principal Core function
Typical Usage Scenarios	<ol style="list-style-type: none"><li>1. Highly virtualized DAA deployment where DOCSIS Core doesn't act as a "primary" core</li><li>2. Mixed DOCSIS Core vendor deployments where a standalone Principal Core can remove the need to integrate OSS with multiple different DOCSIS Cores and associated vendor orchestration/provisioning tools</li></ol>





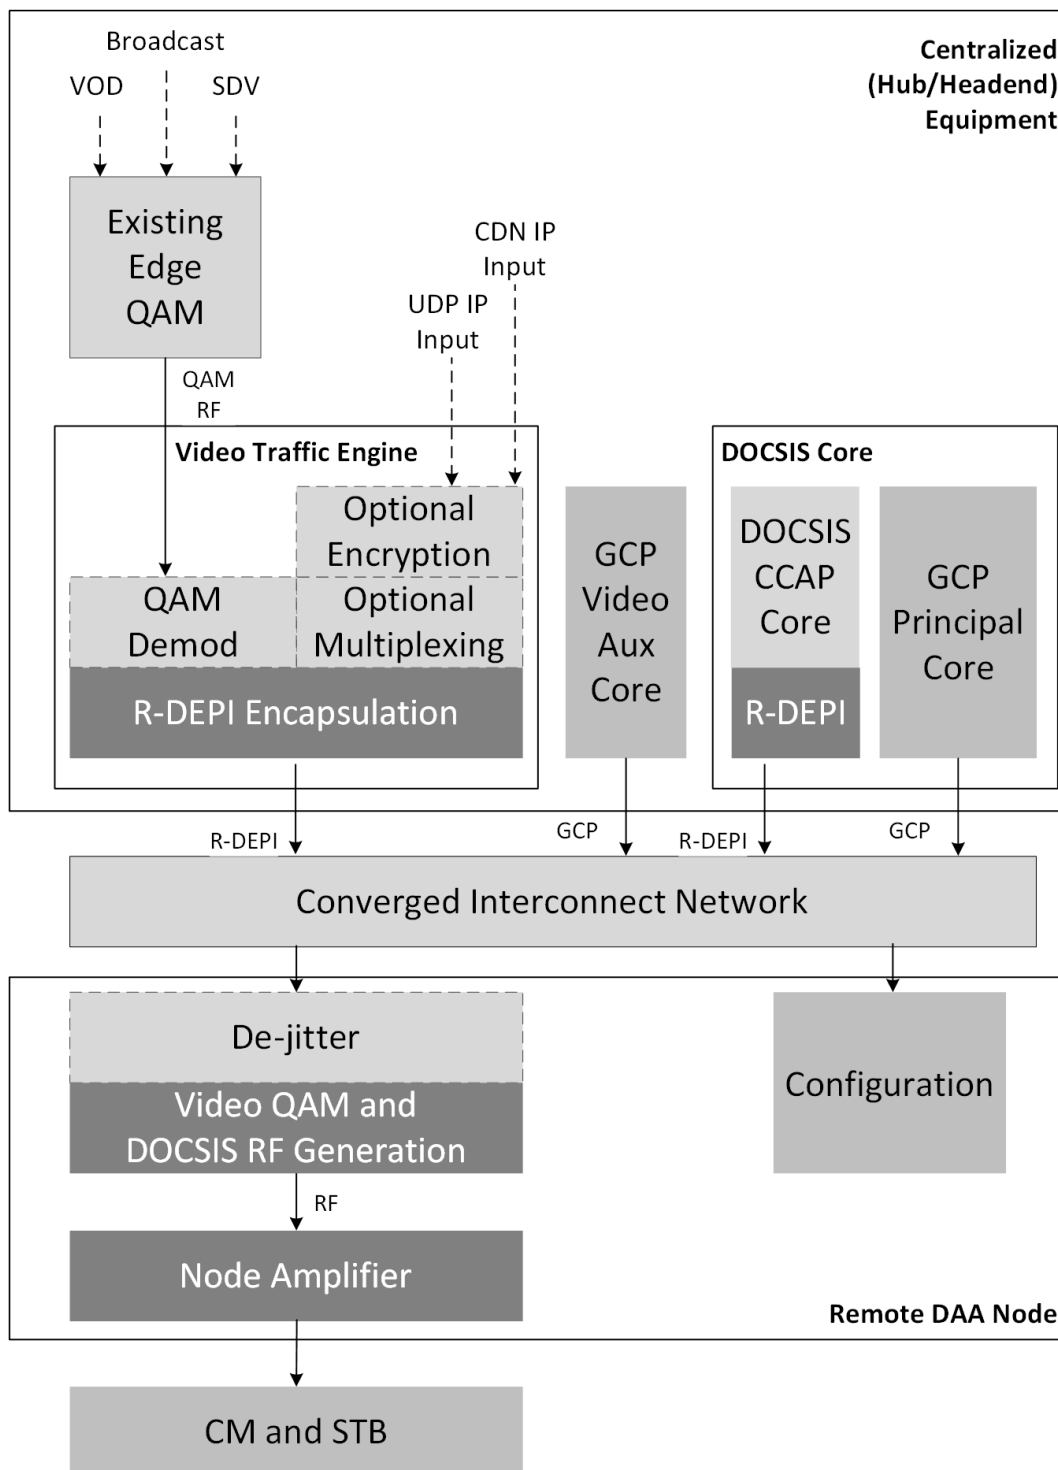
**Figure 9 – Standalone Principal Core + Traffic Engine Architecture**

#### 4.4. Separate Auxiliary Core and Video Traffic Engine

This architecture, shown in Figure 10, is similar to the Auxiliary Video Core architecture, except the configuration functionality is separated from the data plane functionality. This architecture supports a single configuration Auxiliary Core subtending many data plane Video Traffic Engines, allowing the different functions to scale independently. By separating the configuration and traffic responsibilities, this architecture also supports placing the functions at different network topology locations, such as Traffic Engines near a HITS reception location and the Auxiliary Core in a central data center.

**Table 4 – Separate Auxiliary Core + Video Traffic Engine Architecture Capabilities**

Attribute	Capability
Encryption	Good – encryption handled by Video Engine implementation or by existing encryptor investments (edge QAM or bulk)
Mixed Vendor	Good – maintains configuration control in a single entity for flexibility when deploying a mix of DOCSIS and video solutions/vendors
Network Topology	Good - can be located separate from the DOCSIS Core wherever video solution may be needed
All-IP Transition	Good – can support a CDN input Traffic Engine
Organizational	Good – separates video from DOCSIS in both data and control plane
Typical Usage Scenarios	<ol style="list-style-type: none"><li>1. Deployments where DOCSIS Core isn't the same as CAS vendor</li><li>2. Deployments where keeping DOCSIS and video separate is important for Core capacity, licensing, organizational or other reasons</li><li>3. Support of virtualized DOCSIS Core deployment with separate video solution</li><li>4. Mixed DOCSIS Core vendor deployments where a separate single vendor video core can remove the need to integrate video backend with multiple different CCAP Cores</li></ol>

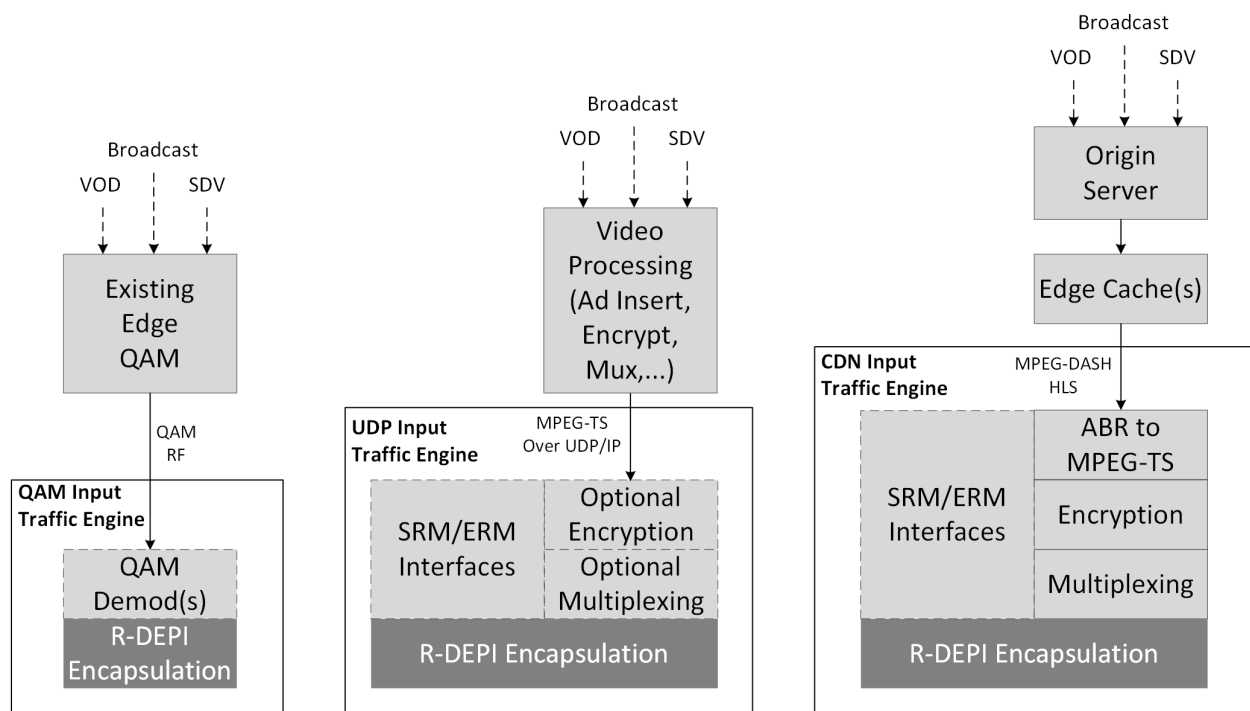


**Figure 10 – Separate Auxiliary Core + Traffic Engine Architecture**

## 5. Video Traffic Engine Options

A functional separation of control plane elements (“Cores”) and data plane elements (“Traffic Engines”) is fully supported by the R-PHY specifications. Architectures utilizing control and data plane separation were discussed in Section 4.3 and Section 4.4. When separating the functions, the data plane element is made simpler by having static pseudowires which use L2TPv3 for encapsulation and do not have a dynamic L2TPv3 control plane and in-band setup. Instead, the Core uses GCP to configure the correct static multicast pseudowire elements on the remote node. In the specific case of QAM video, the static pseudowire is multicast R-DEPI which allows the Traffic Engine to stream continuously and the network takes care of ensuring the packets are delivered to the remote node(s).

Traffic Engines for QAM video delivery all output multicast R-DEPI pseudowires, but may take several options as input, as shown in Figure 11.



**Figure 11 – Video Traffic Engine Options**

### 5.1. QAM RF Input

Video Traffic Engines using QAM RF Input, left side of Figure 11, allow the re-use of existing deployed edge QAMs in a DAA network. QAM channels from the existing edge QAMs are demodulated and then encapsulated in R-DEPI. This option allows for maximum reuse of existing deployed equipment and is ideally suited to environments where the Broadcast encryption technology is proprietary.

Some vendors also support UDP output encapsulation instead of R-DEPI to allow for pre-encrypted streams to be provided to an integrated CCAP Core acting as both control plane and R-DEPI data plane.

This option is deployed widely in North America for broadcast video with operators who have a mismatch between their CCAP Core vendor and their CAS vendor.

**Table 5 – QAM Input Video Traffic Engine Capabilities**

<b>Attribute</b>	<b>Capability</b>
Encryption	Good – encryption handled by existing edge QAM and just passed through
Mixed Vendor	Good – maintains configuration control in a single entity for flexibility when deploying a mix of DOCSIS and video solutions/vendors; supports existing edge QAM vendor and video backend integration
Network Topology	Good - can be located exactly where the video is located today separate from the DOCSIS Core
All-IP Transition	Neutral – maintains existing equipment as a transition instead of new investment in QAM video
Organizational	Neutral – video can be kept separate from DOCSIS if there is a separate Core for remote node configuration but UDP to CCAP Core is common
Typical Usage Scenarios	<ol style="list-style-type: none"> <li>1. Deployments where DOCSIS Core isn't the same as CAS vendor</li> <li>2. Deployments where leveraging existing edge QAMs helps with DAA migration or to avoid new test cycles to integrate new edge QAM functions into an integrated CCAP.</li> <li>3. Deployments where minimizing spend in QAM video over DAA infrastructure is critical</li> <li>4. Support of virtualized DOCSIS Core deployment with separate video solution</li> </ol>

## 5.2. UDP IP Input

Video Traffic Engines using UDP IP Input, center of Figure 11, support SPTS or MPTS from a point further back in the QAM video processing pipeline. A minimal implementation focuses on R-DEPI encapsulation of pre-encrypted and pre-multiplexed MPTS which may be available from a network encryptor or other broadly deployed multiplexing platforms. A maximal implementation provides full edge QAM functionality in the Traffic Engine, possibly as a fully virtualized software instance since no hardware elements are required to generate RF signals.

UDP IP Input Video Traffic Engines are well suited to operators who have high-density narrowcast content, typically due to the use of SDV. In this situation, a QAM Input Video Traffic Engine, as discussed in Section 5.1, requires significant space and RF plumbing to connect to existing edge QAMs. Deployments utilizing high-density narrowcast are commonly deployed with network encryptors, so moving further back in the video processing pipeline and connecting directly to the network encryptors can save significant space and power in the hub.

**Table 6 – UDP IP Input Video Traffic Engine Capabilities**

<b>Attribute</b>	<b>Capability</b>
Encryption	Neutral – good for DVB CAS systems but requires bulk encryptor solutions for proprietary CAS systems
Mixed Vendor	Neutral– may require another video backend integration cycle depending on how the Traffic Engine is integrated with existing video processing pipeline
Network Topology	Good - can be located exactly where the video is located today separate from the DOCSIS Core
All-IP Transition	Poor – doesn’t directly support next generation all-IP video delivery mechanisms
Organizational	Neutral – video can be kept separate from DOCSIS if there is a separate Core for remote node configuration
Typical Usage Scenarios	<ol style="list-style-type: none"><li>1. Deployments where there is high QAM count of narrowcast video (SDV for example) and bulk encryptor solutions are in place or planned</li><li>2. Deployments where keeping video separate from DOCSIS core is important (virtual DOCSIS core, mixed vendor DOCSIS environment)</li></ol>

### **5.3. CDN Input**

Next generation all-IP video delivery solutions utilize CDNs to cache content close to the customer for high quality-of-experience and they leverage distribution mechanisms such as MPEG-DASH to deliver video to clients. As cable operators move to all-IP video services available over consumer non-STB devices (such as tablets and streaming boxes), duplication of the video backend occurs as operators to serve both QAM STBs and newer IP devices.

There are significant operational expenditure benefits to moving to a common video backend based on their new CDN infrastructure investment. The transition to DAA offers an opportunity to integrate DAA delivery of QAM video by adding R-DEPI encapsulation functionality to CDN edge caches, right side of Figure 11, and avoiding deploying a new set of systems tied to the traditional QAM video backend.

**Table 7 – CDN Input Video Traffic Engine Capabilities**

<b>Attribute</b>	<b>Capability</b>
Encryption	Neutral – good for DVB CAS systems but requires bulk encryptor integration for proprietary CAS systems
Mixed Vendor	Neutral – may require another video backend integration cycle depending on how the Traffic Engine is integrated with existing video processing pipeline
Network Topology	Good - can be located exactly where the video is needed due to proximity/integration with edge caches and is separate from the DOCSIS Core
All-IP Transition	Good – directly supports next generation all-IP video delivery mechanisms
Organizational	Neutral – video can be kept separate from DOCSIS if there is a separate Core for remote node configuration
Typical Usage Scenarios	<ol style="list-style-type: none"> <li>1. Operators looking to move to a common modern CDN-based video backend to support all video services</li> </ol>

## Conclusion

The transition to DAA introduces many architecture options to maintain delivery of QAM video to existing STB deployments. The section discussed four video delivery options available to operators, ranging from highly integrated CCAP deployments to loosely coupled Cores and Traffic Engines. Deploying with Traffic Engines opens innovative ways to integrate video into existing infrastructure environments, including long-deployed edge QAM hardware and newly-minted CDN investments. Each architecture option has pros and cons, with no “right size fits all”. Operators have freedom and options to optimize QAM video delivery depending on their specific deployment needs. Thankfully each architecture presented can work in an interoperable and standards-based way with any DAA remote node deployments, allowing operators to deploy the solution that fits their need.

## Abbreviations

ABR	adaptive bit rate
CAS	conditional access system
CCAP	Converged Cable Access Platform
CDN	content delivery network
CM	cable modem
CMTS	cable modem termination system
DAA	distributed access architecture
DASH	dynamic adaptive streaming over HTTP
DOCSIS	Data Over Cable Service Interface Specifications
EQAM	edge QAM

ERM	edge resource manager
FMA	Flexible MAC Architecture
FPGA	field programmable gate array
GCP	generic control plane protocol
HFC	hybrid fiber-coax
HITS	headend-in-the-sky
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet protocol
L2TPv3	layer 2 tunneling protocol version 3
MHAv2	Modular Headend Architecture version 2
MPEG	Moving Pictures Expert Group
MPTS	MPEG transport stream
MSO	multi-system operator
OOB	out-of-band
PTP	precision time protocol
QAM	quadrature amplitude modulation
RF	radio frequency
R-DEPI	remote downstream external phy interface
R-DTI	remote DOCSIS timing interface
R-PHY	remote physical layer
SCTE	Society of Cable Telecommunications Engineers
SDV	switched digital video
SNR	signal-to-noise ratio
SoC	system-on-chip
SRM	session resource manager
STB	set-top box
TS	transport stream
UDP	user datagram protocol
VOD	video on demand

## Bibliography & References

[R-PHY] CableLabs Data-over-Cable Service Interface Specifications, DCA-MHAv2, Remote PHY Specification, CM-SP-R-PHY-I12-190307

[R-DEPI] CableLabs Data-over-Cable Service Interface Specifications, DCA-MHAv2, Remote Downstream External PHY Specification, CM-SP-R-DEPI-I12-190307

[Rahman] Remote PHY Architecture Options – Cores and Engines; Saifur Rahman, Comcast; SCTE Remote PHY Seminar 2017  
<https://www.scte.org/SCTEDocs/Expo/PHY/SCTE%20-%20RPHY%20Seminar%20-%20Comcast-%20cores%20and%20engines.pdf>



# **Offloading Data Using Unlicensed LTE (CBRS)**

A Technical Paper prepared for SCTE•ISBE by

**Loay Kreishan**  
Principal Mobile Engineer II  
Charter Communications  
6380 S Fiddler's Green Circle, Greenwood Village, CO 80111  
303-793-4489  
Loay.Kreishan@Charter.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Data Offloading for MVNOs .....	3
How Conventional DSDS Solutions are not Adequate .....	5
The Dynamic DSDS (D-DSDS) Mode.....	6
Initial Performance Experiments.....	8
Future Work.....	9
Conclusion .....	10
Abbreviations.....	10
Acknowledgements .....	11

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Small Cell Coverage .....	4
Figure 2 - Voice and Data networks .....	4
Figure 3 - PDN Status During Mobility Between MSO network and MNO Network.....	6
Figure 4 - Geolocation Function.....	7
Figure 5 - Initial Test Results .....	9
Figure 6 - Link Aggregation.....	10

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Network Thresholds .....	8

# Introduction

Multiple System Operators, MSOs, are capitalizing on new revenue streams by providing mobile services to their customer segment through mobile virtual network operators, MVNOs. MSOs are motivated to offload mobile data onto their own wireless infrastructure in order to reduce mobile data cost and improve user experience. MSOs can utilize Wi-Fi and LTE networks operating on spectrum bands like 3.5GHz Citizens Broadband Radio Service (CBRS), while keeping users registered with the MNO network to take advantage of broad coverage.

To this end, some MSOs are investigating using dual credential user equipment (UE). A UE with two subscriber identity modules (SIM) enables it to simultaneously stay attached to two networks. Dual SIM dual standby (DSDS) handsets are widely available in the market particularly in Asia.

This does not come without obstacles. The main challenge in the existing DSDS UE implementation is that it requires intervention on the part of the subscriber to designate one network for data and one for voice. The existing DSDS UE functions under the assumption that the subscriber has full control to select which networks will be suitable for data.

Because both networks' overlap, there is no need for the UE to know when it is within coverages and which network to use for data. Put simply, the devices don't need to be smarter in this area.

The standard DSDS implementation does not meet the MSO offloading requirements because the two networks' do not overlap.

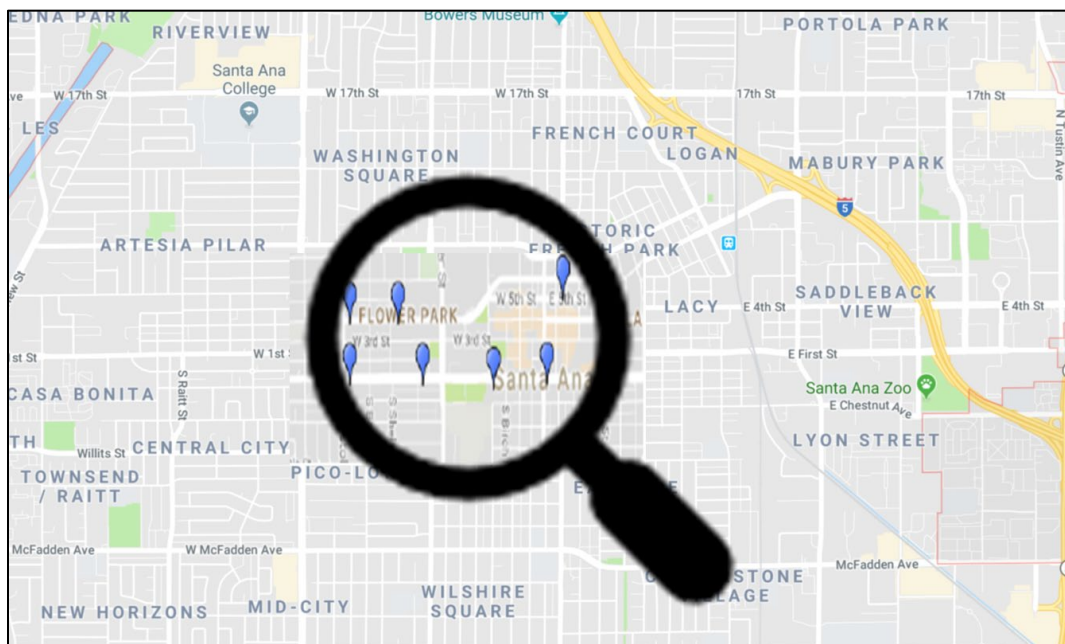
This paper first describes how the solution of using DSDS UE meets the main requirements of offloading data. The existing DSDS UE needs enhancements to seamlessly select a designated data network and dynamically switch data path between networks when available.

Then the paper depicts an improved DSDS solution, which is called dynamic DSDS (D-DSDS) in this paper, that lifts the burden of intervention from the subscriber and elevates his experience. Last, some promising D-DSDS results from early trials is shared that indicates a D-DSDS solution will soon become a reality.

## Data Offloading for MVNOs


For MSOs running MVNOs, one of the main sources of lowering data cost comes from offloading to less expensive, and sometimes free, sources of access networks like Wi-Fi. Along with the introduction of multiple licensed and unlicensed spectra like 3.5GHz CBRS spectrum came the opportunity to provide low cost mobile and fixed wireless broadband. MSOs are capitalizing on the implementation of these small cell CBRS networks to lower their MVNO data costs.

Early stages of deploying CBRS small cell networks will primarily be focused on densely populated, high traffic areas. This will leave the remaining uncovered areas, to be serviced by macro cells. Macro cells coverage can be provided by existing mobile network operators MNOs. The overall view of the coverage map will resemble a small spotted coverage area of CBRS networks (Figure 1). MVNOs using small cell will still depend on MNOs' network to provide complete, continuous coverage for their subscribers.



**Figure 1 - Small Cell Coverage**

There is no benefit for offloading voice services onto MSO network due to patchy nature of CBRS small cell network and voice type services are likely to be disrupted due to potential subscriber mobility. That is the MSO network will only support data services (D) for data offload only. The primary SIM network associated to the MNO will be used for voice and data services (V+D) and will be the only source for providing voice service (Figure 2).



**Figure 2 - Voice and Data networks**

This requires an active UE that can connect to multiple networks and implementing policy for selecting an offloading network. Connecting to multiple networks raises the issues of session continuity and user experience. Since offloading to Wi-Fi, it has been a challenge to move between multiple independent networks while maintaining session connectivity and providing non-interrupted data services.

## **How Conventional DSDS Solutions are not Adequate**

To solve this challenge, some MSOs have considered a variety of solutions which include:

1. Network sharing
2. Roaming
3. Smart SIM
4. Standard DSDS handsets

Of these solutions that achieve network selection, they do not support the required offloading behavior. The main requirements of successful offloading include:

1. Supports multiple offload networks. Multiple networks (carrier profiles) that support data can be provisioned on the UE.
2. The UE will select available network based on priority and session transition threshold provided by the carrier.
3. Seamless and dynamic data plane transition without user intervention.
4. Maintain user experience. The user's experience will be the same or better than with a single SIM device.
5. MSO offload networks will be used to support only data services and not voice services. Voice services shall use the primary network for voice.
6. Data offload shall not depend on integration between MSO and MNO networks, carriers' configuration, or user equipment configuration through a user interface.

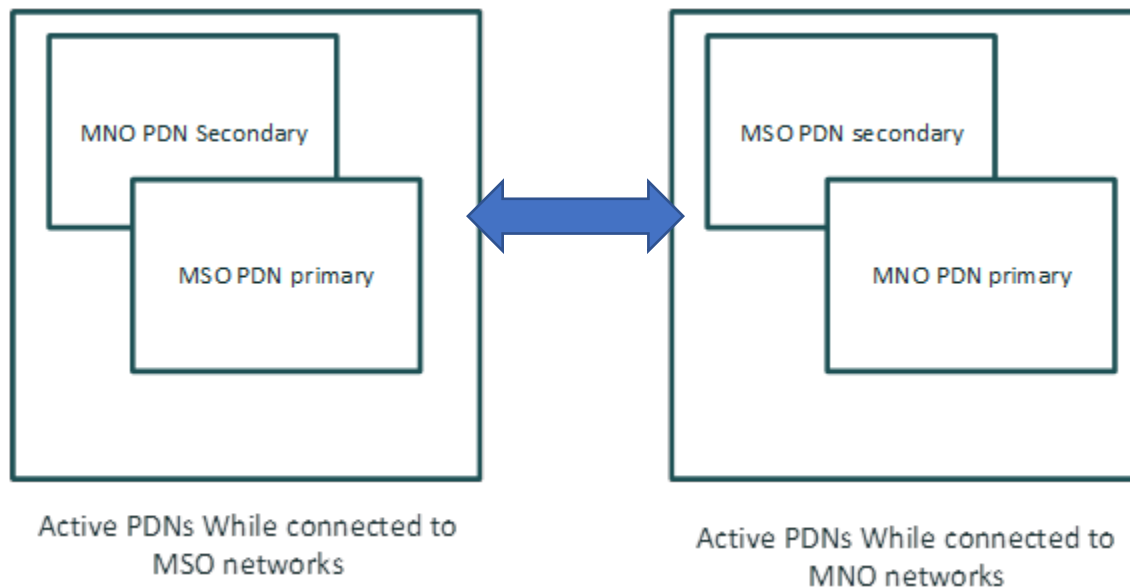
The available solutions fall short of meeting the requirements. The requirements mentioned above for offloading are not met for the following reasons:

1. Network sharing and roaming require support from the MNO to enable certain interfaces between the two networks. This will limit the capability of carrier, MVNO or MSO to dynamically select the desired network based on location.
2. Additionally, roaming requires the carrier to have all services available including voice infrastructure which will increase the cost of the offload network.
3. Solutions stated above assume that either both provisioned networks fully overlap, as in the case of DSDS, or they do not overlap at all, as in the case of roaming. In the case of offload, it may include both scenarios as depicted in Figure 1.
4. In the case of existing DSDS implementation, the user is requested to select the primary and secondary networks for voice and data. The selection is fixed and requires user intervention to change it. User intervention limits the carrier from selecting multiple offload networks.
5. The selection of networks in DSDS is static based on user choices. Changing the default network for data because of availability, quality or cost is not possible.
6. In the case of existing DSDS, the selected primary data network will always support voice services. This scenario may not be desirable if the secondary network used for offloading is out of reach.

The standard DSDS device/UE capabilities in the market today meet most but not all the offloading requirements. There is always room for improvement. To this end, Spectrum Mobile's engineering team has worked with handset OEMs to provide an enhanced version of the DSDS that more comprehensively meets all offload requirements. We will reference this advancement of DSDS as Dynamic DSDS (D-DSDS).

## The Dynamic DSDS (D-DSDS) Mode

D-DSDS UE takes advantage of chipset features that provide dynamic switching. In this case each network will have a registered packet data network (PDN), but only one PDN will be used for data, namely the primary network, at a given time as depicted in Figure 3

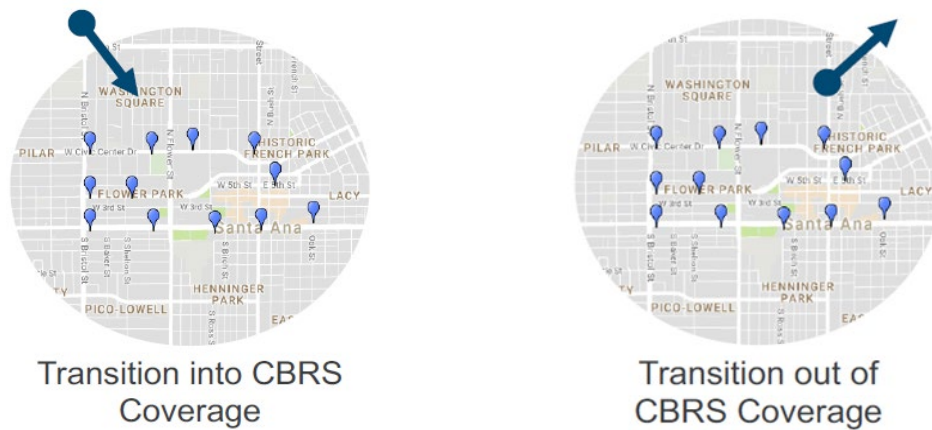


**Figure 3 - PDN Status During Mobility Between MSO network and MNO Network**

When the UE transitions to the MSO network coverage, it will continue tuning to the MNO network to listen to its paging window. If there is a page destined for the UE, it will check the invite message to determine if there are voice services (e.g. MMS, SMS or VoLTE) session or data session. The UE will switch to the MNO network only if there is a voice service session waiting on the MNO network. At that time, the active PDN will be the MSO PDN offload network. If the invite message has a data session on the MNO network, and the UE is not idle on the MSO networks, then the UE will not switch to the MNO network and will keep the MSO PDN as the active PDN.

Another enhancement is to enable UE to seamlessly switch between primary and secondary networks. While the D-DSDS UE follows the 3GPP technical specifications for handover and mobility within a network, it has intelligence and capability to select and connect to the appropriate data network based on location and network selection thresholds provided by the MSO carriers.

Location will provide the UE with information about possibility of available networks and when to start the process of searching and connecting to the MSO network. A geolocation function, part of the connection manager, will process the information provided to the handset and select the best network for the current location and the appropriate time to transition between networks.



**Figure 4 - Geolocation Function**

A major benefit of the geolocation function is to improve the percentage of offloading by minimizing the scanning interval time. Minimizing power drainage from unnecessary scanning and associating to multiple small cell networks means resources are utilized more efficiently.

Network transitioning thresholds are an important factor as well. The standard DSDS has network-independent transitioning thresholds provided by each network and used for events related to mobility for that network. This is true for the D-DSDS UE, but there are additional set of network thresholds that facilitate seamless transitioning between networks. These thresholds are used to decide when and which network to use for data offloading.

The following table lists available network thresholds used by the UE to determine which network to connect to and designating the primary PDN for initiating data session transition. Key threshold values include reference signal received power (RSRP) and average time to trigger (TTT).

**Table 1 – Network Thresholds**

Threshold	Description
DATA_HANDIN_CONNECT ED_RSRP	Average power value of sector/site
DATA_HANDIN_CONNECT ED_RSRQ	Average signal Received Quality sector/site
DATA_HAND OUT_IDLE_RSRP	Average power value of sector/site
DATA_HAND OUT_IDLE_RSRQ	Average signal Received Quality sector/site
DATA_HAND OUT_CONNECTED_RSRP	Average power value of sector/site
DATA_HAND OUT_CONNECTED_RSRQ	Average signal Received Quality sector/site
Average BLERBER/Frequency (MSO, MNO)	Average Block error and bit error rates when connected to MSO and MNO network
Average Time to Trigger (TTT) value	Time duration required for power values of MSO should be higher than power values of MNO (depends on the direction of handoff)

Unlike the standard DSDS, the D-DSDS handset does not require any user intervention to select and connect to the desired offload network. To the user, the device is still connected to the carrier or service provider network, and more importantly, data services are available and not disrupted. To this extent the user does not need to be concerned with the offload network.

Moreover, the D-DSDS handset will not render any information regarding the offload network. This includes a dedicated received signal strength indicator known as signal bars, network name, network selection option and configuration information of the offload network. This is because the offload network will be totally seamless to the subscriber. The subscriber doesn't need to be concerned with the offload network.

In the case with voice services, the UE will transition to the MNO network to establish the voice session, simultaneously switching the active PDN to MNO PDN. When the voice service is terminated, the UE will revert to the offload network if location and key performance indicators for transition are met. For SMS/MMS, the MNO network will provide these services. The handset will not switch the PDN to the MNO network when receiving invite for SMS services. The UE will be able to receive the SMS content without switching.

## Initial Performance Experiments

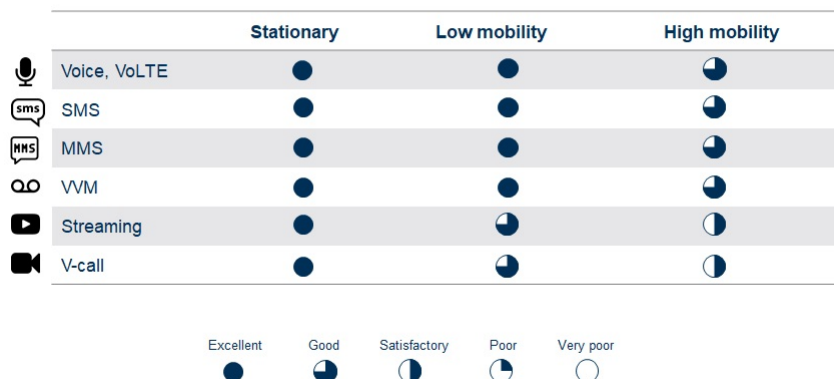
The first version of the D-DSDS has been used in trials and shows promising results. Vital services and applications were tested implementing D-DSDS during mobility between networks. Among the services tested were voice calls and video streaming. As to be expected, all services performed excellently in stationary, low mobility and performed less than expected during high mobility between networks due to delay. In the cases of high mobility, we see some impact, but results are still within the good to satisfactory range.

Voice: results for testing voice were very good at low and high mobility Video streaming: results for testing video streaming were not affected due to buffering. Although service was not interrupted during high mobility, the network transition was noticeable to the subscriber during high mobility.



SMS/MMS: results were very good during mobility and subscriber did not notice any change in service.

### Device feature performance under MNO + MSO coverage



**Figure 5 - Initial Test Results**

The degradation in service during low and high mobility was due to network transitioning time to switch between PDNs. Although we did not notice any interruption with service, the transition was affecting the service during high mobility.

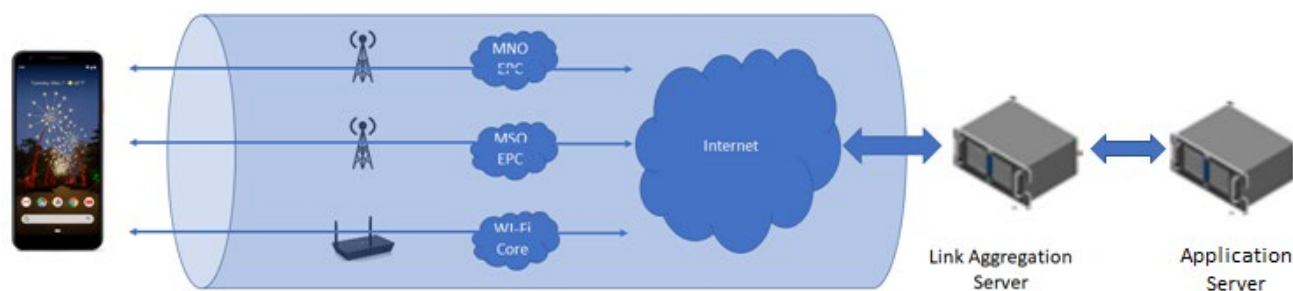
We think a session continuity solution will reserve the service and allow the running applications to be isolated from the effect of network transitioning.

## Future Work

Session continuity and user experience are important performance measurements for successful data offloading. To enhance session continuity and improve user experience a multipath solution like link aggregation will be introduced to the D-DSDS. This important enhancement will allow to unify the user experience when transitioning between available networks MNO, MSO and Wi-Fi.

Link aggregation will utilize a virtual tunnel to provide a single IP address and a virtual gateway IP address for all layer 7 applications, regardless of the network or medium used. Applications will use the virtual IP and the virtual gateway to reach the internet.

Link aggregation will maintain session continuity by providing alternative paths for packets when one path is blocked. It also can provide redundant packet transmission through multiple paths.



**Figure 6 - Link Aggregation**

## Conclusion

MNOs and MVNOs everywhere are trying to overcome the ever-increasing demand for data while lowering the data cost by providing multiple options for offloading. Current DS/SDS solutions are missing the mark, and a more dynamic DS/SDS is on the horizon to enhance connectivity while slashing the costs. Still in its early stages, D-DS/SDS handsets introduce promising features that are valuable to operators providing mobile data services, private LTE networks or offloading data services alike. During limited trials, most test results have exceeded expectations.

We anticipate more work to be completed on the D-DS/SDS handset which includes improvements and new features to be added soon.

## Abbreviations

3GPP	Third Generation Partnership Project
CBRS	Citizen Broadband Radio Service
D-DS/SDS	Dynamic Dual SIM Dual Standby
DDS	dynamic data switching
DS/SDS	Dual SIM Dual Standby
GSMA	Global System for Mobile Communications
MNO	Mobile Network Operator
MSO	Multiple System Operator
MVNO	Mobile Virtual Network Operator
OEM	Original Equipment Manufacturer
PDN	Packet Data Network
RSRP	reference signal received power
RSSI	received signal strength indicator
SIM	Subscriber Identity Module
TTT	time to trigger
UE	user equipment

# Acknowledgements

Ahmed Bencheikh, GVP, Wireless Engineering (Charter)

Will Logan, VP, Wireless Engineering Mobile Engineering & Technology (Charter)

Marcus Greenwood, Dir, Software Development Mobile Engineering & Technology (Charter)

# Characterizing Network Problems Using DOCSIS® 3.1 OFDM RxMER Per Subcarrier Data

A Technical Paper prepared for SCTE•ISBE by

**Ron Hranac**

Technical Marketing Engineer  
Cisco Systems  
9155 E. Nichols Ave., Ste. 400, Centennial, CO 80112  
720-875-1338  
rhranacj@cisco.com

**James Medlock**

Founder and CEO  
Akleza, Inc.  
18695 Pony Express Dr. 2350, Parker, CO 80134  
303-670-7951  
jmedlock@akleza.com

and

**Bruce Currivan**, JJP Development

**Roger Fish**, Broadcom

**Tom Kolze**, Broadcom

**Jason Rupe**, CableLabs®

**Tom Williams**, CableLabs®

**Larry Wolcott**, Comcast

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	6
Content .....	6
1. What is MER? .....	6
1.1. SC-QAM RxMER .....	8
1.2. OFDM RxMER Per Subcarrier .....	9
1.2.1. Description of RxMER per Subcarrier Measurement .....	11
2. Impairment Identification .....	13
2.1.1. Simulated Ingress .....	14
2.1.2. Inverted Plot and Equalized Noise Floor .....	14
2.1.3. Impairment Examples .....	15
3. Lab Testing .....	17
3.1. Amplitude Ripple .....	18
3.1.1. Amplitude Ripple Test Results Discussion .....	20
3.2. RxMER Edge Rolloff .....	23
3.2.1. RxMER Edge Rolloff Discussion .....	25
4. Using RxMER per Subcarrier Data for Maintenance and Troubleshooting .....	26
5. Areas for Further Investigation .....	30
Conclusion .....	31
Abbreviations .....	31
Bibliography & References .....	32
Appendix .....	33
6. Appendix I – Lab Test Results .....	33
6.1. Test Case 1A .....	33
6.2. Test Case 1B .....	35
6.3. Test Case 1C .....	36
6.4. Test Case 2A .....	37
6.5. Test Case 2B .....	39
6.6. Test Case 2C .....	40
6.7. Test Case 3A .....	41
6.8. Test Case 3B .....	43
6.9. Test Case 3C .....	44
6.10. Test Case 4A .....	45
6.11. Test Case 4B .....	47
6.12. Test Case 4C .....	48
6.13. Test Case 5A .....	49
6.14. Test Case 5B .....	51
6.15. Test Case 5C .....	52
7. Appendix II – Acknowledgements .....	53

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1. MER is the ratio of average symbol power to average error power. ....	7

Figure 2. The symbol points in the left constellation are tightly grouped, indicating higher MER (27 dB). The symbol points in the right constellation are diffuse (spread out), indicating lower MER (17.5 dB).....	8
Figure 3. Example test equipment screen shot showing 38.3 dB RxMER for an SC-QAM signal on CTA channel 19. ....	9
Figure 4. RxMER per subcarrier plot for a 96 MHz-wide OFDM signal.....	11
Figure 5. OFDM RxMER computation.....	13
Figure 6. RxMER per subcarrier graph for a 96 MHz-wide OFDM signal. ....	14
Figure 7. OFDM RxMER per subcarrier plot showing simulated ingress and its impact on the affected subcarriers. ....	14
Figure 8. Inverted plot of RxMER per subcarrier.....	15
Figure 9. OFDM RxMER per subcarrier graph showing ingress interference at about 890 MHz. ....	16
Figure 10. "Edge rolloff" visible at the left edge of the RxMER per subcarrier plot. ....	16
Figure 11. Multiple impairments are evident in this RxMER per subcarrier graph (see text), along with an exclusion band from 759.85 MHz to 763.55 MHz. ....	17
Figure 12. RxMER per subcarrier graph with amplitude ripple caused by an impedance mismatch-related reflection. ....	17
Figure 13. Equipment configurations for RxMER per subcarrier amplitude ripple testing.....	18
Figure 14. Test Case 4B spectrum analyzer capture (the frequency domain ripple in the channel was the same for Test Case 3B).....	19
Figure 15. Test Case 3B RxMER per subcarrier graph showing little or no amplitude ripple.....	20
Figure 16. Test Case 4B RxMER per subcarrier graph showing amplitude ripple.....	20
Figure 17. Test Case 4B OFDM channel power (the OFDM channel power was the same for Test Case 3B).....	20
Figure 18. OFDM signal test system block diagram.....	21
Figure 19. Test conditions with Noise 1 injected before the source of an echo (reflection). As seen at observation point (B) both Noise 1 and the OFDM signal have amplitude ripple, but the modem's adaptive equalizer removes the channel effect so that at observation point (C) the RxMER-per-subcarrier plot does not show ripple. In this figure, the horizontal axis is frequency, and the vertical axis is relative amplitude in decibels. ....	22
Figure 20. Test conditions where Noise 2 is injected after the source of an echo (reflection). At observation point (B), Noise 2 does not have ripple because it did not pass through the echo. However, the modem's RxMER-per-subcarrier plot does have amplitude ripple (C) because the noise passed through the receive equalizer H(f)-1. In this figure, the horizontal axis is frequency, and the vertical axis is relative amplitude in decibels. ....	23
Figure 21. Another example of edge rolloff in an RxMER per subcarrier graph. ....	23
Figure 22. Example spectrum capture of an SC-QAM signal adjacent to an OFDM signal. Note the OFDM signal's taper region, which extends into part of the SC-QAM signal.....	24
Figure 23. RxMER per subcarrier graph after reconfiguring $N_{cp}$ and $N_{rp}$ .....	24
Figure 24. Table 75 from Appendix V of the DOCSIS 3.1 Physical Layer Specification [1]. ....	25
Figure 25. Amplitude ripple-to-echo tunnel length calculation. ....	26
Figure 26. System map with RxMER per subcarrier data.....	27
Figure 27. Close-up of RxMER per subcarrier amplitude ripple periodicity.....	28
Figure 28. Technician troubleshooting RxMER per subcarrier amplitude ripple.....	29
Figure 29. Trunk splice removed from network.....	29
Figure 30. RxMER per subcarrier after replacing splice.....	30
Figure 31. Test Case 1A spectrum analyzer screen capture.....	33

Figure 32. Test Case 1A channel estimate.....	34
Figure 33. Test Case 1A RxMER per subcarrier.....	34
Figure 34. Test Case 1A OFDM channel power (nominal -10 dBmV). ....	34
Figure 35. Test Case 1B spectrum analyzer screen capture.....	35
Figure 36. Test Case 1B channel estimate.....	35
Figure 37. Test Case 1B RxMER per subcarrier.....	35
Figure 38. Test Case 1B OFDM channel power (nominal 0 dBmV). ....	36
Figure 39. Test Case 1C spectrum analyzer screen capture.....	36
Figure 40. Test Case 1C channel estimate.....	36
Figure 41. Test Case 1C RxMER per subcarrier. ....	37
Figure 42. Test Case 1C OFDM channel power (nominal +10 dBmV). ....	37
Figure 43. Test Case 2A spectrum analyzer screen capture.....	37
Figure 44. Test Case 2A channel estimate.....	38
Figure 45. Test Case 2A RxMER per subcarrier.....	38
Figure 46. Test Case 2A OFDM channel power (nominal -10 dBmV). ....	38
Figure 47. Test Case 2B spectrum analyzer screen capture.....	39
Figure 48. Test Case 2B channel estimate.....	39
Figure 49. Test Case 2B RxMER per subcarrier.....	39
Figure 50. Test Case 2B OFDM channel power (nominal 0 dBmV). ....	40
Figure 51. Test Case 2C spectrum analyzer screen capture.....	40
Figure 52. Test Case 2C channel estimate.....	40
Figure 53. Test Case 2C RxMER per subcarrier. ....	41
Figure 54. Test Case 2C OFDM channel power (nominal +10 dBmV). ....	41
Figure 55. Test Case 3A spectrum analyzer screen capture.....	41
Figure 56. Test Case 3A channel estimate.....	42
Figure 57. Test Case 3A RxMER per subcarrier.....	42
Figure 58. Test Case 3A OFDM channel power (nominal -10 dBmV). ....	42
Figure 59. Test Case 3B spectrum analyzer screen capture.....	43
Figure 60. Test Case 3B channel estimate.....	43
Figure 61. Test Case 3B RxMER per subcarrier.....	43
Figure 62. Test Case 3B OFDM channel power (nominal 0 dBmV). ....	44
Figure 63. Test Case 3C spectrum analyzer screen capture.....	44
Figure 64. Test Case 3C channel estimate.....	44
Figure 65. Test Case 3C RxMER per subcarrier. ....	45
Figure 66. Test Case 3C OFDM channel power (nominal +10 dBmV). ....	45
Figure 67. Test Case 4A spectrum analyzer screen capture.....	45
Figure 68. Test Case 4A channel estimate.....	46
Figure 69. Test Case 4A RxMER per subcarrier.....	46
Figure 70. Test Case 4A OFDM channel power (nominal -10 dBmV). ....	46
Figure 71. Test Case 4B spectrum analyzer screen capture.....	47
Figure 72. Test Case 4B channel estimate.....	47
Figure 73. Test Case 4B RxMER per subcarrier.....	47
Figure 74. Test Case 4B OFDM channel power (nominal 0 dBmV). ....	48
Figure 75. Test Case 4C spectrum analyzer screen capture.....	48

Figure 76. Test Case 4C channel estimate.....	48
Figure 77. Test Case 4C RxMER per subcarrier. ....	49
Figure 78. Test Case 4C OFDM channel power (nominal +10 dBmV). ....	49
Figure 79. Test Case 5A spectrum analyzer screen capture.....	49
Figure 80. Test Case 5A channel estimate.....	50
Figure 81. Test Case 5A RxMER per subcarrier.....	50
Figure 82. Test Case 5A OFDM channel power (nominal -10 dBmV). ....	50
Figure 83. Test Case 5B spectrum analyzer screen capture.....	51
Figure 84. Test Case 5B channel estimate.....	51
Figure 85. Test Case 5B RxMER per subcarrier.....	51
Figure 86. Test Case 5B OFDM channel power (nominal 0 dBmV). ....	52
Figure 87. Test Case 5C spectrum analyzer capture.....	52
Figure 88. Test Case 5C channel estimate.....	52
Figure 89. Test Case 5C RxMER per subcarrier. ....	53
Figure 90. Test Case 5C OFDM channel power (nominal +10 dBmV). ....	53

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Example tabular list of OFDM per-subcarrier RxMER values reported by a DOCSIS 3.1 cable modem. The original list has been shortened.....	10



# Introduction

Receive modulation error ratio (RxMER) has long been a powerful metric for cable network maintenance and troubleshooting. A limitation to single carrier quadrature amplitude modulation (SC-QAM) RxMER is that the reported value doesn't give an indication of why that value is what it is, or what kind of impairment might exist.

The DOCSIS® 3.1 specifications [1] define several operational measurements that can be reported by the cable modem and cable modem termination system (CMTS) or converged cable access platform (CCAP). One important modem performance parameter is orthogonal frequency division multiplexing (OFDM) RxMER per subcarrier, which can be plotted to show a graph of all subcarriers' RxMER performance. Based on real-world observations of data from production cable networks and subsequent lab testing to recreate and validate the observations, a number of specific impairments can be identified that point to faults in the underlying network. Not only does the identification of these problems assist with maintenance and troubleshooting of the network, but various impairments identifiable in the RxMER per subcarrier plots can impact subscriber service and result in lower throughput and performance than expected. Plus, due to the sensitivity of the RxMER per subcarrier measurement, it can find impairments in the network before they adversely impact customer service, and before repairs become costly.

This paper includes discussions about a number of impairments that have been observed, describes the findings when recreated in a laboratory environment, and explains how the observed results point to potential cable network faults. Examples include:

- Amplitude ripple in the channel in the frequency domain can under certain conditions cause amplitude ripple in RxMER per subcarrier graphs.
- Interference caused by long term evolution (LTE) and other ingress can be correlated to specific frequencies by observing the impact on the RxMER per subcarrier graphs.
- SC-QAM signals adjacent to OFDM signals can cause a rolloff at the edges of the RxMER per subcarrier graph.

Production network examples are presented that show how analysis of RxMER data collected from cable modems can be used to identify and locate specific cable network impairments, resulting in an improved subscriber service performance and experience.

## Content


### 1. What is MER?

For a single QAM carrier or a single OFDM subcarrier, modulation error ratio (MER) is the ratio of average signal constellation power to average constellation error power – that is, digital complex baseband signal-to-noise ratio (SNR).<sup>1</sup> Indeed, MER is often called SNR. From a high-level perspective, the following formula defines MER (refer to Figure 1):

$$\text{MER} = 10\log_{10}(\text{average symbol power}/\text{average error power})$$

---

<sup>1</sup> For more information about MER, see [2], [3], and [6] through [10] in the bibliography.



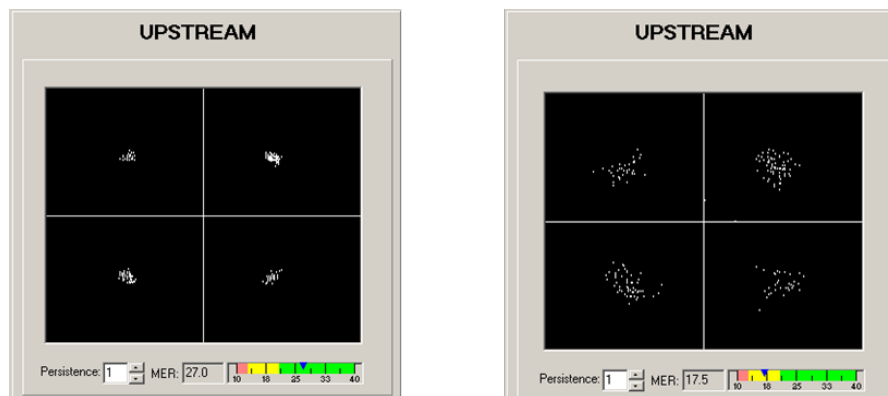
**Figure 1. MER is the ratio of average symbol power to average error power.**

A more precise mathematical definition of MER is

$$MER = 10 \log_{10} \left[ \frac{\sum_{j=1}^N (I_j^2 + Q_j^2)}{\sum_{j=1}^N (\delta I_j^2 + \delta Q_j^2)} \right]$$

where  $I$  and  $Q$  are the real (in-phase) and imaginary (quadrature) parts of each sampled ideal target symbol vector, and  $\delta I$  and  $\delta Q$  are the real (in-phase) and imaginary (quadrature) parts of each modulation error vector. This definition assumes that a long enough sample is taken so that all the constellation symbols are equally likely to occur. Note: The numerator in the above equation can be replaced with a constant if the constellation power is known, as is the case in DOCSIS 3.1 OFDM RxMER where all constellations have average power = 1.

In effect, MER is a measure of how “fuzzy” or spread out the symbol points in a constellation are. For example, Figure 2 shows two quadrature phase shift keying (QPSK) data constellations, one with high MER (left), the other with low MER (right).



**Figure 2. The symbol points in the left constellation are tightly grouped, indicating higher MER (27 dB). The symbol points in the right constellation are diffuse (spread out), indicating lower MER (17.5 dB).**

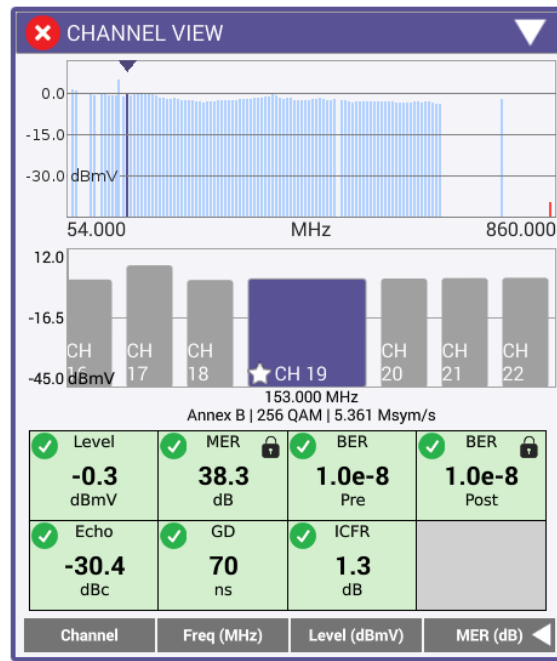
RxMER is the MER as measured in a digital receiver after demodulation of a given QAM carrier or OFDM subcarrier, after adaptive equalization.

RxMER is affected by the carrier-to-noise ratio (CNR); phase noise in the transmitter or receiver; linear distortions such as micro-reflections, amplitude ripple, and group delay; non-linear distortions such as composite triple beat, composite second order, and common path distortion; in-channel ingress; laser clipping; and just about anything else that degrades the channel through which the signal is transmitted. Its usefulness lies in the fact that it is a bottom-line measurement at the receiver slicer, just before forward error correction (FEC) decoding. An RxMER computation based on blind slicer decisions would produce inaccurate results at low SNR due to slicer symbol errors. This would be particularly important with the strong low density parity check (LDPC) coding used in DOCSIS 3.1 OFDM, which allows the link to operate at low SNR values where slicer errors are a normal occurrence, and are corrected by the FEC decoder. As we will see, the DOCSIS 3.1 OFDM RxMER per subcarrier metric uses the pilot subcarriers, which have known modulation values, so no slicer errors occur and the measurement is accurate over a wide dynamic range.

### 1.1. SC-QAM RxMER

When a QAM receiver in a set-top box, cable modem, CMTS upstream burst receiver, or a test instrument computes RxMER for an SC-QAM signal, the value reported is for just that signal – for instance, a 6 MHz-wide downstream DOCSIS signal. Figure 3 shows an example in which the reported RxMER for an SC-QAM signal on CTA<sup>2</sup> channel 19 is 38.3 dB. A detailed explanation of how a QAM receiver computes RxMER can be found in [2].

<sup>2</sup> The Consumer Technology Association's CTA-542-D R-2018 standard [4] defines channel plans and frequencies used for 6 MHz-wide channels in cable networks.



**Figure 3. Example test equipment screen shot showing 38.3 dB RxMER for an SC-QAM signal on CTA channel 19.**

While SC-QAM RxMER is a useful tool for characterizing the health of the signal and/or the network, the reported value doesn't give an indication of why the value is what it is. If the reported RxMER is low, one cannot determine from just the RxMER value what kind of impairment(s) might exist.

## 1.2. OFDM RxMER Per Subcarrier

RxMER is even more useful with DOCSIS 3.1 OFDM signals, because an OFDM signal comprises up to several thousand subcarriers, each of which is a narrow-bandwidth QAM signal with its own RxMER measurement value. However, trying to manage a list of RxMER values for thousands of subcarriers would be unwieldy and impractical. Consider the tabular list in Table 1, from an operational DOCSIS 3.1 cable modem. The far left column is the subcarrier number in hexadecimal notation (hex). The 8-bit hex values to the right are RxMER values in 1/4 dB increments, (two digits represent RxMER for a subcarrier). The zeros at the start and end are nulled for the excluded subcarriers including the taper regions. In practice a list of RxMER values per subcarrier like this would have to be converted from hex to decimal (in dB) to be useful (e.g., subcarrier #1 RxMER = 41.25 dB, subcarrier #2 RxMER = 41.5 dB, subcarrier #3 RxMER = 41.75 dB...subcarrier #7600 RxMER = 41.25 dB, etc.).

**Table 1 - Example tabular list of OFDM per-subcarrier RxMER values reported by a DOCSIS 3.1 cable modem. The original list has been shortened.**

Number of SubCarriers	: 8192
1st Active SubCarrier	: 296
# of Active SubCarriers	: 7600
Tx Time	: 0h:04m:56s ago
Rx Time	: 0h:04m:55s ago
OFDM Profile Failure Rx	: 172h:26m:55s ago
MER Poll Period (min)	: 5
Recommend Timeout (min)	: 120
Unfit Timeout (min)	: 5
Source	: OPT
Sub- RxMER	
Carrier	
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
0x0080	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x00A0	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x00C0	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x00E0	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0100	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0120	00000000 00000000 A5A6A7A6 A7A6A8A6 A4A8A7A6 AAA4A6A3 A6A5A6A8 A8A8A4A7
0x0140	A7ABA6A7 A4A1A6A3 A8A6A6A7 A6A3A5A7 A7A4A6A3 A5A3A7A3 A3A6A3A7 A4A5A8A2
0x0160	A49FA3A6 A7A5A3A7 A8A8A4A4 A4A5A5A7 A7A6A5A7 A79FA7A1 A3A6A5A9 A6A9A5A4
0x0180	A7ABA5A7 A3A6A8A2 A5A7A9A8 A1A8A5A4 A6A5A2AC A7A6A2A5 A7A6A3A5 A4A5A5A6
0x01A0	A8A4A3A6 A3A6A5A7 A5A6A7A4 A8A8A8A6 A8A2A7A4 A8A4A3A5 A6A8A5A8 A4A3A3A2
0x01C0	A6A3A3A5 A7A2A3A3 A6AAA3A4 A7A9A5A5 A6A3A3A7 A5A4A1A8 A3A7A1A8 A7A4A6A6
0x01E0	A0A1A5A5 9FA7A7A5 A7A5A6A3 A5A6A3A8 A4A5A4A4 A7A2A0A3 A1A7A6A5 A7A6A4A7
0x0200	A5A8A5A2 A5A4A7A6 A6A7A5A7 A5A59FAA A6A6A5A3 A7A4A1A5 A5A6A2A6 A2A3A5A4
0x0220	A2A7A3A2 A8A5ABA3 A7A8A4AA A4A4A6A4 A8A3A1A5 A3A6A4A6 9FA7A5A6 AAA4A7A2
0x0240	A5A3A3A6 A5A4A9A2 A7A5A6A6 A8A7A2A5 A2A7A6A6 AAA7A7A6 A5A9A4A2 A7A8A4A5
0x0260	A6AAA5A5 A4A6A9A5 AAA3A7A4 A6A1A8A3 A4A4A8A7 A7A5A4A3 A6A7A8A9 A5A6A4A6
0x0280	A3A4A4A1 A7A4A7A6 A9A5A6A6 A3A2A4A6 A2A7A7A4 ABA5A3AB A2A7A3A4 A5A4A7A4
0x02A0	A3A1A3A5 A3A7A7A0 A7A6A5A5 A7A2A5A8 A7A4A5A5 A9A9A5A4 A4A7A2A6 A4A2A6A2
0x02C0	A4AAA6A4 A0A4AA6 A3A6A6A7 A3AAA4A5 A6A3A8A6 A6A3A4AB A9A2AAA6 A6A5A5A4
0x02E0	A9A5A6A3 A9A4A8AA A6A4A7A5 A8A5A0A6 A4A5A6AA A1A2A5A6 A9A5A3A8 A8A4A3A5
<data deleted>	
0x1D00	A4A5A5AB A8AAA5AB A4A5A3A8 A6A9A6A6 A7A9ABA6 A7A8A4A5 ABA6A8A9 A7A6A6A4
0x1D20	AAA5A7A9 A5A9A6A7 A8A7A8A2 A5AAA9A7 A8AAAAA8 A6A5A5AA A4AAA6A6 A6A6A8A7
0x1D40	A8A7A8A7 A5A9A5A4 A5A69BA9 AAA9A7A4 ACA9A8A7 A6A5A7A9 A4A9A9AB A5A7A7A5
0x1D60	ABA7A4A5 A6A4AE4A A8A9A3A6 A3A4A9AA A6A8A9A8 AAA6A8A9 A9A5A4A7 A8A6A7A8
0x1D80	AAACA9A6 A6A6A6A6 A3A4A6A7 A5A9A5A8 AAA4AA9 A5A6A6A7 A8A7AAA7 A9A7A8A8
0x1DA0	A5A3A8A6 A7A7A7A7 AA6A6A9 A5AAA5A5 A8A7A7A6 A9A7A3AA ACA7A8AA A7A5A9A7
0x1DC0	A9A5ABA5 A7A6A8A6 A6A9ABA8 A7A6A6A7 AA5A5A6B A5A5A8A6 A9A5ACAA A6A6A6A3
0x1DE0	A6A2A39F A7A7A9A8 A6A5A8A8 A6A5A7A7 A9A5A9A9 A7A6A7A8 A3A8A5A4 A4A9A7AA
0x1E00	A7ABA5AA A7AAA8A7 A7A7A7A9 A5A8A7A7 A5A6A7A6 A6A7A8AA A7A5A8A7 A6A3A8AA
0x1E20	A7A7A7A8 A4A8A6A9 A2A9A5A8 A6A4A6A7 A9A6A9A9 A6A7A5AC A8A4A7A6 A7A9A5AA
0x1E40	A9A5A7AA A7A9A3A8 A7A6A6A9 A8A7A4A8 A8A6A7AB AA5A5A8A6 AAAA6A6A A9A8A5A4
0x1E60	A9A9A8AA AAA4A5A3 A7A7A9A6 A7A4A5A3 A6A6A6A5 A8A6ABA8 AAA5A7A8 9DA7A7A7
0x1E80	A9A8A6A8 A5AA8A6 A6A7A8A6 A9A5AAA6 A6A8A4A8 A9A4A5AD A7A6A8A8 A8A9A7A8
0x1EA0	ACA9A7A7 A7A9A8A8 A7A5A8AA A5A3ADA8 A9A6A5A6 AAA6A6A7 A6A5A8AB ACA8A7A9
0x1EC0	A9AAA8A9 A6A7A7AA A5A7A8A7 A7AAA6A9 A7AAA9A7 A4ACA8A5 00000000 00000000
0x1EE0	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x1F00	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x1F20	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x1F40	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x1F60	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x1F80	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x1FA0	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x1FC0	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x1FE0	00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
SC RxMER Distribution (Excluded SCs counted as 0):	
Each *: 2%	
>44dB: 0.03%	
44dB: ** 4.63%	
43dB: ***** 45.32%	
42dB: ***** 44.50%	
41dB: ** 4.81%	
40dB: 0.59%	
39dB: 0.09%	
38dB:	
37dB:	
36dB:	
35dB:	
34dB:	
33dB:	
<33dB:	
-----100	
Percent of Subcarriers	

Instead of dealing with a cumbersome list of RxMER per subcarrier values, it is much more convenient to plot the per-subcarrier RxMER on a graph, showing frequency or subcarrier numbers in the horizontal axis, and RxMER in decibels in the vertical axis. Figure 4 illustrates an example. As will be shown later, plotting OFDM RxMER per subcarrier data on a graph can be used to identify and characterize a number of impairments.




Figure 4. RxMER per subcarrier plot for a 96 MHz-wide OFDM signal.

### 1.2.1. Description of RxMER per Subcarrier Measurement

One of the reasons the OFDM RxMER measurement is useful under a wide range of conditions is the way it is computed. As mentioned earlier, the measurement does not rely on the OFDM signal's data subcarriers, which with their strong LDPC decoding and potentially high constellation densities, are subject to slicer symbol errors and thus cannot provide a reliable measurement of RxMER when a given subcarrier has low SNR. Rather, the cable modem measures the RxMER using pilots and PHY link channel (PLC) preamble symbols, which have known values regardless of SNR.

In the DOCSIS 3.1 OFDM downstream, the scattered pilots scan across all active subcarriers, repeating the scan every 128 OFDM symbols. When the scattered pilots land on a continuous pilot or PLC preamble symbol, they adopt the value of the continuous pilot or PLC preamble symbol. We can use the name "scan pilots" to cover all three cases: scattered pilot, continuous pilot or PLC preamble symbol. The scattered pilots and continuous pilots are BPSK-modulated with real part =  $\pm 2$  and imaginary part = 0. Thus, they have power = 4, or  $10\log_{10}(4) = 6.02$  dB higher than the data subcarrier constellations, which have average power = 1, or  $10\log_{10}(1) = 0$  dB. The PLC preamble symbols are BPSK with real value  $\pm 1$  (and imaginary part = 0). Thus, they have power = 1, which is the same as the QAM data subcarrier constellations; that is, the PLC preamble is not boosted.

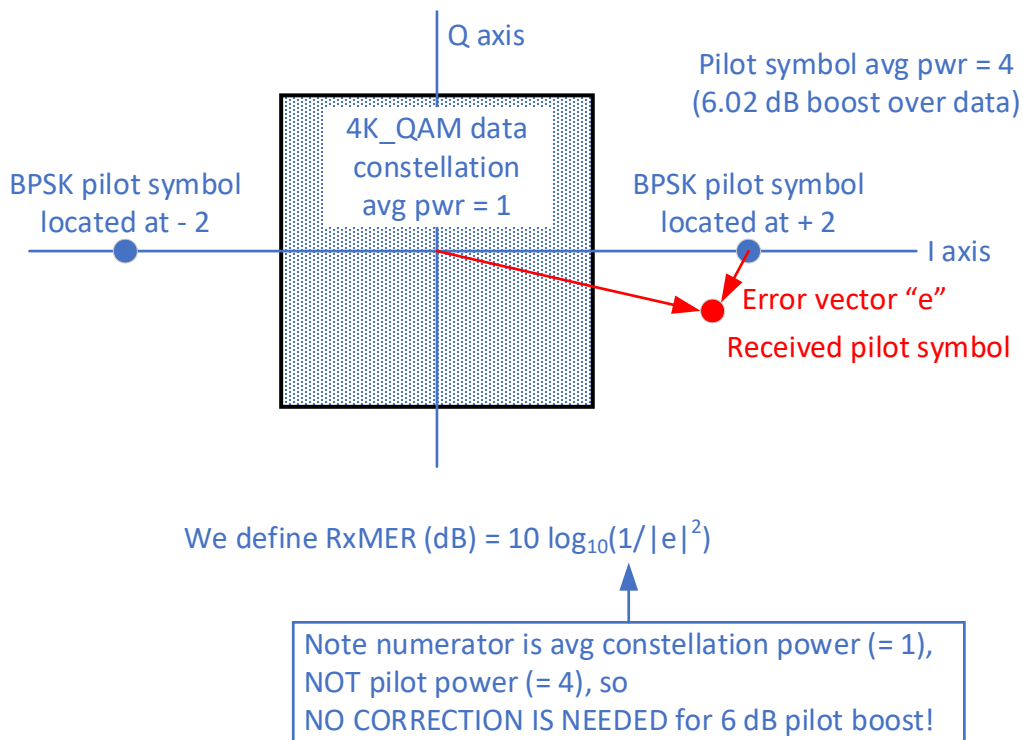
When a scattered pilot – which has a known BPSK value – lands on a data subcarrier location, an accurate measurement of the RxMER of that subcarrier over a wide range of SNR can be performed. For example, if narrowband ingress causes a given subcarrier to have a very low SNR, the RxMER measurement will still be accurate at that subcarrier because the ingress cannot cause the symbol value to

be interpreted in error (recall that the value is known ahead of time), and the error vector, though large due to the strong ingress, will correctly indicate the difference between the received and true symbol.

Thus the RxMER of all active subcarriers across the entire OFDM signal is periodically measured over time. If some subcarriers cannot be measured by the cable modem – for instance, in exclusion bands – the modem indicates that condition by reporting empty values in the measurement data for those subcarriers.

RxMER was carefully defined for the purposes of this measurement as the ratio of the average power of the ideal QAM constellation (numerator of the ratio, always equal to 1) to the average error vector power (denominator of the ratio). The error vector is the difference between the equalized received value and the known correct “scan pilot” value. For additive noise, the noise vector amplitude is not affected by the symbol amplitude, that is, whether or not the symbol is boosted. With this definition, since the numerator is the power of the QAM constellation rather than the “scan pilot” power, the RxMER measurement yields the true QAM RxMER even when the pilots are boosted by 6 dB relative to the data subcarriers. That is, for the case of additive noise, the pilot boost (in the case of scattered or continuous pilots) or lack of boost (in the case of PLC preamble symbols) is taken into account by design and no further compensation of the measurement is necessary to remove the effect of the pilot boosting. For some types of noise, such as phase noise, there may be some dependence of the error vector amplitude as a function of symbol amplitude, and a correction to RxMER may be necessary to reflect the actual noise on the data subcarriers as opposed to the boosted pilots.

The following example will help make this definition clear: For an ideal additive white Gaussian noise (AWGN) channel, an OFDM signal containing a mix of QAM constellations with  $\text{CNR} = 35 \text{ dB}$  on the QAM data subcarriers, will yield an RxMER measurement of nominally 35 dB averaged over all subcarrier locations. That is, RxMER is defined to match the CNR. Figure 5 illustrates how the “scan pilots” are used to compute RxMER per subcarrier. The figure shows the case of scattered or continuous pilots, which are boosted by 6 dB relative to data symbols, and a 4096-QAM data constellation.



**Figure 5. OFDM RxMER computation.**

## 2. Impairment Identification

A graph of RxMER per subcarrier is a useful tool for identifying and characterizing a variety of impairments. Figure 6 shows an impairment-free example of RxMER per subcarrier for a 96 MHz-wide OFDM signal, captured using a DOCSIS 3.1 cable modem-equipped field meter. This particular capture was made by one of the authors [Hranac] on his subscriber drop, which is connected to a 4 dB, two-port end-of-line tap after a node+3 cascade. The OFDM signal is carried in the upper end of the cable network's downstream spectrum.






Figure 6. RxMER per subcarrier graph for a 96 MHz-wide OFDM signal.

### 2.1.1. Simulated Ingress

One impairment of interest to cable operators is in-channel ingress. When ingress is present in an OFDM signal, it can be identified by a reduction of RxMER on the subcarriers that overlap the ingress. Figure 7 shows an example graph of RxMER per subcarrier for a 96 MHz-wide OFDM signal, in which simulated ingress from an idealized 10 MHz-wide LTE signal causes an approximately 10 dB reduction in RxMER on the affected subcarriers (an example with real ingress is included in Section 2.1.3).





Figure 7. OFDM RxMER per subcarrier plot showing simulated ingress and its impact on the affected subcarriers.

### 2.1.2. Inverted Plot and Equalized Noise Floor

An inverted graph of RxMER versus subcarrier frequency (that is, -RxMER per subcarrier) gives a plot of the underlying noise (including ingress) in the channel relative to the signal, after receive equalization. Figure 8 shows the OFDM signal and simulated ingress from Figure 6, but with the RxMER plot inverted.

The y-axis may be labeled “-RxMER (dB)” or “Equalized Noise Floor (dBc)”;

both are equally descriptive of the data being plotted.




**Figure 8. Inverted plot of RxMER per subcarrier.**

The value of inverting the RxMER per subcarrier plot is that the graph now shows the equalized noise and ingress underneath the OFDM channel.

### 2.1.3. Impairment Examples


This section includes examples of RxMER per subcarrier graphs with a variety of impairments observed in production cable networks.

Figure 9 is captured data that shows ingress in an OFDM signal, and its impact on RxMER per subcarrier. This example is more typical of what the effect of real ingress looks like. (Note: In Figure 9 and some subsequent figures, the parallel horizontal red lines are standard deviation plot bars and the horizontal black line in between is the average value.)




**Figure 9. OFDM RxMER per subcarrier graph showing ingress interference at about 890 MHz.**

Figure 10 shows an example of “edge rolloff” in the RxMER per subcarrier graph. The reduction in RxMER at the left edge of the OFDM signal is caused by a combination of the presence of an adjacent SC-QAM signal and the configuration of the cyclic prefix samples ( $N_{cp}$ ) and rolloff period samples ( $N_{rp}$ ), which affect how much the edge of the OFDM spectrum overlaps the adjacent signal. This phenomenon is discussed in more detail later in this paper.




**Figure 10. "Edge rolloff" visible at the left edge of the RxMER per subcarrier plot.**

Figure 11 shows a combination of impairments, and the presence of an exclusion band from 759.85 MHz to 763.55 MHz. Starting at the left side of the graph, RxMER “edge rolloff” is visible. Ingress is evident at several frequencies within the OFDM signal, and the overall RxMER per subcarrier decreases from left-to-right, suggesting possible CNR degradation and/or cable network frequency response problems.



**Figure 11. Multiple impairments are evident in this RxMER per subcarrier graph (see text), along with an exclusion band from 759.85 MHz to 763.55 MHz.**

Figure 12 shows an example in which a reflection caused amplitude ripple to be visible in the RxMER per subcarrier graph. This phenomenon is discussed in detail later in this paper.



**Figure 12. RxMER per subcarrier graph with amplitude ripple caused by an impedance mismatch-related reflection.**

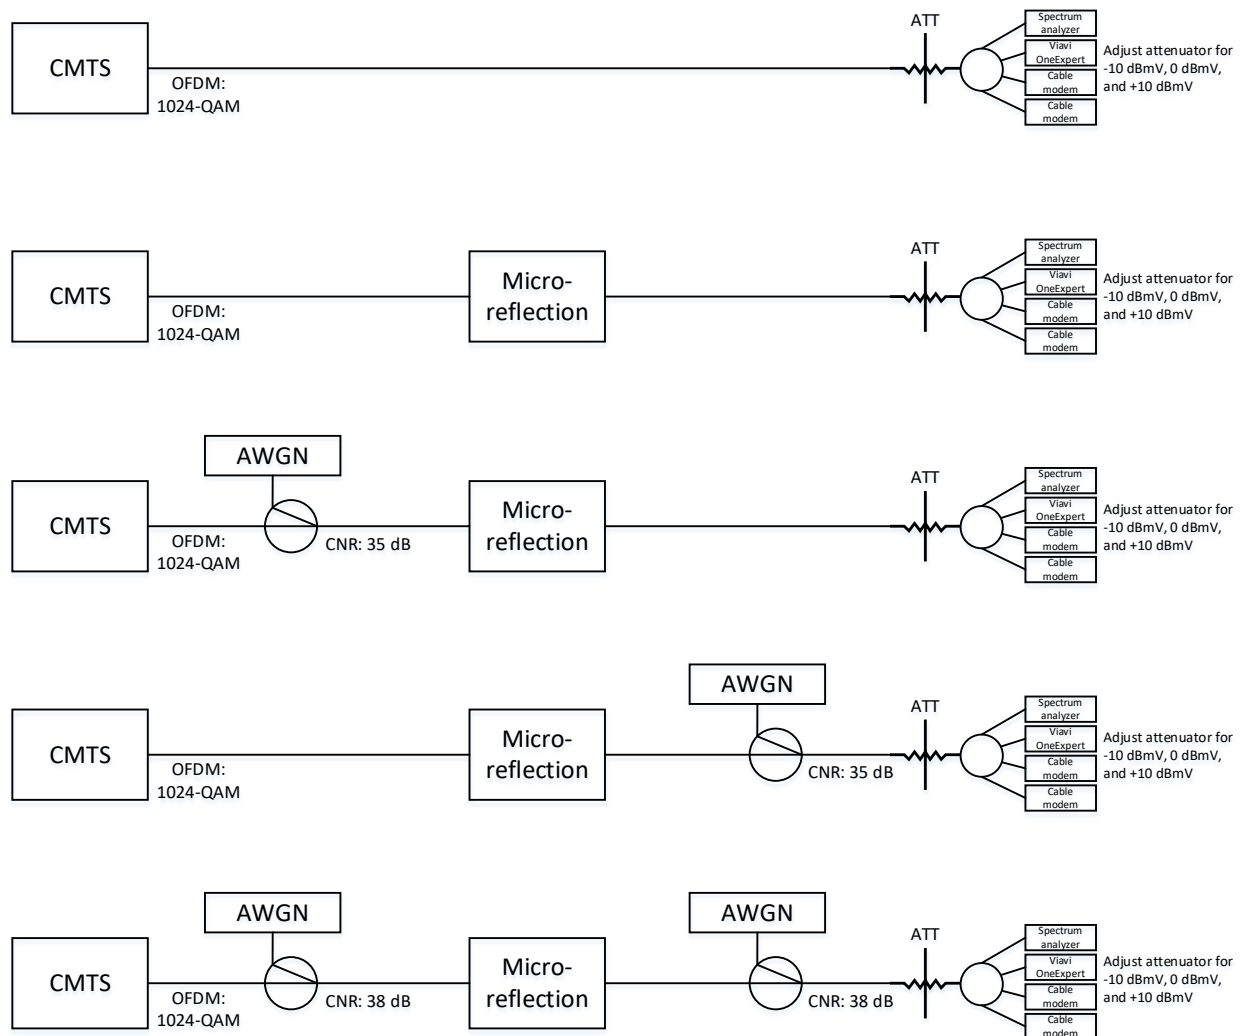
### 3. Lab Testing

A reduction in RxMER per subcarrier when the CNR is low, or when ingress affects certain subcarriers, is expected behavior. Amplitude ripple and edge rolloff in RxMER per subcarrier graphs warranted further investigation, because those phenomena were found to not appear consistently when the underlying mechanisms that cause them occur. Testing was done in CableLabs and Akleza test labs to recreate

amplitude ripple and edge rolloff, and develop a better understanding of why and when those impairments appear.

### 3.1. Amplitude Ripple

The appearance of amplitude ripple in the frequency domain – for instance, as viewed on a spectrum analyzer or a broadband sweep display – occurs when an impedance mismatch (or impedance mismatches) causes a reflection (or reflections). Under some circumstances amplitude ripple can also appear in an RxMER per subcarrier graph. Lab testing was done to better characterize this phenomenon. Figure 13 shows high-level block diagrams of equipment configurations used for five different amplitude ripple test scenarios.



**Figure 13. Equipment configurations for RxMER per subcarrier amplitude ripple testing.**


The micro-reflection in Figure 13 was set up to produce amplitude ripple with approximately 8 MHz spacing. Data captured during the lab testing included the frequency domain display on a spectrum analyzer; and all of the following from one or more of the cable modems (including a field instrument with an embedded DOCSIS 3.1 cable modem): channel estimate, RxMER per subcarrier, and OFDM

channel power – that is, power per 6 MHz.<sup>3</sup> Data was captured with the RF input power (OFDM channel power) to the modems at -10 dBmV, 0 dBmV, and +10 dBmV. The AWGN source was configured for wideband output, and adjusted to produce an aggregate CNR of 35 dB.

The OFDM signal was 192 MHz wide (292 MHz to 484 MHz), with 190 MHz encompassed spectrum,  $N_{\text{rp}} = 256$  (0.975 MHz taper region width), 50 kHz subcarrier spacing, and Profile A set to 1024-QAM.

The captured test results are summarized graphically in the Appendix as Test Case 1A, Test Case 1B, Test Case 1C, Test Case 2A, Test Case 2B, and so on, where “A,” “B,” and “C” refer to RF input levels of -10 dBmV, 0 dBmV, and +10 dBmV per 6 MHz respectively. Test Case 1, 2, 3, 4, and 5 refer to the five configurations shown in Figure 13.

The following figures highlight RxMER per subcarrier measurement results for Test Case 3B and 4B (see the Appendix for all Test Case results). The figures here show how amplitude ripple in the frequency domain (Figure 14) sometimes does not cause amplitude ripple to appear in an RxMER per subcarrier graph (Figure 15), and sometimes does (Figure 16). In both of these test cases, the nominal OFDM channel power at the input to the modem was 0 dBmV (Figure 17). An important takeaway from the lab testing is confirmation that amplitude ripple in the channel in the frequency domain will not always result in visible amplitude ripple in an RxMER per subcarrier graph. For the why behind this, refer to Section 3.1.1.



**Figure 14. Test Case 4B spectrum analyzer capture (the frequency domain ripple in the channel was the same for Test Case 3B).**

<sup>3</sup> OFDM channel power is expressed in terms of the power per CTA channel – that is, the power per 6 MHz. The total power is *Power per CTA channel* +  $10\log_{10}(\text{Number of occupied CTA channels})$  for that OFDM channel. When discussing OFDM signal level (channel power) in this paper, the stated value is the average power per 6 MHz, unless otherwise noted.






Figure 15. Test Case 3B RxMER per subcarrier graph showing little or no amplitude ripple.




Figure 16. Test Case 4B RxMER per subcarrier graph showing amplitude ripple.

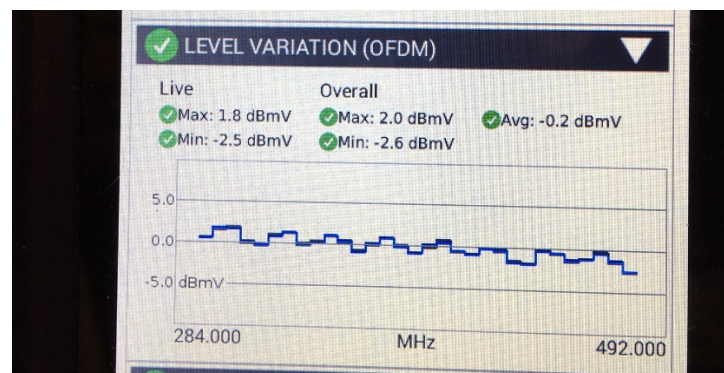



Figure 17. Test Case 4B OFDM channel power (the OFDM channel power was the same for Test Case 3B).

### 3.1.1. Amplitude Ripple Test Results Discussion

When the dominant noise (e.g., from optical fiber links and amplifiers) occurs before the source of a reflection – as was the situation with Test Case 3 – one will usually not see amplitude ripple in a graph of

RxMER per subcarrier. In contrast, when the dominant noise occurs after the source of a reflection, amplitude ripple will be visible in a graph of RxMER per subcarrier. There are some exceptions, as can be seen in the figures in the Appendix. The following explains why the location of dominant noise relative to a reflection affects the visibility of RxMER per subcarrier amplitude ripple.

Figure 18 is a block diagram of a test system for an OFDM signal, showing sources for an OFDM signal and SC-QAM signals; a noise injection point (“Noise 1”) before the source of an echo (reflection); an echo creation circuit with frequency response  $H(f)$ ; and a second noise injection point (“Noise 2”) after the echo circuit. A cable modem is shown as two functional blocks: an analog front end (AFE, including an A-D converter and amplifiers) and an adaptive equalizer, the latter labeled  $H(f)^{-1}$ , indicating that its purpose is to invert the echo channel response. A noise source, “Noise 3”, represents the noise added internally to the modem, often described by the noise figure of the modem<sup>4</sup>. A modem’s typical noise figure can be as good as 5 dB to 10 dB, but could be higher, especially when attenuation is inserted by the AFE at high input signal power. Signal and noise observation points A, B, and C are labeled.

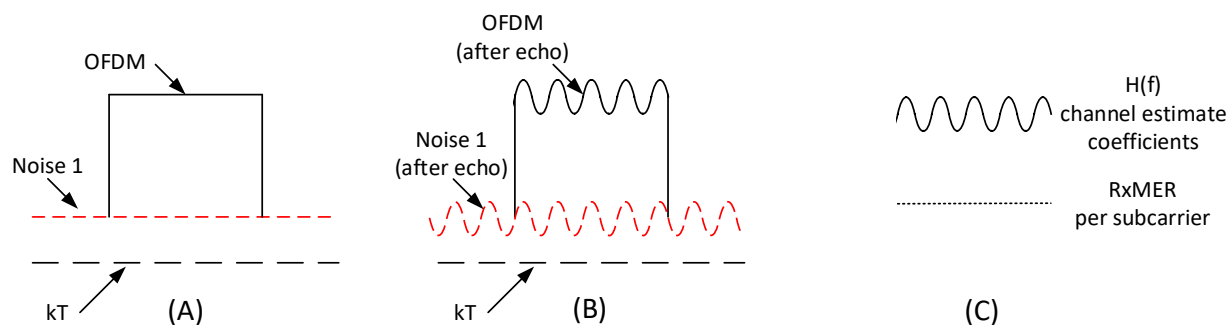


**Figure 18. OFDM signal test system block diagram.**

Figure 19 illustrates signals at points A, B, and C for test conditions with no SC-QAM signals, Noise 1 active, and Noise 2 quiet. At point A, the OFDM signal is flat; Noise 1 is flat, and  $kT$  is ever-present. At point B the echo has put a ripple in both the OFDM signal and Noise 1. At point C the channel estimate coefficients  $H(f)$  and RxMER per subcarrier are both reported with different MIBs. Note that the RxMER has been flattened by the equalizer since Noise 1 has passed through both the echo channel and its inverse, resulting in no net echo at the modem slicer. This happens automatically when the cable modem’s equalizer flattens the OFDM signal. The resulting RxMER per subcarrier plot will be a function of the level of Noise 1 relative to the noise floor of the cable modem. If Noise 1 is very small, the modem’s noise floor may be dominant and some ripple may be seen in the RxMER per subcarrier plot; if Noise 1 is much higher than the internal modem noise floor, the cable modem’s internally generated noise will not contribute appreciably to the RxMER plot, which will be flat. Elements that contribute to the internally generated modem noise are amplifier noise, bits of precision in the A-D converter, phase noise, quantizing error in digital computations, etc.

<sup>4</sup> Noise power spectral density  $kT$ , where  $k$  is Boltzmann's constant ( $1.38 \times 10^{-23}$  joules/kelvin) and  $T$  is the effective noise temperature in kelvin. After the modem AFE,  $T$  is elevated above the standard temperature  $T_0 = 290$  K due to the noise added by the modem.






**Figure 19. Test conditions with Noise 1 injected before the source of an echo (reflection). As seen at observation point (B) both Noise 1 and the OFDM signal have amplitude ripple, but the modem's adaptive equalizer removes the channel effect so that at observation point (C) the RxMER-per-subcarrier plot does not show ripple. In this figure, the horizontal axis is frequency, and the vertical axis is relative amplitude in decibels.**

Figure 20 illustrates signals at observation points A, B, and C for test conditions with no SC-QAM signals, Noise 1 quiet, and Noise 2 active. At point A the OFDM signal is flat and noise is  $kT$  background noise. At point B, the OFDM signal has an echo, but Noise 2 is flat since it was injected after the echo creation circuit. At point C the channel estimate coefficients  $H(f)$  and RxMER per subcarrier are reported. However, equalization has been applied to the injected flat Noise 2, giving it an inverse channel response, which exhibits ripple. The RxMER will exhibit the original (uninverted) channel ripple. This is because RxMER is a signal-to-noise ratio (SNR), with  $N$  in the denominator. So the ripple was inverted twice: once by the inverse channel equalizer and once in the RxMER computation. Since two inversions yield a non-inversion, we see the original ripple signature in the RxMER plot.

Note that in determining whether or not a ripple is observed in RxMER, it matters whether the elevated random noise experienced the echo or not, illustrated by the differences in Figure 18(C) and Figure 19(C). Furthermore, if the OFDM signal level is very weak at the input to the cable modem, the background noise  $kT$  will become the dominant noise component in RxMER, and the ripple in the RxMER-per-subcarrier plot will be diminished.


To summarize: If the noise passes through both the echo and the inverse equalizer, these two filtering operations will cancel and the RxMER-per-subcarrier plot will NOT show channel ripple. If the noise does not pass through the echo but does pass through the inverse equalizer, the RxMER WILL show channel ripple, and the ripple in the RxMER-per-subcarrier plot will be upright (same polarity as the channel estimate coefficients, that is, not inverted).



**Figure 20. Test conditions where Noise 2 is injected after the source of an echo (reflection). At observation point (B), Noise 2 does not have ripple because it did not pass through the echo. However, the modem's RxMER-per-subcarrier plot does have amplitude ripple (C) because the noise passed through the receive equalizer  $H(f)-1$ . In this figure, the horizontal axis is frequency, and the vertical axis is relative amplitude in decibels.**

### 3.2. RxMER Edge Rolloff

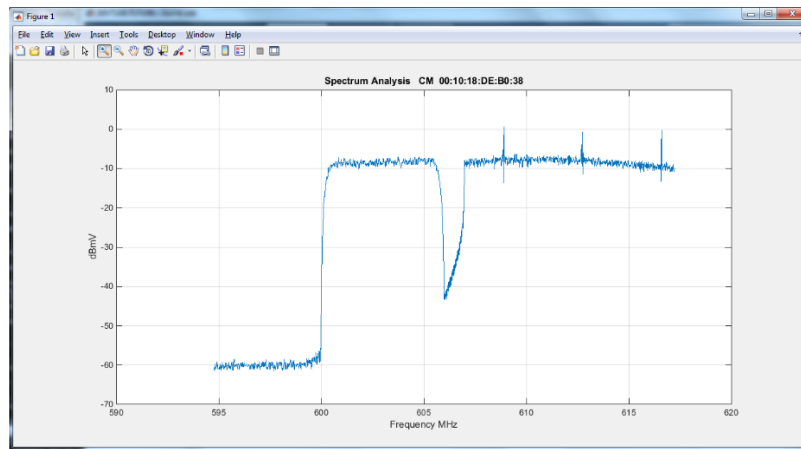
Observations from production cable networks and subsequently recreated in test lab environments have shown that the presence of an SC-QAM signal adjacent to an OFDM signal can cause a degradation of the reported RxMER values in the subcarriers near the edge of the OFDM signal, depending upon the configuration of the cyclic prefix and rolloff period parameters. When looking at a graph of RxMER per subcarrier, this effect can present itself as rolloff at the band edge(s) of the OFDM channel as shown in the highlighted areas in Figure 10 and Figure 21.



**Figure 21. Another example of edge rolloff in an RxMER per subcarrier graph.**


In the aforementioned examples, the signals were aligned on standard CTA channel boundaries with an SC-QAM signal adjacent to the OFDM signal.  $N_{cp}$  was configured to 192 and  $N_{rp}$  to 128. Figure 22 shows an example spectrum capture with the same signal conditions including  $N_{cp}$  and  $N_{rp}$  settings for the OFDM signal as in Figure 21; the taper region width is 1.875 MHz. Even though the RF spectrum shows a guard band of sorts between the SC-QAM and OFDM signals, the OFDM taper region extends into the lower adjacent channel, introducing degradation, although possibly minimal. In this example, energy from

the adjacent SC-QAM signal also leads to interference in the first few subcarriers of the OFDM signal, resulting in the edge rolloff in the RxMER per subcarrier graph visible in Figure 10 and Figure 21.



**Figure 22. Example spectrum capture of an SC-QAM signal adjacent to an OFDM signal. Note the OFDM signal's taper region, which extends into part of the SC-QAM signal.**

While degradation of the RxMER of the first few OFDM subcarriers does not have a meaningful impact on overall performance or throughput, changing the configuration of cyclic prefix and rolloff period can help. In particular, proper configuration of  $N_{cp}$  can sharpen the OFDM signal's spectral edges in the frequency domain (that is, reduce the taper region width). Figure 23 shows a graph of RxMER per subcarrier after adjusting  $N_{cp}$  to 1024 and  $N_{rp}$  to 256. This is data from the same modem as shown in Figure 10, but with the new cyclic prefix and rolloff period settings.



**Figure 23. RxMER per subcarrier graph after reconfiguring  $N_{cp}$  and  $N_{rp}$ .**

One can avoid or minimize adjacent channel interference and OFDM RxMER per subcarrier edge rolloff by properly configuring the OFDM signal's  $N_{cp}$  and  $N_{rp}$  values. Appendix V of the DOCSIS 3.1 Physical Layer Specification includes a table showing taper region width versus  $N_{rp}$  setting (see Figure 24). If the channel(s) adjacent to the edge(s) of the OFDM signal will be occupied, then an OFDM band edge exclusion should be configured in accordance with the table in Figure 23. From the table in

Figure 24, maximum  $N_{rp} = 256$  will yield a taper region width of either 0.975 MHz or 0.9875 MHz, depending on subcarrier spacing. Keep in mind the DOCSIS requirement that the  $N_{rp}$  value must be less than the  $N_{cp}$  value.

FFT	Roll-Off Period Samples ( $N_{rp}$ )	Taper Region (MHz)
4K	64	3.575
	128	1.875
	192	1.325
	256	0.975
8K	64	3.3375
	128	1.7125
	192	1.1625
	256	0.9875 <sup>1</sup>
1. The taper region of 0.9875 MHz is in accordance with the requirement for a minimum taper region of 1 MHz minus half subcarrier spacing. Achieving up to approximately 0.5 dB impact to the noise power in the adjacent spurious emissions integration region would allow a taper region of 0.8625 MHz, if the specification did not mandate the minimum taper region to be larger than this.		

**Figure 24. Table 75 from Appendix V of the DOCSIS 3.1 Physical Layer Specification [1].**

### 3.2.1. RxMER Edge Rolloff Discussion

During the field observations and lab testing done for this paper, the edge rolloff was not always consistently visible in a graph of RxMER per subcarrier. Under some conditions the rolloff was seen, but other times the rolloff was not seen, despite the same configuration (existence of adjacent SC-QAM signal, same  $N_{cp}$  and  $N_{rp}$  settings, etc.). Indeed, field observations indicated that some modems in the same node service area displayed the edge rolloff, while others did not, including two modems that were in homes next door to each other.

Preliminary results suggested that factors such as the total power at the input to the modem, the presence of a micro-reflection, and even cable modem make/model (and silicon vendor) appeared to affect visibility of the rolloff. However, additional testing showed that when performing multiple RxMER captures from a given modem, in some cases the edge rolloff was not consistently present. Further investigation showed that with the intermittent rolloff in a cable modem, the visual impact in the RxMER band edges was notable; quantitatively, the impact to the link was minimal, but this has not been characterized. The impact to the link can be reduced to zero if the bit loading for the affected subcarriers is reduced by one or two bits, and the cost of this reduction in throughput is less than 0.1 bits per Hz (less than 0.5%). For the same transmission parameters, a different cable modem showed consistent band edge rolloff which was more severe than the modem which showed intermittent rolloff. Note that reducing the bit loading for a few (e.g., 100) subcarriers, impacts throughput much less than increasing  $N_{cp}$ .<sup>5</sup>

A DOCSIS 3.1 cable modem, usually employing a system on a chip, constantly adjusts several parameters to optimize throughput of a downstream OFDM channel and minimize frame errors. Adjustments may include tracking loops for receive carrier frequency/phase and timing (selection of the samples used in FFT processing); and channel tracking (adaptive equalization). These dynamic adjustments may result in seemingly intermittent behavior under certain signal conditions while the receiver is actually maintaining tight, near optimal link performance. An analogous situation occurs in

<sup>5</sup> Optimization of  $N_{cp}$  and  $N_{rp}$  settings is important because larger values add overhead, which in turn impacts usable throughput.

SC-QAM where an adaptive equalizer may exhibit some frequency response fluctuation or path wander while maintaining consistent performance throughout.


#### 4. Using RxMER per Subcarrier Data for Maintenance and Troubleshooting

CableLabs' PNM Best Practices [5] document outlines a fault localization process based on correlating pre-equalization or full band capture data from cable modems based on their network topology. The procedure identifies the point along a shared path where a fault or anomaly indicator changes. This is a similar process that field technicians use to determine the area of a fault using field test meters where readings are taken while working upstream from a fault until the impairment is no longer visible. The area where this change occurs gives an indication of the location of the cable plant that must be contributing to the fault. With the advent of PNM software tools capable of taking readings from existing cable modems, the determination of the approximate fault location can now be achieved without having to roll a truck, making the technician's troubleshooting efforts much easier.

As described in the previous sections, OFDM RxMER per subcarrier data can indicate impairments related to cable plant faults such as an echo or reflection resulting in amplitude ripple in the RxMER per subcarrier plot. By analyzing results from cable modems that show amplitude ripple in the RxMER per subcarrier caused by a reflection, the area of the impairment can be isolated when correlated with the shared network path. Comparing the reported RxMER per subcarrier of devices on the same shared path, a boundary can be determined where a problem exists and where it does not. The fault therefore must be between the last device showing the fault and the first device not showing the fault.

Additionally, the peak-to-peak frequency spacing of the amplitude ripple can be used to compute the approximate length of an "echo tunnel," the distance between two impedance mismatches (e.g., and amplifier and a tap full of water). This echo tunnel length can then be used to further narrow down the possible location of the fault along the shared path.

Figure 25 shows the formula for computing the echo tunnel length based on the peak-to-peak frequency spacing of amplitude ripple.



Length in feet:

$$L = 492 * (VF/F_{\text{MHz}})$$

$$L = 492 * (0.85/8.35)$$

$$L = 50.08 \text{ feet}$$

where:

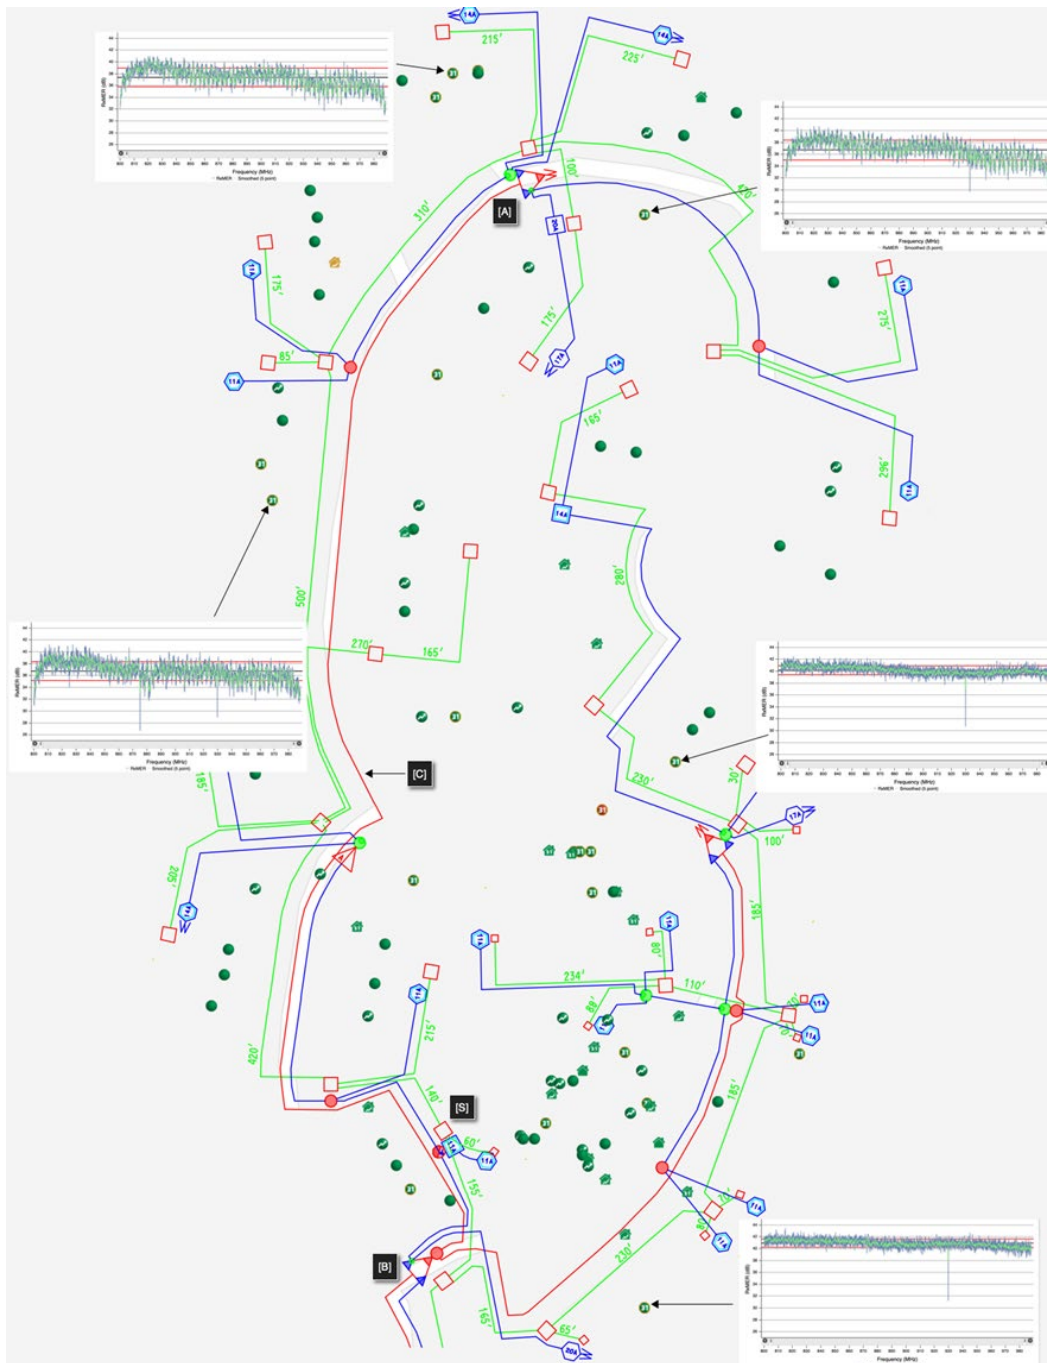
L = length in feet

VF = velocity factor of the cable (velocity of propagation expressed in decimal form)

F<sub>MHz</sub> = frequency spacing of amplitude ripple in megahertz

**Figure 25. Amplitude ripple-to-echo tunnel length calculation.**

Figure 26 shows a real-world example from a production network and the resulting analysis and findings. Cable modem icons are shown along with an overlay of the cable plant in the area. This area shows a number of DOCSIS 3.1 cable modems connected to an amplifier [A] that all show a similar and significant amplitude ripple in a plot of their RxMER per subcarrier data. This amplitude ripple is however not seen on any cable modem connected to the other outputs of the amplifier and directional coupler show at point [B]. This analysis would therefore indicate that the cause of the impairment is located on the interconnecting trunk cable [C].




**Figure 26. System map with RxMER per subcarrier data.**



The trunk cable in this area is underground, however above-ground trunk block splices are located at different points along the path in pedestals. The locations of these pedestals with block splices are indicated with red boxes such as shown at [S].

Zooming into the RxMER per subcarrier and measuring the peak-to-peak distance in MHz shows a periodicity of approximately 3 MHz as illustrated in Figure 27. (There also appears to be at least one other ripple visible in the graph.)



**Figure 27. Close-up of RxMER per subcarrier amplitude ripple periodicity.**

Using the amplitude ripple length calculation formula shown in Figure 25 one can compute the approximate length of the echo tunnel impacting this part of the network. The velocity of propagation for the trunk cable being used in this particular system is 93%. Plugging these values into the formula results in a calculated echo tunnel length of about 152.5 feet.

The cable plant diagram shown in Figure 26 shows a trunk span of 155 feet between the amplifier [B] and the first splice block [S] on the impaired trunk. A field maintenance crew was dispatched to the area and after verifying the RxMER per subcarrier responses reported were consistent with what was being reported by their field meters, they proceeded to inspect and replace the trunk block splice [S] shown in Figure 28 and Figure 29.




**Figure 28. Technician troubleshooting RxMER per subcarrier amplitude ripple.**



**Figure 29. Trunk splice removed from network.**

While no obvious impairment was visually apparent, a refresh of the RxMER per subcarrier data for the impacted cable modems attached to this trunk showed that fault had been corrected and that the average RxMER per subcarrier increased by 2 dB after replacing the splice (see Figure 30). The technicians also replaced a suspect directional coupler and second splice, which helped improve the response somewhat.





**Figure 30. RxMER per subcarrier after replacing splice.**

The same general techniques described in this section to determine the location of an impairment can also be used for other types of interference visible in the RxMER per subcarrier data. For example, if there is an indication of LTE ingress, this is caused by degradation of shielding effectiveness due to a cable crack, loose or corroded connector, etc. By isolating those modems that show this interference the general area of the source of the fault can be found. If the problem is visible on the cable modem from one home but not a neighboring home, then most likely the fault is in the drop cable or in the home itself. If, however, interference is visible across a number of cable modems in an area, then the fault is most likely farther upstream in a section of cable on the common path. As in the previous example, by examining data from a number of points progressively farther upstream on the common path, one can determine the location between where the interference is visible and where it isn't. This then defines the bounds of the area to investigate and locate the problem.

## 5. Areas for Further Investigation

One area for additional investigation is field and lab testing to characterize the impact of nonlinear distortions – composite second order, composite triple beat, and noise-like composite intermodulation distortion – on RxMER per subcarrier. For instance, if an OFDM signal's RxMER per subcarrier was found to unexpectedly improve deeper in an amplifier cascade, that might be an indication of the presence of nonlinear distortion(s).

Another area for further investigation is related to impairment testing using one or more SC-QAM signals under an OFDM signal to simulate ingress. Some testing using this method has shown that as the underlying SC-QAM signal amplitude was increased, the RxMER on the affected subcarriers decreased as expected. However, as the “interference” amplitude increased even more, in one cable modem the RxMER on all subcarriers started to decrease, but on another vendor's modem there was no observed RxMER decrease outside of the interference region. Additional testing could be done to determine whether this behavior is consistent and repeatable.

One other area for further investigation is evaluation of the effectiveness of DOCSIS 3.1's frequency domain interleaving and LDPC FEC when in-channel ingress is present and when the SNR is several dB above the threshold for the modulation order in use. When the overall SNR margin is high, some testing has shown that a block of noise (e.g., 6 MHz to 12 MHz wide) used to simulate ingress has little or no effect on overall throughput until the RxMER on the affected subcarriers is well below the known margin

for the modulation order (profile) in use. There is a tradeoff with respect to the bandwidth of the interference relative to the bandwidth of the OFDM signal and the amount of SNR margin. Additional testing could more accurately quantify interleaving and FEC performance vs the level and bandwidth of the interference..

## Conclusion

SC-QAM RxMER has long been an important and useful metric for cable operators, but it does have limitations. The ability to take advantage of DOCSIS 3.1 OFDM RxMER per subcarrier data is even more useful and goes beyond the limitations of SC-QAM RxMER. Graphs of OFDM RxMER per subcarrier can in many instances be used to identify the type(s) of impairment(s) affecting the network, and nicely complements other tools for maintaining optimum performance, maximum throughput, and overall subscriber satisfaction.

## Abbreviations

ACI	adjacent channel interference
A-D	analog-to-digital
AFE	analog front end
AWGN	additive white Gaussian noise
CableLabs	Cable Television Laboratories
CCAP	converged cable access platform
CMTS	cable modem termination system
CNR	carrier-to-noise ratio
CTA	Consumer Technology Association
dB	decibel
dBmV	decibel millivolt
DOCSIS	Data-Over-Cable Service Interface Specifications
FEC	forward error correction
I	in-phase
ISBE	International Society of Broadband Experts
LDPC	low density parity check
log	logarithm
LTE	long term evolution
MER	modulation error ratio
MHz	megahertz
MIB	management information base
$N_{cp}$	cyclic prefix samples
$N_{rp}$	rolloff period samples
OFDM	orthogonal frequency division multiplexing
PLC	PHY link channel (also physical layer link channel)
PNM	proactive network maintenance
Q	quadrature
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying
RF	radio frequency

RxMER	receive modulation error ratio
SC-QAM	single carrier quadrature amplitude modulation
SCTE	Society of Cable Telecommunications Engineers
SNR	signal-to-noise ratio

## Bibliography & References

[1] Data-Over-Cable Service Interface Specifications DOCSIS 3.1 Physical Layer Specification (CM-SP-PHYv3.1-I16-190121); Cable Television Laboratories

[2] Hranac, R., Currivan, B. “Digital Transmission: Carrier-to-Noise, Signal-to-Noise, and Modulation Error Ratio.” In *Proceedings Manual and Collected Technical Papers*, SCTE Cable-Tec Expo 2007, Orlando, FL

[3] Hranac, R., Currivan, B., “Understanding Real-World MER Measurements.” In *Presentations and Collected Technical Papers*, SCTE Cable-Tec Expo '11, Atlanta, GA

[4] Cable Television Channel Identification Plan (CTA-542-D R-2018); Consumer Technology Association

[5] DOCSIS Best Practices and Guidelines “PNM Best Practices: HFC Networks (DOCSIS 3.0)” (CM-GL-PNMP-V03-160725); Cable Television Laboratories

Copies of the following articles are archived on SCTE’s web site at

[https://www2.scte.org/SCTE2/Areas\\_of\\_Interest/Technical\\_Columns\\_Communications\\_Technology/SCTE2/Areas\\_of\\_Interest/ct\\_archives.aspx?hkey=a9cc9833-247e-4476-a1a9-9b5d4daaa8e8](https://www2.scte.org/SCTE2/Areas_of_Interest/Technical_Columns_Communications_Technology/SCTE2/Areas_of_Interest/ct_archives.aspx?hkey=a9cc9833-247e-4476-a1a9-9b5d4daaa8e8)

[6] Hranac, R., “Is MER Overrated?” (January 2010 *Communications Technology*)

[7] Hranac, R., “Making MER Better: Part 2” (September 2009 *Communications Technology*)

[8] Hranac, R., “Making MER Better: Part 1” (August 2009 *Communications Technology*)

[9] Hranac, R., “Equalized or Unequalized? That is the Question” (February 2007 *Communications Technology*)

[10] Hranac, R., “Modulation Error Ratio” (January 2007 *Communications Technology*)

# Appendix

## 6. Appendix I – Lab Test Results

The following figures highlight data measured during the lab testing. A variable attenuator was set to obtain three nominal levels (power per 6 MHz) at the cable modem inputs (-10 dBmV, 0 dBmV, and +10 dBmV) for each test case. The figures in this section detail the following four parameters.

**Spectrum analyzer screen capture** – A spectrum analyzer was tuned to the OFDM signal under test to show the signal in the frequency domain. The analyzer display was photographed for each test case.

**Channel estimate** – The downstream channel estimate coefficients (a cable modem's estimate of the downstream channel response) were obtained from a modem in the lab test setup and plotted in graph form.

**RxMER per subcarrier** – This data for each OFDM subcarrier was obtained from a modem in the lab test setup (same modem used for channel estimate) and plotted in graph form.

**OFDM channel power** – A screen capture of OFDM channel power (the RF power per CTA channel) was taken from a Viavi OneExpert field meter. Each short horizontal line represents the power per 6 MHz segment of the OFDM signal.

### 6.1. Test Case 1A




Figure 31. Test Case 1A spectrum analyzer screen capture.




Figure 32. Test Case 1A channel estimate.




Figure 33. Test Case 1A RxMER per subcarrier.




Figure 34. Test Case 1A OFDM channel power (nominal -10 dBmV).

## 6.2. Test Case 1B




Figure 35. Test Case 1B spectrum analyzer screen capture.




Figure 36. Test Case 1B channel estimate.




Figure 37. Test Case 1B RxMER per subcarrier.






Figure 38. Test Case 1B OFDM channel power (nominal 0 dBmV).

### 6.3. Test Case 1C




Figure 39. Test Case 1C spectrum analyzer screen capture.




Figure 40. Test Case 1C channel estimate.




Figure 41. Test Case 1C RxMER per subcarrier.

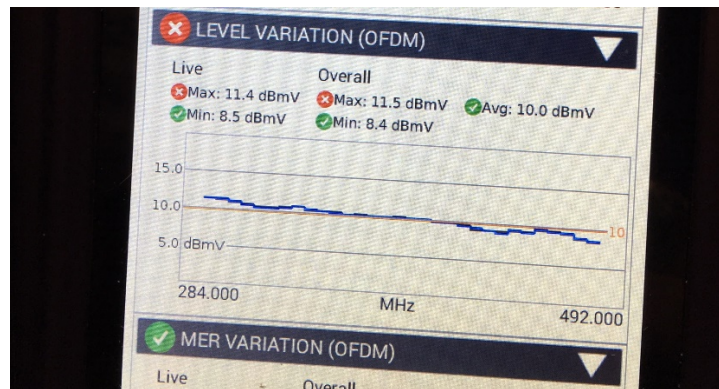


Figure 42. Test Case 1C OFDM channel power (nominal +10 dBmV).

#### 6.4. Test Case 2A

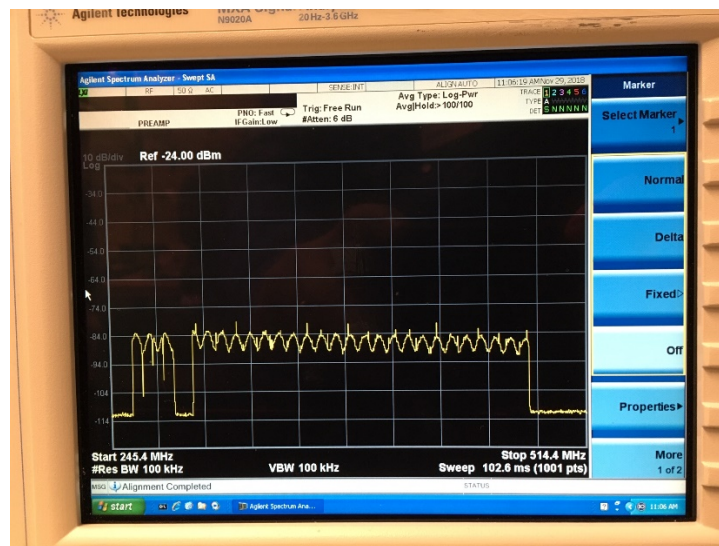


Figure 43. Test Case 2A spectrum analyzer screen capture.



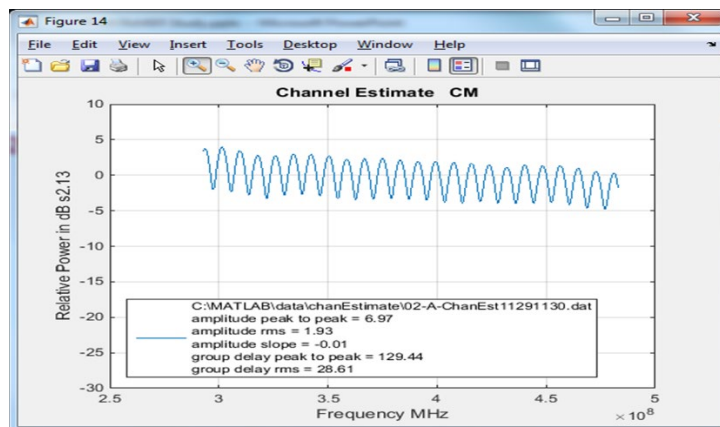


Figure 44. Test Case 2A channel estimate.




Figure 45. Test Case 2A RxMER per subcarrier.

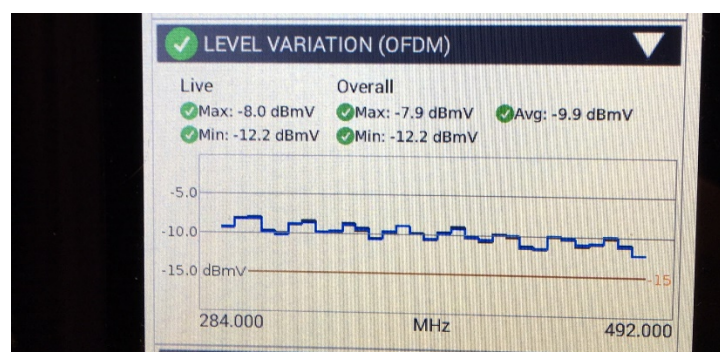


Figure 46. Test Case 2A OFDM channel power (nominal -10 dBmV).

## 6.5. Test Case 2B




Figure 47. Test Case 2B spectrum analyzer screen capture.




Figure 48. Test Case 2B channel estimate.




Figure 49. Test Case 2B RxMER per subcarrier.

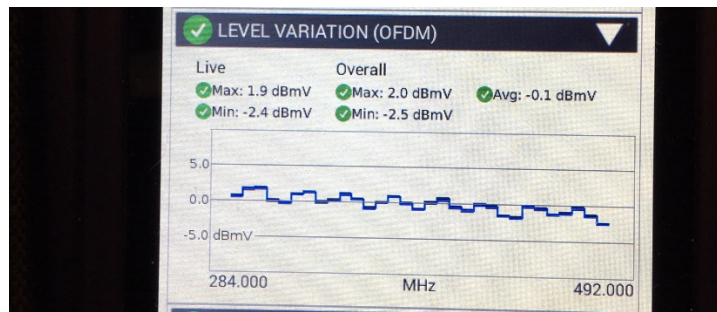


Figure 50. Test Case 2B OFDM channel power (nominal 0 dBmV).

## 6.6. Test Case 2C

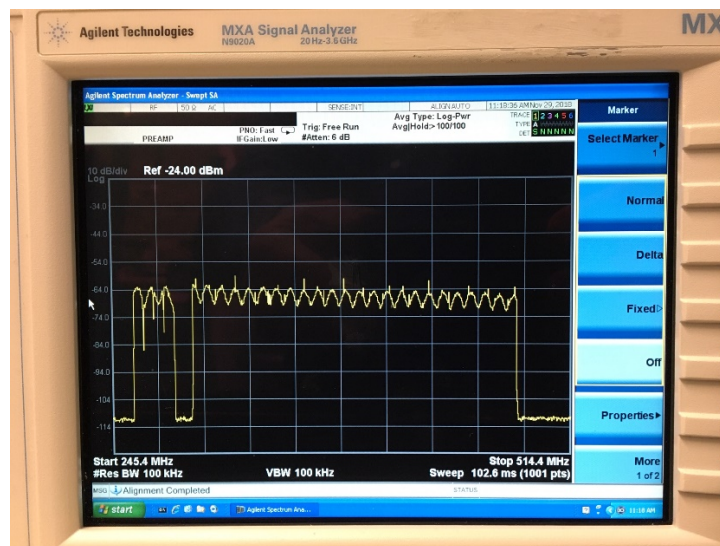


Figure 51. Test Case 2C spectrum analyzer screen capture.




Figure 52. Test Case 2C channel estimate.




Figure 53. Test Case 2C RxMER per subcarrier.




Figure 54. Test Case 2C OFDM channel power (nominal +10 dBmV).

## 6.7. Test Case 3A




Figure 55. Test Case 3A spectrum analyzer screen capture.






Figure 56. Test Case 3A channel estimate.




Figure 57. Test Case 3A RxMER per subcarrier.

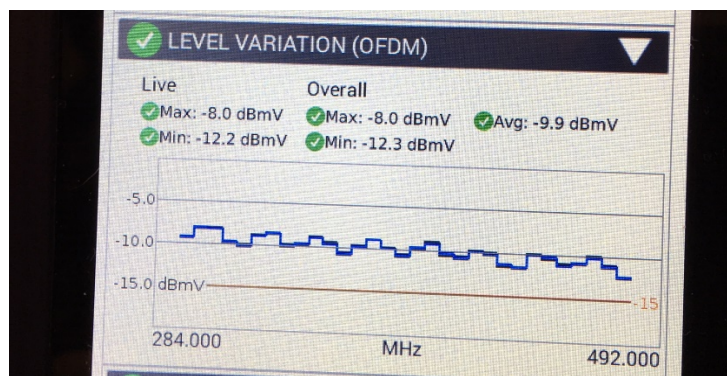


Figure 58. Test Case 3A OFDM channel power (nominal -10 dBmV).

## 6.8. Test Case 3B




Figure 59. Test Case 3B spectrum analyzer screen capture.




Figure 60. Test Case 3B channel estimate.




Figure 61. Test Case 3B RxMER per subcarrier.




Figure 62. Test Case 3B OFDM channel power (nominal 0 dBmV).

## 6.9. Test Case 3C



Figure 63. Test Case 3C spectrum analyzer screen capture.

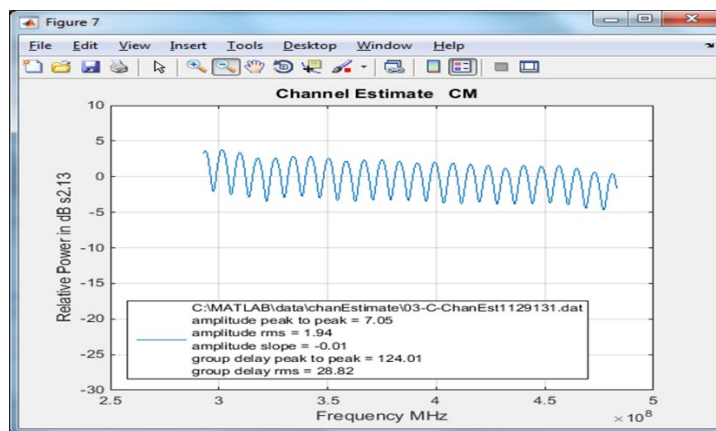


Figure 64. Test Case 3C channel estimate.




Figure 65. Test Case 3C RxMER per subcarrier.




Figure 66. Test Case 3C OFDM channel power (nominal +10 dBmV).

## 6.10. Test Case 4A




Figure 67. Test Case 4A spectrum analyzer screen capture.






Figure 68. Test Case 4A channel estimate.




Figure 69. Test Case 4A RxMER per subcarrier.

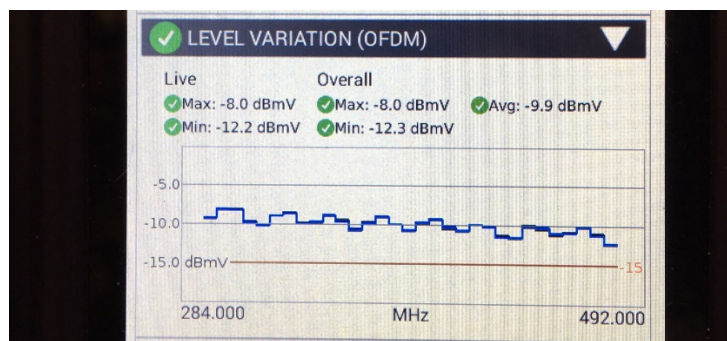


Figure 70. Test Case 4A OFDM channel power (nominal -10 dBmV).

## 6.11. Test Case 4B




Figure 71. Test Case 4B spectrum analyzer screen capture.




Figure 72. Test Case 4B channel estimate.




Figure 73. Test Case 4B RxMER per subcarrier.

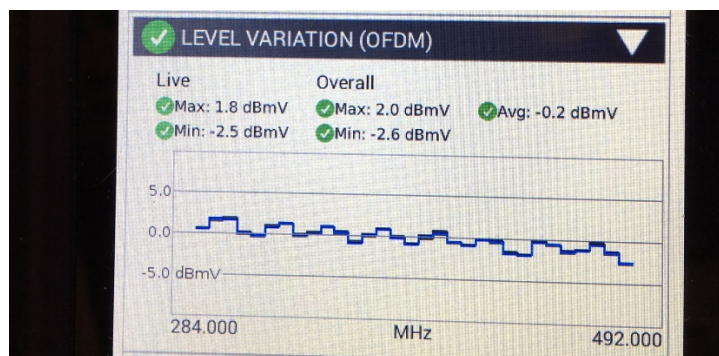


Figure 74. Test Case 4B OFDM channel power (nominal 0 dBmV).

## 6.12. Test Case 4C

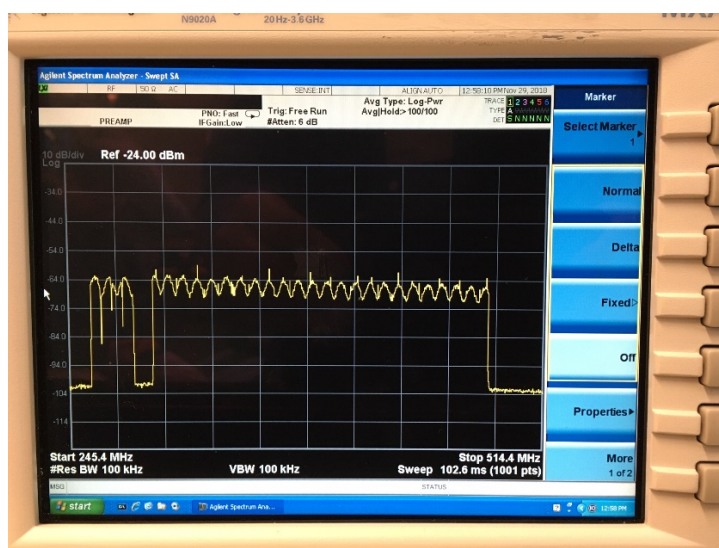


Figure 75. Test Case 4C spectrum analyzer screen capture.




Figure 76. Test Case 4C channel estimate.




Figure 77. Test Case 4C RxMER per subcarrier.

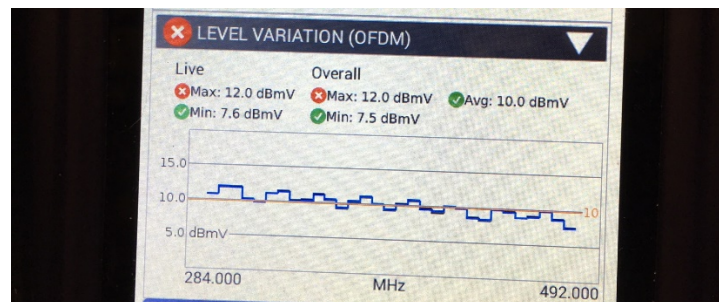


Figure 78. Test Case 4C OFDM channel power (nominal +10 dBmV).

### 6.13. Test Case 5A




Figure 79. Test Case 5A spectrum analyzer screen capture.



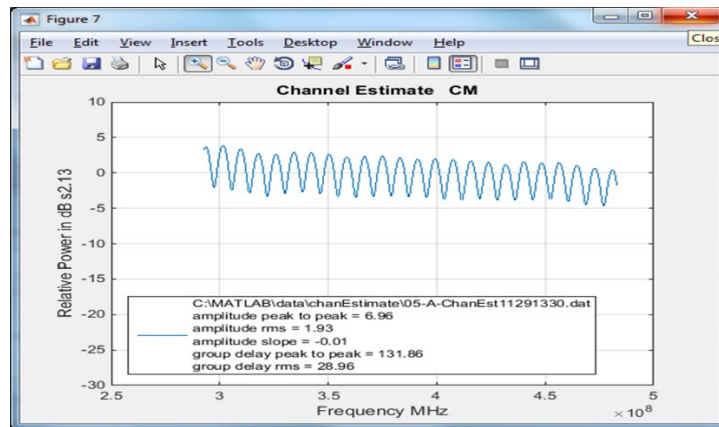


Figure 80. Test Case 5A channel estimate.




Figure 81. Test Case 5A RxMER per subcarrier.

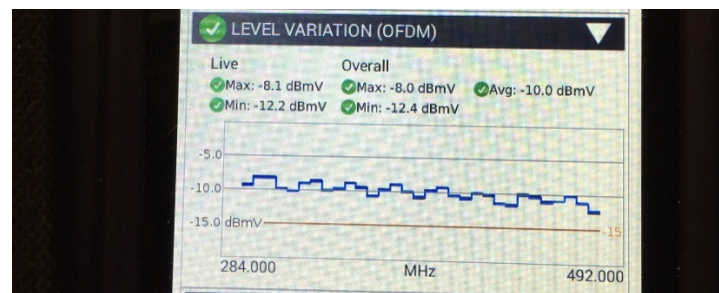


Figure 82. Test Case 5A OFDM channel power (nominal -10 dBmV).

## 6.14. Test Case 5B




Figure 83. Test Case 5B spectrum analyzer screen capture.

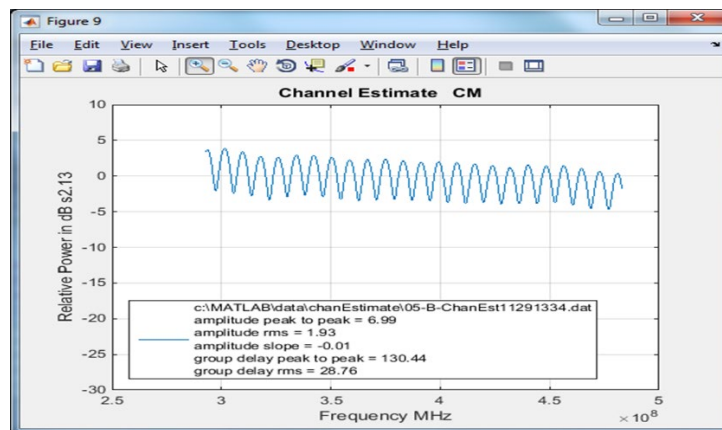


Figure 84. Test Case 5B channel estimate.




Figure 85. Test Case 5B RxMER per subcarrier.

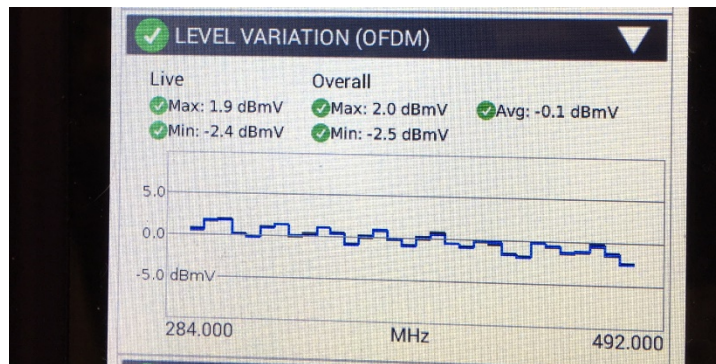


Figure 86. Test Case 5B OFDM channel power (nominal 0 dBmV).

### 6.15. Test Case 5C




Figure 87. Test Case 5C spectrum analyzer capture.

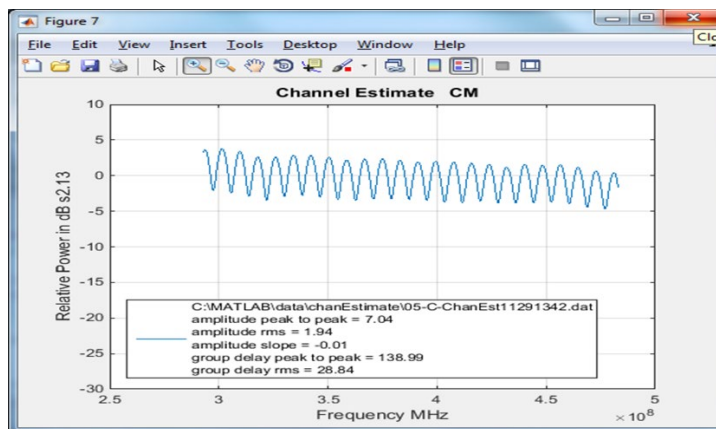




Figure 88. Test Case 5C channel estimate.



**Figure 89. Test Case 5C RxMER per subcarrier.**



**Figure 90. Test Case 5C OFDM channel power (nominal +10 dBmV).**

## 7. Appendix II – Acknowledgements

The authors wish to acknowledge the following:

- CableLabs and Akleza for providing facilities and equipment to conduct lab testing for this paper.
- Viavi Solutions for providing a loaner OneExpert CATV (ONX-630) meter for some of the testing.
- Blue Ridge Communications, a subsidiary of Pencor Services, Inc., for allowing the authors to collect field data from subscriber cable modems.



# **Maximizing the Capacity and Reliability of a DOCSIS Network**

A Technical Paper prepared for SCTE•ISBE by

**Matthew Olfert**  
Senior Network Architect II  
Shaw Communications  
2728 Hopewell Place NE, Calgary, AB T1Y 7J7  
+1 (403) 538-5210  
Matthew.olfert@sjrb.ca

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
Building the Mindset .....	4
DOCSIS Network Configuration Options .....	4
1. Downstream Channel Configuration .....	4
1.1. Preparing for your first OFDM deployment .....	5
1.2. Deployment Options for OFDM .....	5
1.2.1. Channel Width (Start and End Frequencies) .....	5
1.2.2. Subcarrier Spacing .....	5
1.2.3. Primary Channel Capability .....	5
1.2.4. PLC Frequency Selection .....	6
1.2.5. Next Codeword Pointer .....	6
1.2.6. Cyclic Prefix .....	6
1.2.7. Downstream Profiles .....	6
1.2.8. Exclusion Bands .....	6
1.2.9. Variable Bit Loading .....	7
1.3. Legacy Downstream Channel (SC-QAM) Configuration .....	7
1.3.1. Primary and DSG Channel Reduction .....	8
2. Upstream Channel Configuration .....	8
2.1. Modifying Cable Plant Split .....	8
2.2. Deployment of OFMDA .....	9
2.2.1. Spectrum Choice .....	9
2.2.2. Channel Width .....	9
2.2.3. Interval Usage Code .....	9
2.2.4. Exclusion Bands .....	9
2.2.5. Variable Bit-loading .....	10
2.3. Legacy Upstream Channels (SC-QAM) .....	10
2.3.1. Channel Width .....	10
2.3.2. Modulation Profiles .....	10
2.3.3. Dynamic Upstream Configuration Changes .....	10
3. Global Configuration Options .....	11
3.1.1. Cable Modem Load-balancing .....	11
3.1.2. Transitional MAC Scheduler .....	12
4. Cable Modem Distribution .....	14
5. Data Analysis Applications .....	15
5.1. Proactive Network Maintenance .....	15
5.2. Profile Management Application .....	17
DOCSIS Network Deployment Example .....	17
1. Starting Conditions .....	17
2. Deployment Senario .....	18
3. Deployment Plan .....	18
4. Ending Conditions .....	19
Conclusion .....	20
Abbreviations .....	21
Bibliography & References .....	21

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - OFDM Channel Spectrum with an Exclusion Band .....	5
Figure 2 - OFDM Channel Exclusion Band .....	7
Figure 3 - Variable Bit Loading Example .....	7
Figure 4 - Return Plant Upgrade Gains .....	8
Figure 5 - Traffic Utilization without using load-balancing .....	12
Figure 6 - Traffic Utilization after load-balancing .....	12
Figure 7 - MAC Scheduler Differences – Single CM .....	13
Figure 8 - MAC Scheduler Differences – 100 CMs with 1 DOCSIS 3.1 CM (Mbps) .....	13
Figure 9 - MAC Scheduler Differences – 100 CMs with 1 DOCSIS 3.1 CM (%) .....	14
Figure 10 - Impacts of CM Distribution of New 85 MHz of Return Plant .....	15
Figure 11 - In-channel Frequency Response Showing An Impairment Caused By In-home Drop Amplifier In A 85 MHz Return Plant .....	16
Figure 12 - Time Domain Reflectometry Showing A In-home Issue .....	16

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Minimum Guard Band for A OFDMA Exclusion Band .....	10
Table 2 - Recommended SNR Levels for Modulation Type .....	10
Table 3 - Example of Dynamic Upstream Configuration States .....	11
Table 4 - DOCSIS Device Configuration – Starting Conditions .....	17
Table 5 - DOCSIS Device Configuration – Post Configuration .....	20

# Introduction

With the arrival of DOCSIS® 3.1 operators have a fantastic new asset in their arsenal to provide industry leading broadband experiences. However, with great power comes great responsibility (and complexity). As we build and upgrade our DOCSIS networks, we need to also take the time to maximize their configurations to effectively utilize the new advances in technology, while also increasing resiliency and reliability. To get the most out of a DOCSIS network today, operators need to balance capacity and reliability to bring high-quality customer experience and services. As you maximize the DOCSIS network, you will build a strong foundation. With a strong foundation you can provide a solid framework for future DOCSIS technologies and services. This framework will ease future deployments to improve services.

## Building the Mindset

Before you start the process of maximizing your DOCSIS network, you need to build a mindset of capacity building and customer experience. Understanding the future capacity goals for the next 2-3 years can be a huge benefit. Coming in with an understanding of hardware upgrades, plant upgrades, cable modem deployment plans, future packaging, network traffic CAGR, licensing and capital plans provides the information necessary for making strategic choices.

Upgrades, CCAP hardware or Cable Plant, will enable new features or capabilities that can provide solutions for capacity needs. With DAA, FMA, and FDX significant hardware and plant upgrades are coming. These will come with high capital costs and will take time to deploy. Understanding the timelines for upgrades can provide insights for current network changes.

Cable Modem deployment plans need to include the roadmap for the continual upgrade cycle to the latest cable modem version. As an example, plans to upgrade your return plant to 204 MHz should also include a plan to start pre-seeding high-split capable cable modems years in advance. This will provide a higher initial gain when the upgrade takes place. In addition, old generations of CPE may have compatibility issues with a change of return plant and have the need to be removed prior.

Understanding CAGR, or traffic growth, and future customer packages will help operators recognize what future congestion and overhead needs will be. With this forecast, you need to build in the remaining costs of licensing. With all this information you can prioritize the next changes in your network.

Finally, you must always keep customer experience top-of-mind. Beyond the congestion risks, this also includes other performance indicators like latency and packet loss risks.

## DOCSIS Network Configuration Options

There are many options for configuration of the DOCSIS access network. I will discuss these options in five sections: downstream, upstream, global, cable modem distribution, and data analytics applications.

### 1. Downstream Channel Configuration

With OFDM channel deployment in full swing, many lessons have been learned during these deployments. From personal experience, there are many items to consider. Also, we have additional work we can proceed with on the SC-QAM channels to aid with the transition to DOCSIS 3.1 capable CM.

## 1.1. Preparing for your first OFDM deployment

Before you turn on your first OFDM channels, you need to ensure your operational staff is ready for this new technology. Do they have meters, training, and other tools need to troubleshoot impairments within this carrier? A prepared staff will provide you great feedback to the success of your deployment.

## 1.2. Deployment Options for OFDM

The CableLabs specification provides many options for OFDM configurations. CCAP vendors have been continuously adding feature through their firmware upgrades over the past 2-3 years, but it is not yet in a full-feature set. Due to this limitation, you may have access to some or all these features.




Figure 1 - OFDM Channel Spectrum with an Exclusion Band

### 1.2.1. Channel Width (Start and End Frequencies)

As you choose the starting and ending frequency, you will need to be aware of all plant attributes in between these end points. Also, be aware for every 6 MHz you deploy of OFDM may impact your highest output level of your downstream port on your CCAP chassis. The minimum size of an OFDM channel is 24 MHz with a maximum channel size of 192 MHz.

### 1.2.2. Subcarrier Spacing

Subcarrier spacing can be modified between 25 and 50 kHz. This choice provides one side that favors efficiency or one side that support more robustness to impulse noise. I recommend 50 kHz to gain resiliency over capacity until you understand spectrum performance.

### 1.2.3. Primary Channel Capability

Primary channels are utilized for cable modem registration, station maintenance, and CM STATUS messages. For OFDM, the DOCSIS specification provides additional CM STATUS messages that are critical for DS Profile switching. If CM are unable to communicate their CM STATUS message reliably, they will be unable to manage the change RF plant conditions that can occur with 192 MHz over one

channel. Another downside with a primary channel is that cable modems do not impair it until they can no longer receive and replay to station maintenance. This requires the cable modem register, which can cause a constant cycling cable modem since cable modems prefer to go back to their previous primary channel. Most OFDM deployments happening in the higher spectrum which likely has not been used for cable modems services before. This will bring unknown states about the reliability of this spectrum which has an add risk of roll-off. With this is mind, you may want to disable the primary channel capability and remain relying on SC-QAM channels for primary channel functions.

#### **1.2.4. PLC Frequency Selection**

The Physical Layer Link Channel (PLC) is a requirement for a CM to receive its channel information. This channel selection is critical for CM's usage of the OFDM channel. From my testing we have found that the PLC carrier can be place almost anywhere in the OFDM channel. It is critical that the PLC remains free of noise. Be aware of recent 600 and 700 MHz wireless spectrum deployments when you are planning this channel placement.

#### **1.2.5. Next Codeword Pointer**

The next codeword pointer (NCP) is a critical piece for the cable modem to perform forward error correction within the OFDM channel. You can choose it to modulation at QPSK, 16-QAM, or 64-QAM. It should be modulated below your lowest OFDM modulation to keep this critical data clean.

#### **1.2.6. Cyclic Prefix**


This setting allows for greater protection from micro-reflection in your cable plant. The higher the sample rate the more protection you will receive. This will cause a loss of capacity as you increase the samples. Depending on your placement of the OFDM channel you may want greater protection from micro-reflections. From our experience 2.5  $\mu$ sec cyclic prefix (512 samples) provides enough protect for majority of deployments.

#### **1.2.7. Downstream Profiles**

Downstream profiles are your greatest assets to increase capacity and resiliency of an OFDM channel. With multiple profiles, you can gain diverse groupings of different capabilities for each fiber-node. With downstream profile A, you want to select the minimum modulation you want for a flat OFDM profile. Typically, you will want to run 256-QAM if your currently SC-QAM channels are running the same. Any additional downstream profiles should be utilized to address noise or additional capacity needs. Based on your average downstream MER data from the fiber-node value you can assign higher level of modulation to additional DS profiles. Ideally analysis from PMA (Profile Management Application) would give the best DS profile configuration for maximized capacity.

#### **1.2.8. Exclusion Bands**


In order to maximize your OFDM channel width you may have to build it around existing SC-QAM channels. The use of exclusion bands allows you to create gaps within the spectrum of the OFDM channel. Be aware that you will be required to pad 1 MHz on each side of the exclusion band. Also, the OFDM channel must have 22 MHz or greater of contiguous spectrum before the first exclusion. Exclusion bands are great, but they do cause overhead to the full potential of the OFDM carrier. Your long-term strategy should include ways to not have exclusion bands within your OFDM channels.



**Figure 2 - OFDM Channel Exclusion Band**

### **1.2.9. Variable Bit Loading**

Once you have your DS profiles setup you can further adjust the spectrum by changing the modulation for parts of the OFDM channel. If you have a noisier section of frequency you can lower the modulation in the section only in place of change the whole OFDM channel. This can be very useful to mitigate roll-off or other channel impairments without sacrificing the capacity of the good spectrum of your OFDM channel.



**Figure 3 - Variable Bit Loading Example**

## **1.3. Legacy Downstream Channel (SC-QAM) Configuration**

For the first time in the history of DOCSIS we have created a channel type in the downstream that is not backwards capable. We have now started the transition of a DOCSIS network that is working towards the removal of old SC-QAM channels. This transition will take many years but will result in higher capacity per MHz. You can use the below features to help you transition capacity and to lengthen the timeframe of SC-QAM capacity without adding more channels.

### 1.3.1. Primary and DSG Channel Reduction

Primary channels are necessary for all cable modem registration, but we don't require a primary channel on every SC-QAM channel. With the movement towards bonding capable CMs has provided us the opportunity to remove this service from some channels. By removing this feature from a DOCSIS channel, you can gain additional capacity by reducing the overhead cause by cable modem registration, station maintenance, and CM STATUS messages. In addition to primary channel reduction, we can also reduce the DSG services to the primary channels. We have seen a gain of ~5% capacity (1.84 Mbps) on a SC-QAM channel that has the primary and DSG channel features disabled. This value will change vendor-to-vendor I have seen about 1-2 Mbps on different platforms or configurations.

If you are moving towards a R-PHY DAA solution you may already have noticed a requirement to reduce your primary channels. My vendor is only offering five primary channels for a fiber-node. My experience with integrated CCAP platforms a reduction to eight primary channels can easily be supported.

## 2. Upstream Channel Configuration

Upstream channel configuration has gained the option of OFDMA channel type. A lot of companies are also exploring Mid (5-85 MHz) and High (5-200 MHz) split return cable plants. In addition to these changes we also have some great upstream features to mitigate noise issues and maintain high upstream capacities.

### 2.1. Modifying Cable Plant Split

Gaining higher splits of the forward and return spectrum can enable a powerful way to expand your upstream capacity. In the 5-42 MHz plant, you have about 27 MHz of usage upstream frequency but when you expand to 85 MHz you gain an additional 43 MHz of usable spectrum. This provides 159% more upstream spectrum for DOCSIS capacity. This also enables new spectrum for OFDMA which enables higher capacity per MHz than SC-QAM. Additionally, high-split will bring an addition of 119 MHz of capacity (or a 600% gain). With each split change, be very aware of all equipment to ensure it supports the forward and return plant changes.




Figure 4 - Return Plant Upgrade Gains



## **2.2. Deployment of OFMDA**

OFDMA channels provide higher orders of modulation and less channel spacing than the traditional SC-QAM channel. This provides high upstream capacities and high noise resiliency. The challenge for most operators will be where to place this new carrier since is not backwards capable. In addition, likely all available upstream spectrum of already taken by SC-QAM channels to meet current capacity needs. I expect most operators will wait for return plant upgrades before implementing an OFDMA channel.

### **2.2.1. Spectrum Choice**

First you need to understand where you are going to place the OFDMA carrier before you look at options for this channel. There are four possible choices:

1. “Dirty” Deployment – Deploying OFDMA in the 5-20 MHz. This would provide more MHz of spectrum used by CMs, but this would be at low modulations.
2. SC-QAM Sacrifice Deployment – Turn off existing SC-QAM channel(s) to free up spectrum for OFDMA. This would provide more capacity for DOCSIS 3.1 CMs but would sacrifice DOCSIS 3.0 or later CMs capacity. CM distribution would have a large factor to how many channels you can sacrifice.
3. High Frequency Deployment – Deploying OFDMA in new high frequency is the ideal way to deploy this new channel. This provides an addition to your current upstream deployment. In order to deploy this way, cable plant split must be modified to have more upstream frequency capacity (85 or 204 MHz).
4. Hybrid Deployment – This deployment would include multiple options from the previous three choices.

### **2.2.2. Channel Width**

The minimum size for a 2K FFT in 10 MHz that can be a challenge to find in a 5-42 MHz plant. You can expand this channel size up to 95 MHz of usable spectrum with 500 kHz of guard bands on each side.

### **2.2.3. Interval Usage Code**

Interval Usage Code (IUC) allows us to manage multiple modulation levels within the same OFDMA modulation profile. You set IUC 13 to the lowest modulation you want to run within the OFDMA channel. Then you can add additional IUCs to offer multiple modulation levels within the same channel similarly to DS profiles for OFDM channels.

### **2.2.4. Exclusion Bands**

We can also utilize exclusion bands for OFDMA channels. This can assist with managing OFDMA spectrum within your SC-QAM channels. This allows you to create gaps in the OFDMA band to fit other channels within. Be aware the guard band size changes based on the SC-QAM channels you have next to the OFDMA channel.

**Table 1 - Minimum Guard Band for A OFDMA Exclusion Band**

<b>SC-QAM Channel Width</b>	<b>Minimum Guard Band</b>
1.6 MHz	550 kHz
3.2 MHz	300 kHz
6.4 MHz	0 kHz

### **2.2.5. Variable Bit-loading**

This provides a method to deploy OFDMA channel in different quality of spectrum and maintain maximum capacity. This becomes very useful when running OFDMA in the noisier areas of the upstream spectrum with some clean spectrum. It has all the same benefits that variable bit-loading offers to an OFDM channel.

## **2.3. Legacy Upstream Channels (SC-QAM)**

The current SC-QAM channels are still very much important to offer upstream capacity for DOCSIS 3.0 or later CMs. We have far less options than the new OFDMA channels but if configured correctly you can offer the capacity needed for legacy CMs.

### **2.3.1. Channel Width**

There are three common channel widths for SC-QAM channels: 1.6, 3.2, and 6.4 MHz. To maximize capacity the wider the channel the better. The downside is the noise resistance decreases with a wider channel. If you deploy 3.2 MHz wide channels you can fit two inside the same space as one 6.4 MHz wide channel. This will increase your costs and increase the maximum TX level for your CMs.

### **2.3.2. Modulation Profiles**

There are also three common modulation orders: QPSK, 16-QAM, and 64-QAM. 64-QAM provides the highest level of capacity for SC-QAM channels; OFDMA is required to modulate higher. The higher the modulation level the greater capacity but the higher required SNR level as well.

**Table 2 - Recommended SNR Levels for Modulation Type**

<b>Modulation Type</b>	<b>Recommended SNR Level</b>
64-QAM	27
16-QAM	23
QPSK	18

### **2.3.3. Dynamic Upstream Configuration Changes**

Dynamic upstream configuration change feature goes by different names depending on your vendor of your CMTS/CCAP. This feature allows you to automatically change upstream channel configuration based on RF performance thresholds. Typically, you have access to average SNR level, Correctable FEC errors, and Uncorrectable FEC errors as thresholds. Based on these thresholds, you may want to change modulation, channel width, or even change the center frequency. By using this feature, you can quickly

mitigate a customer impacting noise events automatically. This is done at the cost of capacity, but this is better than a loss or degradation of service.

By using this feature, you can configure your upstream channels at the maximum possible setting and have the CCAP lower the modulation and channel width. This feature also comes with traps so you can report on these configuration changes to assist with PNM or incident activities.

**Table 3 - Example of Dynamic Upstream Configuration States**


State ID	Capacity (Mbps)	Modulation	Channel Width	Center Frequency	Degradation Trigger ID	Improvement Trigger ID
1	30	64-QAM	6.4	Default	1	N/A
2	20	16-QAM	6.4	Default	3	2
3	15	64-QAM	3.2	Default	1	2
4	15	64-QAM	3.2	High	1	2
5	15	64-QAM	3.2	Low	1	2
6	10	16-QAM	3.2	Default	3	2
7	10	16-QAM	3.2	High	3	2
8	10	16-QAM	3.2	Low	3	2
9	5	QPSK	3.2	Default	3	4
10	5	QPSK	3.2	High	3	4
11	5	QPSK	3.2	Low	N/A	4

### 3. Global Configuration Options


Global configuration brings together downstream and upstream features to provide options to run your DOCSIS network better. These features, for the most part, are vendor specific, so I will cover a couple of the bigger features that most vendors have or should have.

#### 3.1.1. Cable Modem Load-balancing

Cable Load-balancing is a powerful tool to ensure traffic on each channel is operating at similar traffic levels as the other channels within the fiber-node. With different configuration of load-balancing you can dramatically change traffic loading between channels. This also supports load-balancing bonded and non-bonded CMs. The closer the channels are at the same traffic levels the more load-balancing you will need to perform. Each load-balancing event comes with a risk of a CM reset event. You need to balance the need of per channel capacity with the risk of a CM reset event.



**Figure 5 - Traffic Utilization without using load-balancing**




**Figure 6 - Traffic Utilization after load-balancing**


### **3.1.2. Transitional MAC Scheduler**

From my experience MAC schedulers try to balance bonded CMs over all their bonded channels. This was a great system for MAC domains that are only SC-QAMs. As you turn up OFDM/OFDMA interfaces with a low distribution of DOCSIS 3.1 capable CMs a different MAC scheduler is needed. On a traditional MAC scheduler, a DOCSIS 3.1 CM will end up consuming SC-QAM bandwidth with lots of remaining OFDM capacity remaining. This becomes a more pressing issue on high utilization on SC-QAM channel within a fiber-node. Moving to a MAC scheduler that prioritizes the DOCSIS 3.1 CMs to use OFDM/OFDMA channels over the SC-QAM channels becomes a better configuration. This MAC


scheduler should ideally be one you can switch to and back to a traditional scheduler on a per fiber-node's need. With changing CM distribution over the coming years, you may find the need to go back to the traditional MAC scheduler again.



**Figure 7 - MAC Scheduler Differences – Single CM**



**Figure 8 - MAC Scheduler Differences – 100 CMs with 1 DOCSIS 3.1 CM (Mbps)**




**Figure 9 - MAC Scheduler Differences – 100 CMs with 1 DOCSIS 3.1 CM (%)**

## 4. Cable Modem Distribution

With the introduction of new DOCSIS channel types (OFDM/OFDMA) we have created a barrier of non-capability in order to further gain capacity and resiliency. With cable plant upgrades such as mid-split will also create these barriers of non-capability. With upcoming future technologies such as Full-duplex DOCSIS, 1.8 and 3 GHz cable plants will also continue this trend.

At Shaw, we started production testing of 85 MHz return plant in 2014. During this time, Shaw also selected a cable modem for our customer that would support 85 MHz. Shaw didn't immediately deploy 85 MHz in mass deployment, but tests showed the benefit of additional upstream capacity. In 2018, Shaw started our deployment of upstream channel deployment above 42 MHz. Based on the results of traffic offloading to the higher frequency upstream channels we can see a correlation of 5-85 MHz supporting cable modems to effectiveness of traffic offloading.




**Figure 10 - Impacts of CM Distribution of New 85 MHz of Return Plant**

## 5. Data Analysis Applications


CableLabs has been providing data analysis tool designs to provide DOCSIS network designers additional tools to maximize capacity and resiliency of the network. These applications can be a great benefit to you and your customer experience.

### 5.1. Proactive Network Maintenance

Proactive Network Maintenance (PNM) is a powerful tool to identify oncoming negative RF issues. It can also identify current outstanding RF issues. This is a great way to monitor your cable plant from every cable modem in your DOCSIS network. By doing this, you can maintain higher capacity levels due to good plant conditions. PNM commonly looks at pre-equalization, in-channel frequency response, full band capture, and time domain reflectometry data. In DOCSIS 3.1 items like RxMER for OFDM/OFDMA channels have been added to PNM. This addition within PNM drives the support for additional tool called PMA.



**Figure 11 - In-channel Frequency Response Showing An Impairment Caused By In-home Drop Amplifier In A 85 MHz Return Plant**



**Figure 12 - Time Domain Reflectometry Showing A In-home Issue**



## 5.2. Profile Management Application

Profile Management Application (PMA) is a tool to assist with OFDM/OFDMA channel configuration. This tool becomes increasingly more powerful as you are utilizing more profiles (IUC and DS Profile) and variable bit loading features within OFDM/OFDMA channel types. PMA utilizes the new PNM feature of RxMER to assess the capability of each CM on a fiber-node.

With testing at Shaw, we see the benefits of ~40% more capacity with this tool. We are configuring DS profile A as a default flat modulation (256-QAM) for the whole channel. Then every additional DS profile added to support a correlated group of CMs that share the same or similar spectrum characteristics to maximize their capacity for these conditions. The more correlated groups you have, the more DS profiles you require to gain this capacity. Based on our tests we believe DS profile A and seven additional DS profiles per OFDM channel will be needed to optimize your downstream capacity. We also see the need for many variable bit loading sections to further support the capacity of the OFDM channel.

PMA also brings additional channel resiliency into your network like other PNM features. The advantage is that it can be automated instead of a technician being onsite prior to any mitigation. The data from this tool can also provide correlated groups of cable modems with the same channel impairment. This data could also be found within your PNM tool if it performs analysis of the RxMER data too.

# DOCSIS Network Deployment Example

Now that we have covered the available options, let's go over a fictional example of deployment of changes in a DOCSIS network. By going through this exercise, will provide an example of the thought process required to build a design change. It will also demonstrate the additional capacity or resiliency gains you can achieve with these options.

## 1. Starting Conditions

The baseline we will begin our deployment of new options will be the common 24x4 network. This means we will have 24 downstream SC-QAM channels plus four upstream SC-QAM channels. In order to keep this simple, I will display the rest of the setup in a few tables.

**Table 4 - DOCSIS Device Configuration – Starting Conditions**

Feature	Current Setting
<b>Downstream</b>	
SC-QAM Channel Count	24
SC-QAM Channel Modulation	256-QAM
SC-QAM Channel Frequency Range	300-444 MHz
SC-QAM Primary Channel	Enabled on all channels
SC-QAM DSG Service	Enabled on all channels
OFDM Channel Count	0
OFDM Frequency Range	-
OFDM DS Profile Count	-
OFDM DS Profile Modulation	-

Feature	Current Setting
OFDM Subcarrier Spacing	-
OFDM Primary Channel	-
OFDM Cyclic Prefix	-
OFDM PLC Channel	-
OFDM Exclusion Bands	-
OFDM Variable Bit Loading	None
Total Downstream Capacity (PHY Rate)	1029 Mbps
<b>Upstream</b>	
SC-QAM Channel Count	4
SC-QAM Modulation	64-QAM
SC-QAM Channel Width	2 x 6.4 MHz and 2 x 3.2 MHz
SC-QAM Center Frequencies	21.8, 25, 30.1, 36.8 MHz
OFDMA Channel Count	0
OFDMA Frequency Range	-
OFDMA Interval Usage Code(s)	-
OFDMA Exclusion Bands	-
OFDMA Variable Bit Loading	-
Total Upstream Capacity (PHY Rate)	92 Mbps
<b>Global</b>	
Dynamic Upstream Configuration Changes	Disabled
Cable Modem Load-balancing	Enable
MAC Scheduler	Default

On top of the DOCSIS network configuration, it should also be stated the current RF plant capabilities. We will start with are 750 MHz plant with a 5-42 MHz return. The current cable modem distribution for the internet services is 30% DOCSIS 2.0, 65% DOCSIS 3.0, and 5% DOCSIS 3.1. Our current top tier consumer product is 500 Mbps download with 30 Mbps upload speeds.

## 2. Deployment Senario

It's important to set a deployment goal to build towards. For this fictional deployment example, we are going to build a new top tier package. The product team has stated a desire to achieve at least 1 Gbps download with a 50 Mbps upload.

## 3. Deployment Plan

The first thing we need to maximize our capacity would be a DOCSIS 3.1 cable modem for this package. This will allow us to utilize OFDM and OFDMA channels to increase our capacity. If we don't use DOCSIS 3.1 features yet it would build a cable modem pool for this in the future. The first requirement will be a DOCSIS 3.1 cable modem.

We could also look at cable plant upgrades to assist with additional spectrum. Looking at our current video deployment, we have spectrum available between 444-492, and 696-750 MHz. Also, on the return spectrum we have spectrum below 20 MHz available as well. Since we have unused spectrum for use, let's hold off at looking at cable plant upgrades.

In order to offer a 1 Gbps of speed it would be ideal to have 2 Gbps or more of capacity on the downstream. Since we have a lot of CMs that are not on DOCSIS 3.1 cable modems expanding SC-QAM would be a good idea first. This would provide additional growth for these cable modem and offer additional capacity for DOCSIS 3.1 CMs. We have eight channels worth of frequency right by our 24 SC-QAM channels. If we added these eight channels, it would increase our PHY rate by 343 Mbps to a total of 1372 Mbps. Now that leaves us with 54 MHz of free spectrum between 696-750 MHz. If we deploy OFDM in this space which is 54 MHz with 2 MHz of guard band leaving us 52 MHz of data spectrum. 52 MHz @ 256-QAM would provide 320 Mbps more. This would bring us to a total of 1692 Mbps which still doesn't provide us enough. Looking at the higher modulation orders on our OFDM channel will provide more capacity. Setting the channel to 1024-QAM bring the PHY rate to 400 which only get us 80 Mbps more (or a total of 1772 Mbps). Then going up to 4096-QAM get another 80 Mbps or 1852 Mbps in total. To achieve the capacity required for this package, additional spectrum is needed.

Looking at upstream, capacity let's examine what we can get out of it. The spectrum below 20 MHz is not as clean as our other channels. If we changed over to 4 x 6.4 MHz wide channels @ 64-QAM, we would have 123 Mbps in total capacity. This may be enough to services 50 Mbps service. Another option will be OFDMA, for which we need at least 10 MHz of spectrum. One option would be to turn down the lowest 3.2 MHz carrier and go below 20 MHz. This would deploy an OFDMA channel between 13.4 - 23.4 MHz @ 256-QAM to get the 10 MHz. This channel would provide 39 Mbps of capacity bring us to a total of 131 Mbps. Since we have reduced DOCSIS 3.0 upstream capacity the majority of our cable modems have lost capacity with an OFDMA channel.

Additional information was needed to determine the direction. Based on current upstream traffic levels, we can't lose SC-QAM capacity on the upstream so OFDMA is out as a choice. But a conversation with our video team has yielded an additional spectrum. After a channel reorganization it has yielded an additional 30 MHz of spectrum alongside of our possible OFDM channel. This would bring the OFDM capacity to 707 Mbps @ 1024-QAM. This bring the total of 32 SC-QAM + 82 MHz OFDM to 2079 Mbps. With the 123 Mbps of upstream capacity we could support a 1 Gig / 50 Mbps service with this configuration.

We could gain additional efficiency with a primary downstream channel reduction. This would reduce the DOCSIS overhead on the channels with this removed. Looking at the traffic levels of the DOCSIS 2.0 cable modems we can likely go down to 12 primary channels. This would provide 30 Mbps of data for customers bonded-capable CMs.

The requirement for dynamic upstream configuration changes on the upstream, where going below 20 MHz become increasingly important. Also, any upstream channel can benefit from this type of autonomous configuration system. By utilizing trigger events of SNR, uncorrectable, and correctable FEC errors, we can detect when RF impairments are occurring. Then by reducing modulation, channel-width are good first steps. Once the 6.4 MHz channel has been reduced to 3.2 MHz channel width, the center frequency can be moved around to avoid the noise issue. This configuration will also improve the customer experience.

## **4. Ending Conditions**

After the above deployment plan is implemented, the new DOCSIS device configuration would be the following:

**Table 5 - DOCSIS Device Configuration – Post Configuration**

Feature	Current Setting
<b>Downstream</b>	
SC-QAM Channel Count	32
SC-QAM Channel Modulation	256-QAM
SC-QAM Channel Frequency Range	300-492 MHz
SC-QAM Primary Channel	Enabled on all channels
SC-QAM DSG Service	Enabled on all channels
OFDM Channel Count	1
OFDM Frequency Range	666-750 MHz
OFDM DS Profile Count	1
OFDM DS Profile Modulation	1024-QAM
OFDM Subcarrier Spacing	50 kHz
OFDM Primary Channel	No
OFDM Cyclic Prefix	2.5 $\mu$ sec
OFDM PLC Channel	711 MHz
OFDM Exclusion Bands	None
OFDM Variable Bit Loading	None
Total Downstream Capacity (PHY Rate)	2079 Mbps
<b>Upstream</b>	
SC-QAM Channel Count	4
SC-QAM Modulation	64-QAM
SC-QAM Channel Width	4 x 6.4 MHz
SC-QAM Center Frequencies	16.7, 23.4, 30.1, 36.8 MHz
OFDMA Channel Count	0
OFDMA Frequency Range	-
OFDMA Interval Usage Code(s)	-
OFDMA Exclusion Bands	-
OFDMA Variable Bit Loading	-
Total Upstream Capacity (PHY Rate)	123 Mbps
<b>Global</b>	
Dynamic Upstream Configuration Changes	Enabled
Cable Modem Load-balancing	Enable
MAC Scheduler	Default

This configuration would be able to provide internet speeds of 1 Gbps download with a 50 Mbps upload. In addition, we gain resiliency on our upstream with dynamic upstream configuration changes. With the introduction of an OFDM channel, modems on these channels should have more reliable service as well.

## Conclusion

As you upgrade and deploy new features on your DOCSIS network, you will gain capacity and the ability to add resiliency features. Ensure you have the correct mindset to design the DOCSIS network you want.

As MSOs move towards more DOCSIS 3.1 services, the need to manage the transition to OFDM/OFDMA carriers are a huge benefit. Features like transitional MAC scheduler ease the pressure of this transition. PNM/PMA applications enable resiliency and the ability to maximize our OFDM/OFDMA capacity.

## Abbreviations

CAGR	Compound Annual Growth Rate
CCAP	Converged Cable Access Platform
CM	Cable Modem
CPE	Customer Premise Equipment
DAA	Distributed Access Architecture
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DSG	DOCSIS Set-top Gateway
FDX	Full-duplex DOCSIS
FEC	Forward Error Correction
FMA	Flexible Mac Architecture
FFT	Fast Fourier Transform
ICFR	In-Channel Frequency Response
IUC	Interval Usage Code
kHz	Kilohertz
MAC	Media Access Control
Mbps	Megabits Per Second
MHz	Megahertz
MSO	Multiple System Operator
NCP	Next Codeword Pointer
OFDM	Orthogonal Frequency-division Multiplexing
OFDMA	Orthogonal Frequency-division Multiple Access
PHY	Physical
PLC	Physical Link Channel
PMA	Profile Management Application
PNM	Proactive Network Maintenance
QPSK	Quadrature Phase Shift Keying
RxMER	Receive Modulation Error Ratio
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SNR	Single-to-Noise Ratio
TDR	Time Domain Reflectometry
TX	Transmit

## Bibliography & References

CM-SP-PHYv3.1-I16-190121: *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1 Physical Layer Specification*; CableLabs

CM-SP-MULPIv3.1-I18-190422: *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1 MAC and Upper Layer Protocols Interface Specification*; CableLabs

*E6000 CER I-CCAP Release 7.0 User Guide Preliminary v1.0; Arris*

NimbleThis PNM Tool

CableLabs PMA Tool

# **The Role of Lean in Shaw**

## **Our Technical and Operational Journey**

An Operational Practice prepared for SCTE•ISBE by

**Noé Morales**

Director - Lean Enterprise Office  
Shaw Communications Inc.  
630 3rd Avenue SW Calgary, Alberta, Canada, T2P 4L4  
1(403) 234-6203  
Noe.Morales@sjrb.ca

**Aston Fenby**

Portfolio Manager - Lean Enterprise Office  
Shaw Communications Inc.  
630 3rd Avenue SW Calgary, Alberta, Canada, T2P 4L4  
1(403)781-5185  
Aston.Fenby@sjrb.ca

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Abstract.....	4
Content .....	4
1. Lean Introduction .....	4
2. Shaw Beginnings .....	7
3. Learnings.....	7
4. Pivot Points.....	8
5. Results .....	8
6. Journey to a Modern Shaw.....	11
7. Results .....	12
8. Learnings.....	13
9. Lean Enterprise Office Introduction .....	13
10. Strategic Cultural Plan.....	14
11. Tactical Processes.....	15
11.1. Tactical - Focus Areas .....	16
11.2. Tactical - Services Breakdown .....	17
11.3. Tactical – Yellow Belt Accelerator .....	18
11.4. Tactical – In-house Green Belt Training .....	19
11.5. Tactical – Kata Vision and Target Condition Identification .....	19
11.6. Tactical – Benefits Realization .....	19
12. Results .....	21
Conclusion .....	22
Abbreviations.....	23
Bibliography & References .....	23

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – TPS Key Individuals.....	5
Figure 2 – Principles of Lean .....	5
Figure 3 – CPE Fulfillment - Batch Processing.....	8
Figure 4 – CPE Fulfillment – Single Piece Flow.....	9
Figure 5 – 2014 Results.....	9
Figure 6 – Lean Maturity Index .....	11
Figure 7 – Opportunity Matrix.....	12
Figure 8 – Network Build Process Results .....	13
Figure 9 – Hub and Spoke Model.....	14
Figure 10 – Focus Areas – High Level .....	17
Figure 11 – Yellow Belt Accelerator Program .....	19
Figure 12 – Benefit Syntax Example .....	20
Figure 13 – LEO Results – F19.....	21
Figure 14 – Training and Engagement – F19 Totals.....	22



## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Three-Pillar Approach - Defining Principles .....	15
Table 2 – Focus Area Vision Statements .....	16
Table 3 – Services Breakdown .....	18

# Abstract

Cable operators are facing unprecedented levels of competition and technology advancements. Major shifts in technology such as FDX, Fibre Deep, IoT, and 5G/FTTH, competition will only accelerate. To adapt to this competitive and fast paced environment, operators must take effective steps to increase efficiency, identify waste, and empower teams to resolve challenges at all levels of the organization. This paper will present a methodology for identifying, prioritizing and resolving inefficiencies as well as an approach to foster a Lean culture at all levels of the organization.

Lean methodologies are an invaluable tool in identifying and remediating inefficient and wasteful processes in Cable operations. However, the deployment of Lean must be carefully planned and executed. The paper will outline how Lean was deployed at Shaw to improve the efficiency and effectiveness of network build processes and teams, and the outcomes that were achieved to date.

To make Lean truly sustainable and effective though, it must also be engrained in the culture of the organization. Benchmarks from various companies have shown the adoption of new methodologies such as Lean and Agile having varying degrees of success. To increase the effectiveness of adoption, we approached implementation using a three-pillar approach of: People, Purpose, and Process. The pillar approach was augmented with a tactical and strategic lens across the enterprise to maximize cultural adoption by leadership and increase operational improvements. Techniques such as Lean Accelerator, Adaptive Training, Cultural Pull Methodology, Servant Leadership; and avoidance of “One Size Fits All” using an Adaptive Standardization technique has improved exposure to Lean.

Our results demonstrate that a highly focused implementation of the three-pillar approach increases adoption of Lean at the enterprise level. Additionally, it brings a proactive method to identifying Lean process improvement opportunities with rapidly evolving technologies and operational practices.


# Content

## 1. Lean Introduction

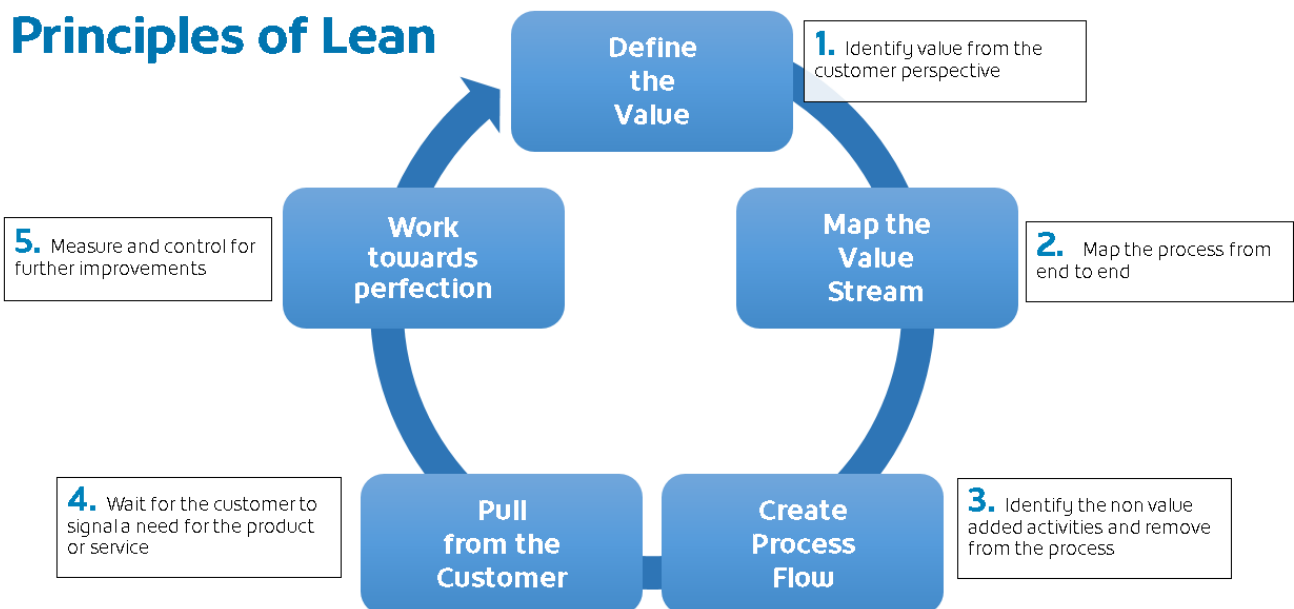
Lean focuses on ongoing process improvement with the obsession of eliminating waste. Lean is a systematic, elegant approach which can be applied to any process that enables you to change your workflow for the better. Everything from how you make your morning coffee, to your daily emails, to how you build a rocket.

Lean at its heart is about respect for our employees, it’s about delivering a valuable product or service with quality, and managed costs, in an efficient manner; always keeping the customers’ needs at the top of our minds. Lean focuses on removing waste and non value-added activities by improving processes, standardization, and fostering problem-solving capabilities at all levels of an organization.

Lean originates from the Toyota Production System (TPS). Several individuals were instrumental in the development of TPS. Some of the key individuals were:



**Figure 1 – TPS Key Individuals**



**Figure 2 – Principles of Lean**

The five Lean principles founded by James P. Womack and Daniel T. Jones [1] encourage the practice of continuous improvement. They provide a framework for creating an efficient organization by focusing on creating a better flow in work processes and building culture.

The principles break out into various tools and techniques to help better understand the value for the customer, the process (either high-level or detailed), and demand. The final step of working towards perfection helps to inspire teams to look deeper into each process after an improvement has been made.

Some of the tools used within Shaw are:

Gemba walks - Gemba is a Japanese word translated to; "the real place" or "where the work happens". Gemba walks allow management to "go & see", show respect for the workforce, and ask "why" while walking the process. The objective is to understand the people, purpose, and process. The more we know about the situation the better we can find solutions to solve with the team. [2]

Voice of the Employee - Voice of the Process/Employee (VOP/E): The voice of the process is the expression or wastes, rework or any other observed challenges or issues experienced by the people within a process.

Kaizen event - Kai in Japanese is translated to "change" and Zen means "for the better". In other words, "change for the good", "continuous improvement". Kaizen events are team led workshops with a set of goals to be achieved to optimize an specific area/process. Process owners and management come together and participate to encourage the team on their future state goals. A typical kaizen event will take between 3 - 5 days. This method allows for the team to experience incremental improvements as they optimize their workflow.

"Sort", "Set In order", "Shine", "Standardize" and "Sustain" (5S) - 5S is a five-step process that leads to workplace organization. Each step in the process is named using a word that starts with the letter "S".

The five steps are:

- Sort (organization, clearing) - Separate what is needed and what is not needed and keep only those things that are needed in the workplace.
- Set in Order (orderliness, configure, simplify) - A place for everything and everything in its place.
- Shine (cleanliness, sweep for abnormalities)- Identify abnormalities by visually sweeping the area.
- Standardize (stabilize, create standards for conformity) - Arrange items that they can be found quickly by anybody.
- Sustain (self-discipline, practice) - Leaders are responsible to sustain the first 4S steps by encouraging the practice of workplace organization.

Leaders are responsible to sustain the first 4S steps by encouraging the practice of workplace organization.

Kata approach – Utilizing the two linked behaviours of the improvement kata and the coaching kata. The improvement kata follows the problem-solving methodology of plan, do, check, act (PDCA) in a repeating cycle with learnings from last target condition applied to the next cycle always heading towards the vision. The coaching kata is a mentor/mentee approach to guiding the student in the right direction but not providing solutions. The most important part of the kata approach is the repetition of the behaviours which helps to build culture.

## 2. Shaw Beginnings

Lean was introduced at Shaw's National Distribution Center (NDC) in August 2013 by a small team sponsored by the Chief Procurement Officer (CPO) and staffed with one Lean Practitioner and two Lean aware individual contributors. The primary focus was to work with the Reverse Logistics Department to improve efficiencies.

For the remainder of 2013, until February of 2014, the team focused on developing staff training and work stream analysis. Buyers, Managers, and Project Managers were placed in Green Belt training with the goal of expanding Lean skills and leadership to work through selected areas within the NDC. Warehouse Technicians were provided with Lean Yellow Belt training coined "The Shaw Way" to increase awareness of waste and learn problem solving skills.

March to August 2014 saw moderate success in the four selected projects with a heavy lift of resources from the small Lean team. A reduction of Total Process Lead Time (TPLT) averaging 15% was achieved across four selected work streams. However, in following months TPLT began to increase, leveling out at a 6% reduction after six months.

## 3. Learnings

While initial results were positive at 15%, the reduction to 6% after two months was not fully understood, requiring a deeper analysis. Data collected through Lean tools (such as Value Stream Maps, Spaghetti Diagrams, and Time/Motion Studies) provided excellent historical context, which would not have been available without the implementation of Lean. However, the introduction of leader Gemba walks and voice of the employee sessions provided better insight into the root cause of the 9% loss of efficiency since implementation.

Discoveries included:

- No control plan accountability (Leader Gemba walk)
- Low engagement from warehouse technicians (Voice of the employee)
- Warehouse technician proposals not implemented (Voice of the employee)
- Problems were people-intensive, and resources were already limited with day-to-day operations (Voice of the employee)
- Low buy-in from middle management (Leader Gemba walk)

The five primary discoveries started our journey towards building a culture of Lean. Moving away from "doing Lean" to "being Lean". Our initial culture efforts across the organization followed a Western approach of aligning everyone to think the right (Lean) way. The NDC Lean team embraced the cultural approach described in the article "Lessons from NUUMI" [3]. NUMMI stands for New United Motor Manufacturing Inc. NUMMI was a joint venture between GM and Toyota in the early 1980s. Both companies wanted to learn from each other; as Toyota was interested in getting into the American market and GM wanted to learn more about the Toyota Production System that brought success to Toyota in quality and cost effective ways to produce vehicles. The NUUMI approach focuses on changing how people behave with what they do and increasing employees' trust in management.

## 4. Pivot Points

Implementation of the TPS approach from the “Lessons from NUUMI” learnings required a complete change to the training and implementation. Instead of providing training to front office staff, focus was shifted to the warehouse technicians working with only the Lean tools that would benefit the specific problem areas identified, exclusively from the Voice of the Employee.

Methodology changes included:

- *Right Size Approach to the training model*
  - *"Sort", "Set In order", "Shine", "Standardize" and "Sustain" (5S) Workshops* - Targeting smaller groups in localized areas to help team “see” results in a shorter time frame.
  - *Kaizen Events* - Engaging the Warehouse Technicians in targeted training and hands-on improvements.
- *Culture as an action instead of a label* - Increasing middle management buy-in through on-the-floor process walks with warehouse technicians, and Lean improvement projects becoming part of the management standard work with a sponsorship instead of a directive approach.

## 5. Results

The revised approach to training and culture resulted in significant improvements across the NDC Test & Repair, and Customer Premise Equipment (CPE) fulfillment areas.

Figure 3 below shows the CPE fulfillment area prior to the three-day kaizen event with the warehouse technicians shown in the image. As evidenced in the photo, there is a clear lack of flow in the process, there are multiple stations completing the same task, space is not used efficiently, and large amounts of inventory are on the floor.



**Figure 3 – CPE Fulfillment - Batch Processing**

Figure 4 below showcases the results of the three-day kaizen event within the CPE fulfillment area, all improvements below were designed, implemented, and executed by the warehouse technicians shown above in figure 3. These individuals received training on day one, planned on day two, and executed the improvements on day three.



**Figure 4 – CPE Fulfillment – Single Piece Flow**

Results from the three-day kaizen event shown in figure 3 and figure 4 with the warehouse technicians yielded significant improvements in inventory savings (cost avoidance), through-put of units (efficiency), and reclaimed floor space (cost avoidance). Highlighted below in figure 5



**Figure 5 – 2014 Results**


Following the successes at NDC, Shaw's Regional Distribution Centers (RDC) adopted improvements implemented at NDC and began development of their own Lean programs. The original NDC Lean team moved to a support and reporting role, refining "The Shaw Way" training with vendor augmented classroom training, expanding tools and templates, and implementing a robust 14-point Lean maturity index report for Supply Chain.

The Lean Maturity Index shown in figure 6 below, was used to capture current state of results and activities at a glance through a spider chart format. The higher the number, the better for each monthly period. The chart was posted within the NDC and RDC work areas to highlight progress to our individual contributors. The fourteen highlighted areas were:

1. Management Support (Results)
2. Culture (Results)
3. 5S (Activity)
4. Value Stream Mapping (Activity)
5. Setup Reduction (Activity)
6. Total Productive Maintenance (Results)
7. Pull Systems (Activity)
8. Production Flow (Results)
9. Plant Layout (Results)
10. Standard Work (Activity)
11. Safety (Results)
12. Quality (Results)
13. Workplace Audit (Results)
14. Continuous Improvement (Activity)

Using a combined result and activity based chart helped to paint a picture of our Lean maturity. While the charting was beneficial in our journey, providing us focus areas, there was no vision defined for the end result. In essence, "doing Lean", not "being Lean". For example, culture was graded as a high number (5.0). But this came from the notion of doing Lean just to get a score. Instead of truly understanding that being Lean was the overall outcome through the application of thinking and practicing the philosophy.





**Figure 6 – Lean Maturity Index**

Results from NDC and RDC showed targeted improvements with positive results. However, it was felt that a “Continuous Improvement” culture had not yet manifested. Most improvement activities were still executed under a project-based approach. Additional learnings from the metrics reported through the Lean maturity index showed they provided awareness to management and individual contributors. But, the supporting processes around these 14 tracked points and actionable paths to improvement were not clearly defined.

## 6. Journey to a Modern Shaw


With the continued successes of Lean at NDC (driven by implementation of the pivot points) and adoption of these principles at RDC, Lean methodologies were implemented into Shaw’s wireline network build teams, starting with the creation of a Lean Networks team in November 2017. This team focused on identifying and removing waste within network build processes through targeted training and Black Belt staff augmentation in support of broader strategic objectives.

With the complexities associated with the network build process, initial work focused on gathering a 30,000-foot view of processes. Simply put, gaining a better understanding of what activities each department or team performed, followed by secondary review with upstream and downstream partners. This provided the ancillary benefits of: introducing Lean in a positive manner to departments and teams who had no previous exposure to Lean methodologies, better understanding of work that moved between teams, and gaining top-down and bottom-up support.

From these discoveries, the Lean Network team organized their approach using three workstreams focusing on the strategic objectives of: Node Activity, Network Operations, and targeted Green Belt training which complemented the strategic objectives. Green Belt Training was supported and mentored by the NDC Lean team.

Utilizing the workstream approach allowed for a first level prioritization. However, resource constraints required the Lean Network team to implement additional strategies to allow for further prioritization of opportunities. It was recognized that while all opportunities were important, the ability to translate qualitative comments on “the problems” to a quantitative approach was critical for narrowing the focus to drive actionable improvements.

The adoption of an Opportunity Matrix methodology which allowed for quantitative charting of potential benefits and anticipated difficulty provided a clear agnostic direction to senior leadership on which opportunities to action, aligned to the strategic objectives.

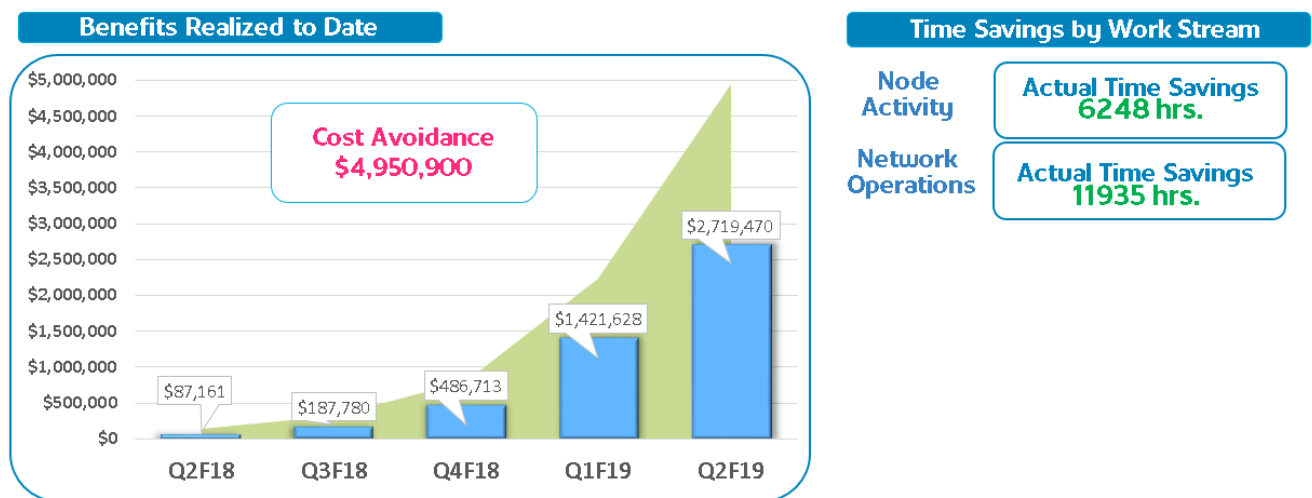


**Figure 7 – Opportunity Matrix**

## 7. Results

A total of 41 strategic network build process initiatives across the node activity and network operations workstreams have been implemented since the second fiscal quarter of 2018 (December 2017). Improvements to the overall network build process have resulted in significant improvements in areas such as: fibre splicing cost consistency, inventory management, reduced tech call volume, distributed access architecture (DAA) installation, truck standardization, move, add, change, delete (MACD) process improvements, optical node radio frequency (RF) module upgrade processes, return band upgrade support processes, and numerous quality and data accuracy improvements.

Figure 8 shows the cumulative savings from Q2 F18 to Q2F19 (December 2017 to December 2018). The optical node radio frequency (RF) module upgrade process was responsible for the majority of the cost avoidance. This Green Belt project in the network operations team reduced the total process lead time on each upgrade from 37 minutes to 13 minutes per unit by updating workstations, improving flow from station to station, and removing unnecessary testing. This improvement is now on its third iteration reducing the need of new equipment purchases on 3 of our optical node lines. The return band upgrade support processes and reduced tech call volume provided the majority of time savings within the network operations workstream, accounting for approximately 6000 of the hours saved.



**Figure 8 – Network Build Process Results**

## 8. Learnings


Integrating Lean principles into the network build process yielded better-than-expected results, which remained steady and did not experience the dip observed during the initial NDC implementation. This was credited to a significantly more robust control plan for each improvement, engagement at the subject matter expert (SME) level, leadership support and sponsorship to implement SME recommendations, and dedicated Black Belt resources allowing for a more robust mentorship of Green Belt projects.

Discoveries included:

- Improvements were project-based, secondary improvements from trained Green Belts were below expectations.
- Second wave of external vendor-based training pressured teams to fill seats, instead of selecting candidates based on strategic vision, leading to long lead times while projects were found.
- Establishing a strategic direction was critical for success. However, a unified vision was missing leading to Black Belts being deployed in a trusted advisor role similar to Business Analysts.
- Results of improvements were heavily highlighted driving pressure to focus on bigger opportunities instead of all opportunities, constantly looking for the “big win”.
- Benefit reporting required significant effort (80+ Hours/Month)
- Discovery activities with larger opportunities took too long, leading to lost sponsorship and unrealized efforts from the team.

## 9. Lean Enterprise Office Introduction

With a shift in the corporate direction to maintain a competitive advantage the Total Business Transformation Office was formed in January 2018. Spearheaded by the CPO, the Lean Enterprise Office (LEO) was also created at this time with the purpose of becoming the central “hub”, or Center of Excellence to the “spokes” like the Lean networks team and existing Lean teams within Supply Chain.



**Figure 9 – Hub and Spoke Model**

With a significant amount of learnings gathered from the previous Lean implementations, the first five months of focus of LEO were used to define the goals, vision, cultural strategy, and tactical processes to launch Lean to the organization. The initial step was to determine a clear set of goals to strive towards.

LEO established the following goals:

- Guide Shaw towards becoming a learning organization using Toyota Production System methodologies (Kata Approach)
- Enhance cultural engagement at the company
- Drive a culture of continuous improvement versus solely project based improvement

To support these goals, the team defined its vision: to empower and inspire everyone at Shaw to live into a continuous improvement mentality by challenging the status quo and always asking, “Is there a better way to do this?”

## 10. Strategic Cultural Plan

The next steps involved creating the strategic cultural plan and tactical processes to increase adoption. This presented unique challenges as current Lean methodologies used within Shaw were showing positive results. However, the current methodologies were tailored to the reverse engineered approach outlined in the Toyota Kata, with a managing by results (MBR) approach instead of a managing by means (MBM) approach. [4] To introduce new behaviors and actions, while maintaining positive momentum and remaining operationally relevant, a significant amount of effort was invested in the strategy of culture and tactical processes.

With the experience of our learnings and benchmarking of other organizations we realized the importance of a strong cultural message. Cultural adoption requires that the tactical processes implemented to achieve improvements reinforce the culture.

To support our cultural plan, LEO implemented the three-pillar approach (widely used in other organizations) with a focus on what that meant for LEO in terms of delivery and anticipated outcomes for the departments and teams.

**Table 1 - Three-Pillar Approach - Defining Principles**

Area	People	Purpose	Process
Lean Enterprise Office ( <i>Delivery</i> )	Servant Leadership	What do our customers want?	Adaptive Governance - Guardrails to Train Tracks
	Support Focused		Adaptive Training
	Cultural Pull Methodology	Clear brand message delivery: “Is there a better way to do this?” and “Always asking Why?”	Clear step by step approach to improvements
Departments & Teams ( <i>Outcomes</i> )	Established Trust	Keeping customer needs at the top of our mind	Tools to fit the opportunities
	It’s all about you!		Accessible Training
	Respect for Employees	Consumable brand message	Developing Problem solving capabilities

The primary purpose of the three-pillar approach was to establish a brand for Lean within Shaw and establish clear messaging to drive our Lean culture. As we had learned earlier in our journey of making culture an action instead of just a label, we applied the same methodology to our three-pillar brand with well-defined delivery actions expected from the LEO team and outcomes to strive towards with our customers instead of just words in a presentation, as shown in table 1 above. This actionable brand provided a tie-in and benchmark to the tactical processes which compliments the three pillars.

## 11. Tactical Processes

Prior to development of our tactical processes it was important to evaluate our defined goals against identified challenges to achieve the most value from a cultural and improvement delivery perspective. Each tactical process developed reinforces the cultural adoption by individual contributors and teams.

We determined the challenges to address were:

- Maintaining strategic big-picture improvements while executing smaller improvements quickly at an enterprise level.
- Implementing a Mentor/Mentee to Plan, Do, Check, Act (PDCA) with limited resources.
- Aligning the organization to a benefit realization language that shows operational relevance but highlights a MBM approach.
- Increasing waste identification and Lean engagement and excitement at all levels of the organization enterprise wide.
- Black Belts being deployed in a trusted advisor role similar to Business Analysts.
- Discovery activities with larger opportunities took too long, leading to lost sponsorship and unrealized efforts from the team.
- Second wave of external vendor-based training pressured teams to fill seats, instead of selecting candidates based on strategic vision, leading to long lead times while projects were found.

- Results of improvements were heavily highlighted driving pressure to focus on bigger opportunities instead of all opportunities, constantly looking for the “big win”.
- Benefit reporting required significant effort (80+ Hours/Month)

### 11.1. Tactical - Focus Areas

To maintain the strategic big picture improvements, increase the velocity of smaller improvements, increase waste identification, and engagement LEO focus areas were developed. The concept for our focus areas came from the learnings of the workstream approach implemented within the Lean Network team. The three key focus areas created were:

- Strategic Improvements
- Standards
- Training & Engagement

While our focus areas aligned to our overall vision, additional granularity was required with a secondary vision for each area shown below in table 2.

**Table 2 – Focus Area Vision Statements**

Focus Area	High Level Description	Vision Statement
Strategic Improvements	Doing things right	Leave behind the tools and culture for teams to live into the “Culture of Why”
Standards	Doing the right things	Engaging the senior leadership team (SLT) with the right story of the “means” instead of the “results”
Training & Engagement	Building Lean culture	“Being Lean” instead of “Doing Lean” The Kata mindset is part of our Deoxyribonucleic Acid (DNA)

Each focus area concentrates on delivering specific services which laterally align to services provided within the other focus areas. Implementing this highly focused approach encourages ownership and accountability from a LEO team member to the services required by our customer. An additional benefit was the ability to make simple lateral moves with efficient hand-offs of complimentary services between the focus areas.

We refer to this as “Adaptive Standardization”. Internally, the LEO team has clearly defined standards and processes with expected outcomes for each service, aligned to the overall vision, and the vision of each focus area the customer may be part of.

However, each department (or team within a department) brings unique improvement opportunities and challenges to LEO. This could be attributed to our Training & Engagement focus area services, or simply a desire to access training for their team. The customer roadmap, or ask, dictates the LEO team response, which is the heart of our cultural pull methodology. No matter how a team or department engages LEO, there is a clear internal process to help them achieve their goal and clear roadmap to move forward. We manage the means to help our customers achieve outcomes.

This approach also mitigates most resource challenges within LEO. If an improvement is categorized as large (crosses multiple departments) it can be divided across our training channels of: Yellow Belt Accelerator (YBA) and in-house Green Belt training with strategically targeted areas of improvement. This allows our Black Belts to provide support and mentorship to the training channels, ensure the overall strategic vision is on track, address the improvements which require more advanced Lean skills, and correctly prioritize the improvements. Also, leaving behind the tools and culture across the departments to continue improvements.

Figure 10 below provides a high-level overview of the focus areas. LEO team members are assigned to spearhead a program or process within the focus areas. A weekly agenda driven touch base with the team provides visibility into each focus area between the team members.



**Figure 10 – Focus Areas – High Level**

## 11.2. Tactical - Services Breakdown

Service provided by LEO aligned to the focus areas detailed above in figure 10 allow us to deliver multiple improvements concurrently while leaving behind the culture and tools for teams to continue improvements on their own. Each service within a focus area listed below in table 3 aligns across all focus areas to increase efficiency of hand-offs between LEO team members.

**Table 3 – Services Breakdown**

<b>Focus Area</b>	<b>Services</b>
Strategic Improvements	Strategic Yellow Belt Accelerators Strategic Green Belt Projects Black Belt Projects Strategic Improvement Vision Alignment Mentorship
Standards	Lean Tools & Templates Enterprise Project Tracking & Training Metrics Resource Capacity Management LEO Adaptive Standardization Process Control Project Governance Benefit Realization Reporting
Training & Engagement	Online White, Yellow, and Green Belt Training In-House Green Belt Training Workshops Yellow Belt - Self Led Improvements Yellow Belt Accelerators Enterprise Communication & Brand Management Training Pathway Maintenance

### 11.3. Tactical – Yellow Belt Accelerator

To increase the speed of smaller improvements. The development of a new program was required, this is our Yellow Belt Accelerator Program. This flexible program and approach allows us to engage teams in identifying improvement opportunities, provides training and skills to teams ranging in size from five to forty people, while introducing PDCA in a Mentor/Mentee approach. YBA improvements are targeted towards specific teams controlling the sphere of influence to no more than one outside team. This provides a secondary benefit of increased engagement within the team as results are delivered in one to three weeks.

The YBA program also allows for implementation in Strategic Improvements with teams working through an identified opportunity that benefits the Strategic Improvement.

Figure 11 below showcases our YBA program process. Each YBA takes approximately two to four weeks to complete the training cycle with the participants. Identified opportunities from the workshop through submitted problem statements, the stupid rule game, or through the supplier, input, process, output, customer (SIPOC) activity are either actioned immediately through a just do it (JDI) activity or prioritized with the group leader (sponsor) for action in the coming weeks. Upon completion of the initially identified opportunities LEO continues to support and guide future improvement opportunities with the team.



# Lean Accelerator

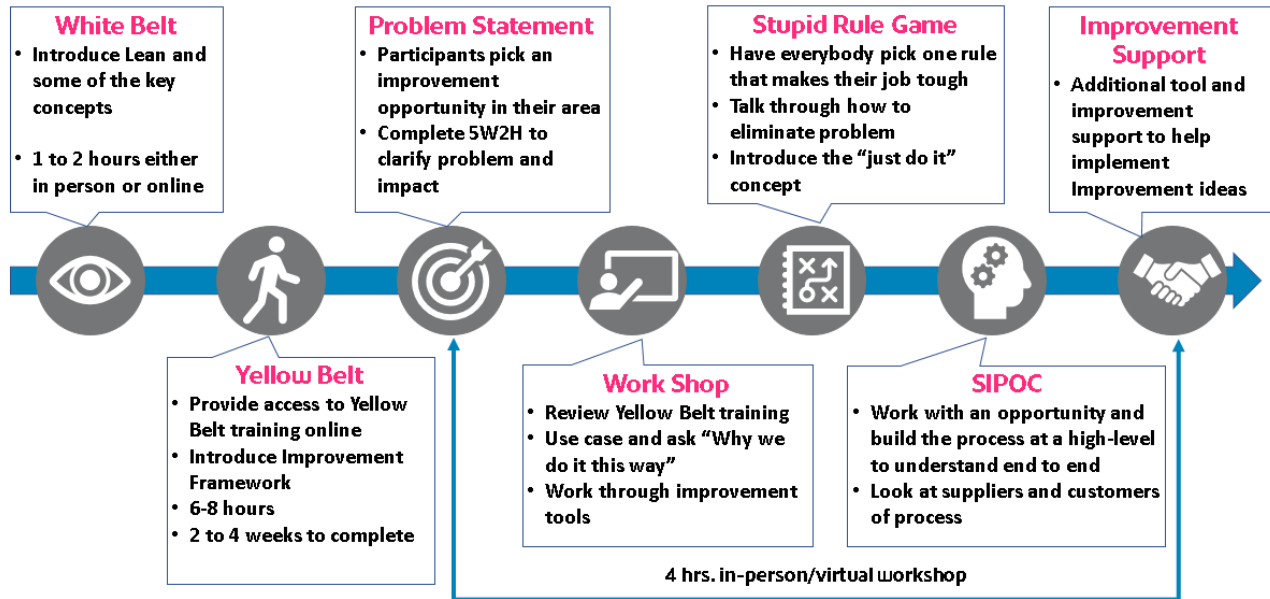


Figure 11 – Yellow Belt Accelerator Program

## 11.4. Tactical – In-house Green Belt Training

The previous approach to Green Belt training with outside vendors has been replaced with vendor-supplied online training augmented by in-house Green Belt training workshops. This change mitigated earlier challenges and pressures to fill seats in the class, and eliminated the long lead times associated with finding a project. All participants in the Green Belt training only begin training when an improvement project has been identified to benefit a Strategic Improvement. This change has also significantly decreased the low secondary improvements conducted from trained Green Belts due to the increased number of identified opportunities discovered via the Strategic Improvement or through YBA programs.

## 11.5. Tactical – Kata Vision and Target Condition Identification

To mitigate Black Belt deployment in a business analyst role, we implemented the requirement for strategic improvements to define a vision prior to discovery of opportunities. This approach allowed for the overlay of the Kata approach using target conditions moving towards a vision. At this time, the methodology to reach each target condition is left to the discretion of the Black Belt either using Define, Measure, Analyse, Improve, and Control (DMAIC) or PDCA. However, we encourage using PDCA and a Mentor/Mentee approach as much as possible. Secondary benefits from the vision first approach is the reduction in lost sponsorship and unrealized efforts as the sponsor is engaged in developing the vision. Should priorities shift elsewhere, time invested is reduced to two-to-four hours instead of months on discovery activities.

## 11.6. Tactical – Benefits Realization

Our biggest challenge was influencing the culture of MBR towards MBM. Previously highlighted results had focused on either time savings or dollar savings, both of which we considered appropriate for roll-up reporting. However, the use of multiple loaded head count rates, value creation forms, limited finance

oversight, calculating efficiency into dollars, multiple reporting teams, and saving categories with vague descriptions influenced our initial decision to not add an additional report using LEO defined categories and calculations from activities in our scope.

Using our own methodology of asking “Is there a better way to do this?” we developed a multi-level plan to standardize benefits realization across the enterprise and begin the cultural shift towards MBM for Lean reporting. Following our vision to “engage the senior leadership team with the right story of the ‘means’ instead of the ‘results’”.

Our first discovery was that various teams had started to implement benefit realization approaches as they faced the same challenges identified above. Following a simplified pareto analysis, we aligned to the team with the greatest amount of support within the delivery and execution teams.

Our next step was to identify the data points required to implement the aligned benefits realization approach, and to identify additional data points that would help move us towards MBM. Our primary focus while identifying MBM based data was to incorporate actionable metrics and leave vanity metrics out of future reporting. Our criteria for actionable metrics involved whether a split test could be conducted on the data point.

With an aligned team-based approach and actionable metrics gathering in place, we incorporated a benefit language statement and syntax approach matched to the categories of: cost avoidance/risk mitigation, cost savings, efficiency, and revenue instead of category descriptions. The benefits of this methodology include less room left for interpretation compared to a category description, simplicity to contributors less familiar with expressing savings, requiring less rework and follow-up from LEO, and reducing time spent collating reporting with automation techniques allowed due to common syntax.

Category	Plain Language Statement	Benefit Syntax	Example Benefit Statement
Efficiency	We will reduce the number of truck rolls	E: Reduce <del>###</del> truck rolls per <u>month/year</u>	E: Reduce 500 truck rolls per year

**Figure 12 – Benefit Syntax Example**

Our final step entailed a transitional period of introducing MBM based data while maintaining previous time and dollar savings on improvements already underway, followed by the incorporation of the new benefits realization approach on improvement initiatives kicked off after the cut-off period. This revised approach has reduced benefit reporting from 80+ Hours/Month to an average of 5.25 Hours/Month.

## 12. Results

Improvement opportunity results since the implementation of LEO shown in figure 13 below have continued to remain steady with the benefit of learnings from the Lean Network team. However, a significant increase in adoption of Lean across the enterprise has been realized with 6 new departments in F19, 30,000 hours saved, and 40% of improvement opportunities solutioned immediately through JDI's. This trend based on current statistics indicates we have achieved our goal of increasing velocity of improvements. From a Training and Engagement perspective, we have experienced a 96% increase in belt acquisition through online and F2F training from individual contributors compared to F18 and the 65 participants in Green Belt certification is a 15% increase from F18.

### Lean Enterprise Office | Outcomes - F19



#### Strategic Improvements

- ✓ Average of 40 Yellow Belt Accelerator Opportunities discovered each month
  - 40% solutioned immediately
- ✓ 8 Departments (Spokes) with active Continuous Improvement teams and initiatives
  - Increase of 6 Departments in F19
- ✓ 30,000 Hours saved (F19 to Date)




#### Training & Engagement

- ✓ 20.6% of Enterprise with White or Yellow Belt
  - 96% increase from F18
- ✓ 65 Participants in Green Belt Certification
  - 34 Continuous Improvement initiatives



Figure 13 – LEO Results – F19

Figure 14 below highlights the training and engagement statistics. As of June 2019 20.6% of the enterprise has received training in white belt via face to face (F2F) training, or through our online training platforms. This totals approximately 2000 individual contributors. Upon completion of yellow belt training, each graduate is required to complete a small improvement within their area. These improvements range from 5S of their workstation or desktop, up to organizing their trucks. This small improvement helps cement the learnings and provides a small efficiency increase to the enterprise.



**Figure 14 – Training and Engagement – F19 Totals**

## Conclusion

Through learnings and benchmarking over our seven-year journey in Lean we have found that sustainment of a Lean program requires a focus on culture that is integrated with all aspects of the strategic approach and tactical processes put in place. By analyzing the challenges and learnings within an organization and defining a clear vision and goals, a tactical plan can be formulated to reinforce the culture of continuous improvement through behavior enhancing processes and programs.

The tactical processes in place at Shaw like the Yellow Belt Accelerator program, vision statements and target conditions, and adaptive standardization ladder up to the three-pillar approach principles by way of clear actions, defined outcomes, and processes that reinforce the culture of continuous improvement. The methodologies for changing cultural approach outlined here are built on Shaw principles and existing culture, adaption to your unique environment may be required. It is important to implement an approach as described in this paper based on the learnings, discoveries and challenges within your organization.

## Abbreviations

5S	“Sort”, “Set In order”, “Shine”, “Standardize” and “Sustain”
CI	continuous improvement
CPE	customer premise equipment
CPO	chief procurement officer
DAA	distributed access architecture
DMAIC	define, measure, analyse, improve, control
DNA	deoxyribonucleic acid
F2F	face to face / face 2 face
JDI	just do it
LEO	lean enterprise office
MACD	move, add, change, delete
MBM	managing by means
MBR	managing by results
NDC	national distribution center
PDCA	plan, do, check, act
RDC	Regional Distribution Centers
RF	radio frequency
SIPOC	supplier, input, process, output, customer
SLT	senior leadership team
SME	subject matter expert
TPLT	total process lead time
TPS	Toyota production system
YBA	yellow belt accelerator
VOP/E	voice of the process/employee

## Bibliography & References

- [1] J Womack, D Jones, “The Machine That Changed the World” Simon & Schuster, October 10, 1990
- [2] M Bremer, “How to Do a Gemba Walk: A Leader’s Guide” CreateSpace Independent Publishing Platform, January 30, 2016
- [3] J Shook, “How to Change a Culture: Lessons from NUMMI” MITSloan Management Review, Vol 51 No.2, Winter 2010
- [4] M Rother, “Toyota Kata: Managing People for Improvement, Adaptiveness and Superior Results” McGraw-Hill, August 4, 2009

# **New Generation Data Governance for Charter Network:1**

A Technical Paper prepared for SCTE•ISBE by

**Jay Liew**

Advanced Analytics Architect  
Charter Communications  
14810 Grasslands Dr. Englewood, CO 80112  
(720)518-2277  
Jay.Liew@charter.com

**Mark Teflian**, Charter Communications

**Bruce Bacon**, Charter Communications

**Jay Brophy**, Charter Communications

**Randy Pettus**, Charter Communications

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
The Deluge of Streaming Data in Telecommunications .....	3
Data Governance for Streaming Events – Lessons from Total Quality Management.....	4
Overview of Data Governance .....	4
Data Governance Shortcomings .....	6
N:1 at Scale Data Governance for Services and Applications.....	7
N:1 Data Product Catalog .....	10
Network:1 Use Case – Optical Network .....	10
Conclusion .....	13
Abbreviations.....	14
Bibliography & References .....	15

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - The DGI Data Governance Framework .....	6
Figure 2 - High Level Architecture of the Data Distribution for N:1 Data Governance.....	12

# Introduction

Charter Communications' Network:1 (N:1) strategic network architecture creates a transformed, unified network for product and application services. It enables strategic increases in capacity and performance, creates adaptable, scalable, reliable, and secure network design patterns, and enables network modeling, abstraction, and orchestration. While facilitating the commoditization of network and systems functions, it leverages programmability and the disaggregation of traditional telecommunications offerings.

Foundational to N:1 is its advanced network analytics services that rest upon a modern data plane infrastructure producing vast amounts of data that drive intelligence and decisions for both humans and machines. Numerous and disparate devices, equipment, software technologies, and applications create and drive data in disparate formats for near-real time analytic model execution and decisions. In addition, various consumers have unique needs to make use of these data. Applications involve joining customer experience, network quality of service, traffic engineering, and several other data sources, which create extraordinary challenges and opportunities for unified network intelligence. These demands and complexities beg the question whether it is possible to govern data in this environment. This question only heightens with skepticism that surrounds data governance today and, often, its inability to achieve the benefits organizations expect from it.

We examine the need for new at scale data governance for near real-time streaming and event data for human and machine actionable analytics within N:1. We also provide an overview and common framework for current data governance while addressing its shortcomings. While data governance encompasses a broad array of processes and governing bodies, we focus on the various technical aspects critical for success within N:1.

We assert that current data governance methods must evolve to enable the complexity of near real-time data streams from multiple sources. By relating to Total Quality Management (TQM), we define the new technical data governance components that are necessary to maintain data integrity, control, and value for intended consumers as data move in this environment. Finally, we show how a curated and collaborative Data Product Catalog helps address today's governance challenges, enabling responsible data production, consumption, and joins using big and small data.

## The Deluge of Streaming Data in Telecommunications

Charter Communications' Network:1 (N:1) architecture enables a transformed, unified, software centric, single image IP services network for strategic increases in capacity and performance. It will leverage network analytics, modeling, abstraction, and orchestration, and seize upon the commoditization of network and systems functions while disaggregating traditional telecommunications offerings.

A critical aspect of the N:1 architecture is a modern data plane separated from the control plane that accommodates numerous virtual devices, software technologies, and applications. The data produced from these sources will continue to see explosive growth in the near future. Global Internet of Things (IoT) IP traffic is predicted to grow more than sevenfold by 2021 while global IP video traffic is expected to increase threefold from 2016 to 2021, at which it will account for 82 percent of all IP traffic[1]. Within Charter's vast network, these data will easily amass Exabytes of scale by this timeframe.



Upon this data plane is the advanced analytics that service network architects, engineers, and various other consumers across business units that need data to meet business objectives. These needs include obtaining access to raw, unaltered streaming or event data that are produced from its origination point to Analytics Data Sets (ADSs), which are data products with enhanced intelligence. Finally, consumers need model outputs and algorithms, which can provide descriptive, diagnostic, predictive, or prescriptive outputs that drive decisions and machine intelligence.

The network services and product data produced in this environment must be accessible and usable in near real-time, and thus, governed in near real-time. For instance, models for subscriber and policy management, bandwidth optimization, network service theft, security mitigation may all require near real-time decisions. In addition, optimizing network Quality of Service (QoS) will require near real-time data processing and model scoring for effective allocation of network resources. This is in contrast to a traditional data governance setting for data at rest, where data persist in a data warehouse as a reduced data set of certain facts and dimensions and is then utilized for reporting at alter time, perhaps even months later[2].

While vast opportunities exist to monetize and make use of these data, challenges for governing data proliferate in this environment. The volume and velocity of disparate data can overwhelm consumers of the data. Near real-time applications will fail without accessible, reliable, and accurate data. In addition, numerous non-traditional sources will produce disparate semi-structured and unstructured data with diverging use cases.

## **Data Governance for Streaming Events – Lessons from Total Quality Management**

Before addressing data governance and its relation to N:1, we first examine how data governance aligns more broadly with quality management practices. Quality management has been traditionally capitalized upon in manufacturing settings, where it has evolved to the various forms and derivatives of Total Quality Management made popular in the 1980s. A comprehensive approach to quality exists under TQM with a focus on managing the organization for what is important to the customer [3]. Two primary goals exist within TQM [3]:

1. Design of the product
2. Ensuring the organization's processes can consistently produce this design

We can relate managing the quality of the nearly limitless amount of data in the N:1 architecture to these TQM goals. The customers or consumers in this environment are the machines and humans that use data to drive the intelligence and decisions for the network. Data are the valuable products these customers desire, and the organization must carefully design these data for their use. Meanwhile, the organization must ensure processes can consistently and reliably use and move the data at any point as it flows through the chain. This includes from its source to various downstream applications, including additional data products and model outputs for various consumers.

## **Overview of Data Governance**

Data governance has been traditionally defined as a framework for decisions and accountabilities within corporate structures to manage data and enable desirable business outcomes from its use [4]. While data governance is customarily associated with IT governance, organizations are pushing towards more data

governance initiatives from a variety of angles, including regulatory requirements, business needs, social considerations, and privacy concerns.

One side of the data governance spectrum consists of organizations that face high regulations, such as banks. These organizations must push for solid data governance frameworks to meet regulatory requirements. On the other side of the spectrum, there are web-scale organizations that have built invaluable data assets. Data governance in this sense is often driven by business needs centered on unlocking the value of these assets. We would argue that Charter aligns more closely with these web-scale organizations in the context of N:1 due to its scale, complex services, and unique value of the data.

While various data governance frameworks exist, at the core of modern data governance programs is the view that data are strategic assets and should support a specified mission [5]. The organization, policies and standards, governance metrics, processes, technology and data architecture all play a role in structuring a data governance initiative around these data assets [5].

Data quality includes defining standards for data and having a means of assessing data quality and measuring its performance. Defining standards includes detailing end-to-end data quality, including definitions, controls, and adoption. This also includes maintaining business and technical definitions for consistency and integrity. Meanwhile, assessing quality and measuring performance includes proactive and reactive assessment, and cleansing and remediation of data quality issues that are aligned with business processes [5].

A simple common data governance framework from the Data Governance Institute (DGI) is shown in Figure 1[6]. This framework shows how the Mission (Why) is the first step taken, while People Organizational Bodies (Who), Rules and Rules of Engagement (What) and Processes (When & How) are enacted to support this mission [6].

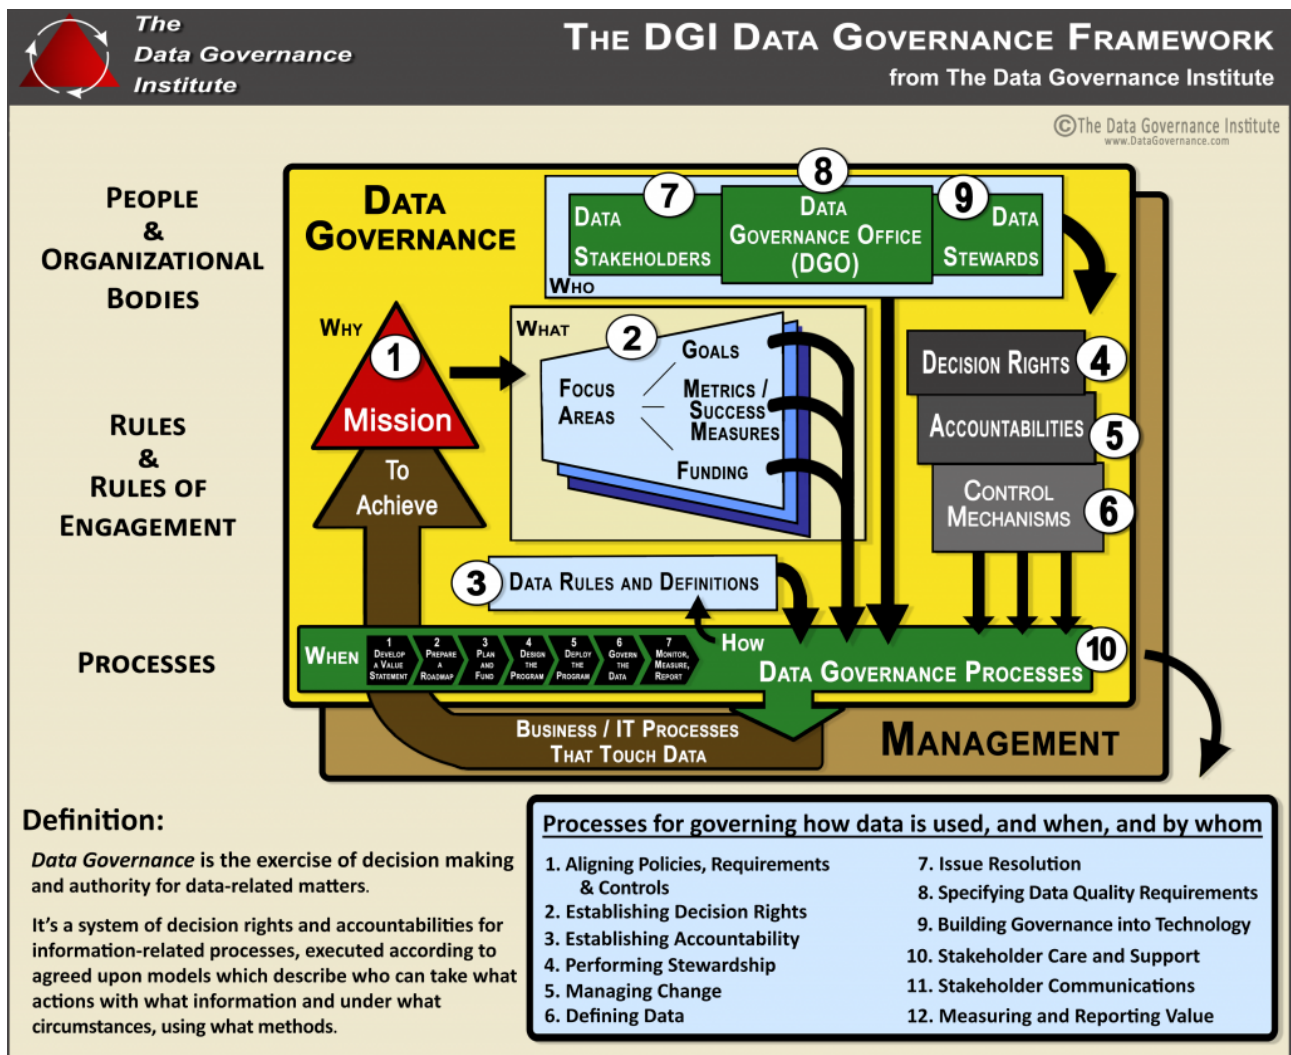


Figure 1 - The DGI Data Governance Framework

## Data Governance Shortcomings

Before applying data governance concepts to near real-time streaming and event data, it is important to understand where data initiatives fall short, especially in how these shortcomings relate to data governance. In examining various surveys on the topic, a lack of data clarity, ownership, and accountability drive many of these failures.

- **Data Quality & Clarity**

- For many organizations, the lack of clear data definitions plagues data initiatives. Only 10 percent of organizations in one study have full documentation for their front-to-back data flows across the organization [7]. Thirty-six percent of companies do not have a data dictionary, and those that do admit the dictionaries are mostly held internally within divisions [7]. Data quality remains a major concern as well. Based on a 2017 Harvard Business Review study, only 3 percent of executives found that their departments had a minimum acceptable rate of data accuracy in their data [8]. The top reasons for poor data

quality include a lack of uniform definition of data fields (no “Golden Source”) [7]. Similarly, business intelligence initiatives are often less than optimal, with 61 percent of respondents noting inconsistent data sources and 57 percent noting incomplete data as reasons [9]. Fifty-three percent say that their BI processes deliver inconsistent or unreliable conclusions [10].

- **Ownership & Accountability**
  - While 71 percent of organizations have some ownership around data, most admit it is not clearly defined and responsibilities are not clearly documented [7]. Silos only increase the challenge, as 56 percent noted data silos as one of the reasons data initiatives fail to achieve their value [7]. A lack of ownership and accountability lead data to be inaccurate, uncontrolled, or dormant. Without clear ownership, many initiatives lack a focus on the potential ethical, social, and regulatory implications [11]. Finally, data governance can often become a gate that just controls access to datasets, often creating hurdles for business users in getting access to data when needed, even for the company’s highest priorities.
- **Hurdles for Data in Motion**
  - Simply applying traditional data governance approaches to data in motion poses significant challenges for data efficacy. Under traditional systems for data at rest, data is controlled and monitored for access, use policy, and reliability in this central data store. However, in a complex streaming environment with near real-time uses, this system will impede results. In addition, it will fail to realize that data are assets even when they are not persisted in a data warehouse or a data lake. We will further emphasize this challenge and propose solutions in the coming sections.

## **N:1 at Scale Data Governance for Services and Applications**

Data produced in N:1 must be governed to meet today and tomorrow’s near real-time applications. Overall, traditional data governance components of data standards and quality testing, performance, and measurement still apply. However, these aspects must be evolved and modernized for scale. We provide an overview of these components below, and provide additional detail later in how these can be applied in a real protocol at scale use case.

We focus our discussion on the most critical technical processes to ensure quality with near real-time streaming and event data and the effects of decision latency. While we do not focus on the softer components of a data governance, such as governing bodies and structures in this discussion, we do note that they are extremely crucial to maintaining governance. A function within an enterprise must exist to provide guidelines, standards, and make governing policies, such as those for access, security, and retention.

For this discussion, we walk through the data value chain as it proceeds from production through distribution for various applications. This data can be consumed for near real-time applications, and transformed or joined to other data all while being in motion. Data in this chain can also be stored or used to produce automated Analytics Data Sets (ADSs), where intelligence is instrumented. Finally, analytics and data science tools can consume this data at various points through this chain. The point of emphasis is that the data is an asset throughout this value chain and not just once it is at rest. Accordingly, these processes ensure quality throughout the chain so that intended consumers receive their intended data products at whatever point in this chain.

- **Data Production Controls**
  - Aligning with our discussion of TQM, governing data at the source involves ensuring a design is created to produce the data. The traditional trace-to-source practices are obsolete in themselves. Initial data engineering must include detailed specifications on how the data design can be consistently produced. To have a clear definition of produced data at the source requires subject matter experts (SMEs) to produce and properly test systems so that data are fully understood and consistently produced according to design. This stage should include any necessary components to address quality at the source so that potential consumers know expected results as well as limitations.
- **The N:1 Data Distribution Bus**
  - The data distribution stage is a crucial point in enacting data governance with near real-time streaming and event data for key reasons. First, placing data governance controls during distribution enforces the notion that data is an asset while in motion and not just once it is at rest. Second, near real-time applications must rely on the speed of obtaining and utilizing data. For instance, a near real-time fraud detection application cannot wait for data to become at rest or it will fail to deliver its ability to effectively prevent a threat. Finally, this stage represents a branch in the data value chain where it can move to various consumers. Similarly, to how a manufacturing company under TQM addresses possible quality concerns at the point of high-value activities, this stage allows an enterprise to address data quality before errors magnify as data moves to various other phases, applications, or storage. The criticality of this phase cannot be emphasized enough, and is a shift in thinking from the traditional view of governing data further downstream.
  - To govern data and the various touchpoints in this stage, the Data Distribution Bus brokers the data by topics among various producers and consumers in a streaming environment [12]. Producers and consumers of data on this data distribution bus, and similarly those able to read and write data, should be registered, known, and controlled. Producers should be on-boarded with controlled requirements to produce topics while consumers should also be on-boarded to understand their role in maintaining traceability, meeting requirements, distributing messages, and ensuring company policies [12]. Meanwhile, topics should be properly marked for their availability, integrity, security, and sensitivity, especially regarding personal identifiable information (PII) [12]. Finally, this system should log the Who, What, and When components of producing and accessing data for full audibility [12].
- **Schema Registry**
  - Data originating from multiple sources can lead to multiple formats containing ill-defined and malformed data. As a result, the quality of analysis directly correlates with the quality of data. During data distribution, schemas, the “grammar” of the data, ensure that structured data meets desired designs of quality [13].
  - Schemas not only define the structure of data payload, but more importantly, they define the contract between producers and consumers of such data. By defining schemas, data integration between producer and consumers is simplified.
  - Data assets evolve with respect to time along with the schema. Therefore, versioning of schema is necessary for data to evolve and yet be accurate. The evolution of data directly affects consumers. Without proper understanding of the evolution, integration time between consumers and data sources increases. In addition, schemas prevent bad data from infiltrating into topics by ensuring data types, formats, and design expectations are met. As data evolves, backwards compatibility ensures that new schemas can read old schemas. No breakage of functionality should occur with multiple consumers on multiple systems consuming data.

- **Data Lineage**
  - As data moves downstream from its source, it can evolve and transform to provide additional value to consumers. For instance, during data engineering, data assets can be joined with other data assets for additional intelligence, including through adding new and enriched data features. To ensure veracity as this data is curated and transformed, the lineage of these data must be clear from source to where it is at any point along the value chain. This includes knowing the various touchpoints of the data, including who changed it and what changed while tracing the data from its source.
  - Data lineage tracking ensures quality for various upstream and downstream uses. For instance, this lineage provides an understanding to work backwards or downwards for privacy, ethical, or regulatory aspects. It can also be important for consumers, which could include analysts, data scientists, business users, machines, and/or applications, to know the downstream derivatives of data. For instance, a consumer might benefit from knowing that instead of ingesting raw data for a business problem, an ADS, with its enriched contents, may better align with their application.
- **Data Dictionary**
  - Similar to understanding the data structure, it is important for the various consumers to understand the contents of the data. A curated data dictionary provides appropriate data asset information to data scientists, analysts, and architects so they can understand the technical aspects of the data and its underlying meaning. Meanwhile, a data dictionary allows business users to understand the business definitions, enabling these users to align the data contents with strategic initiatives. While the cost to produce and maintain this level of information might seem high, the benefits exceed this cost, since it reduces the time to understand and act on the data. This is especially true for large, complex enterprises where vast resources are often spent understanding and relating data.
- **Quality Testing, Performance & Measurement**
  - Measuring quality and performance of data in this environment is important to assure accountability from producers while ensuring reliability to consumers. For instance, if a near real-time model is scoring or inferencing streaming data, consumers need guarantees regarding the uptime of the stream and overall reliability of the data flow. Thus, critical measurements must be performed, including measuring the accuracy, completeness, validity, timeliness, and reliability at core touchpoints in the data stream. This includes at production, during data distribution, and through any additional value-add tasks during the data value chain. Additional data quality checks should also be considered, such as handling of Null values to detecting anomalies, depending on downstream business applications. Various algorithmic approaches can be utilized to monitor and control data quality in this phase [14].
- **Decision Outcome Collaboration**
  - It is important to keep in mind that data in this environment exists for the purpose of enabling decisions for humans and machines, and thus a governance program should keep this front and center to facilitate such decisions. One way to enable this is through curating decision outcomes and aligning these with the various data products, including both the data streams and ADSs. This method allows both strategic and near-term decision science for network engineers and architects to have a common framework to visualize outcomes coupled with the necessary data. Users in this environment can validate and improve content, and raise alerts or concerns. They can also view descriptions of use cases of a particular data set along with pre-built artifacts containing code or other solutions for decision outcomes that can enable users to get a head start on certain problems.

# N:1 Data Product Catalog

These data governance components are realized through a curated next generation Data Product Catalog (DPC). It is estimated that by 2020, organizations that provide a curated data catalog will realize twice the business value from analytics than organizations that do not [15]. Much like Yelp™ [16] provides a curated catalog to help people solve the problem of finding a business in their area, a data product catalog can be modernized to ensure governance for data in motion on Charter's vast core and access network, thus aligning data products with network services.

The N:1 Data Product Catalog (N:1DPC) includes information on producers and consumers, core information on topics, including recovery and topic monitoring requirements, encryption, corporate data classification and other security and retention requirements. Schema information is populated including version history, effective dates, formats, message sizes, error rates, persistence, replication, and priority. In addition, users can preview the schema and have the ability to download for additional analysis or investigation.

A data dictionary provides further visibility by providing the metadata for each ADS. Environmental aspects are included, such as the ADS name, business description, source, creation and updated dates with any additional lineage information to trace to the source. Other information, such as the owner, size, protection modes, storage formats, and refresh cadence allow prospective or current users with the ability to understand high-level aspects of the data. Meanwhile, metadata is provided to convey the dataset schema, core descriptions of each field within the ADS, datatypes, and calculations for derived fields. Additional statistics for contents are provided to show items, such as key descriptive statistics and percentage of NULLs.

Finally, at the core of N1:DPC is a new collaborative environment that links decision outcomes with the various datasets. Decision outcomes are populated with the ability to leverage crowdsourcing techniques for users to search and tag content with additional feedback mechanisms to validate and improve content. The N1:DPC also incorporates Machine Learning feature engineering to accelerate model developments, as well as unfinished blocks of reusable human and machine algorithms to accelerate problem solving. Role based access policies connect SMEs and consumers.

## Network:1 Use Case – Optical Network

For years, Multiple System Operators (MSOs) have been talking about the advantages of advanced analytics and automation. These efforts have been focused on service creation. Advanced Engineering Core Optical has challenged our vendor community to broaden this definition to include the reduction of manual effort and the need to provide “raw” data for analytic models in all areas of Charter's business.

From the Optical Network perspective, various challenges arose without a common data and governance architecture.

- Each vendor product only addresses data out of its own Management Systems.
- Any analytics solutions and outcomes are based on vendor models only; vendor solutions are a subset of problems and solutions use cases engineering needs to solve.
- There is no way to visualize the entire Optical Network holistically when operators have to manage multiple heterogeneous EMS environments.
- It is difficult to get raw data from vendor systems to fully contextualize the data.

- The need for data from all vendors to fully understand the state of the network, or the future state of the network.

## **Background**

The optical network forms layer 0 of the network stack and spans across Charter’s Backbone, Regional and Metro markets. Telemetry data from these optical network devices, have traditionally required very low level polling, such as TL1, and data acquired is vendor specific and cumbersome to make use of. MSO’s are in need of all telemetry data from these devices, to properly trend key characteristics of the network.

In most cases, every vendor has its own EMS (Equipment Management System) that provides vendor specific visualizations as well as analytics. MSOs are therefore beholden to the vendor provided analytics. Charter is not only interested in canned vendor analytics, but has aspirations of building in-house analytics models, which require data that is traceable to the source. Charter has been diligently working with vendors for a modern data solution that makes use of API calls to the EMS, and event streaming techniques.

Acquiring telemetry data from all vendors, enables data discovery that drives high ROI analytics use cases, such as traffic engineering and capacity planning. By addressing data governance at the onset of the data pipeline architecture, it has enabled Charter in the following areas.

- Vendors are starting to expose telemetry data through REST APIs, streaming, GRPC etc.
- Data quality of raw data feeds is addressed from the onset.
- Continuous feedback, evaluation, and iteration of data with vendors.
- Data model of streaming data is important, to address quality of data. The use of the Avro™ file format ensures that streaming data adheres to the data model.
- By addressing data quality and governance at the start of the process, enables the visualization of streaming telemetry data in a very short timeframe.

## **Data Distribution**

During data distribution, data governance is enacted across optical network telemetry data that is in motion, ingesting through Apache Kafka™. All data used by consumers of the optical network telemetry data is of the Avro™ file format. In parallel with and underlying Apache Kafka™ distribution, Confluent® Schema registry ensures schema governance while also improving the developer and consumer experience [18]. The Schema Registry stores a versioned history of all schemas and allows for the evolution of schemas according to the configured compatibility settings and expanded Apache Avro™ support [19].

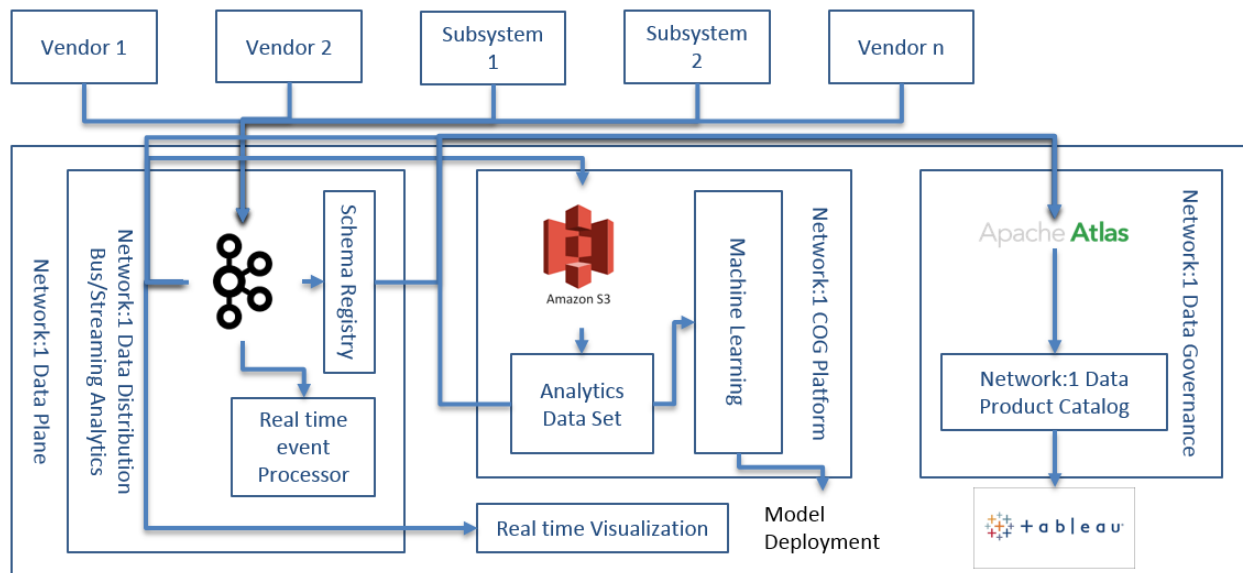
Avro™ is in binary file format and is more efficient and compact compared to JSON [20]. The schemas for any topic can evolve with respect to time, which is essential as products also evolve [21].

A question is often presented for why Avro™ encoding is favorable to JSON. In addition to the file format benefits, Avro™ has strong governance benefits over JSON. In a world of big data, “data lakes” often become “data swamps.” These large amounts of data, if not governed properly, can quickly become overrun by data errors and become hard to use, or in some cases, unusable. Worse yet, these erroneous



data can be used to produce decision outcomes that have a deleterious effect to N:1. This Avro™ constraint prevents data stores from becoming the next swamp for ungoverned data.

Figure 2 depicts the high-level architecture of the data distribution for N:1 Data Governance.



**Figure 2 - High Level Architecture of the Data Distribution for N:1 Data Governance**

### **Real-time event processing**

Data distributed from vendors in the optical network use case is raw in nature consisting of telemetry and inventory data. In many cases, the telemetry data ingested needs to be enriched with inventory data or data from other sources to make telemetry data more useful. Since vendors visualize data on their vendor specific EMS, there was no mechanism to view data from an optical network perspective as a whole, vendor agnostic. Vendor data is diverse and represented differently from system to system.

Real-time event processing enables the enrichment of all the telemetry streams to create streams that are vendor or non-vendor agnostic, and enable MSOs to view the optical underlay from a holistic network perspective.

Real-time event processing for the N:1 Data Distribution Bus is implemented on the KSQL™ [23] from the Confluent® Platform. KSQL™ enables the joining of streams of data that are in motion. The powerful concept enables data exploration and discovery, and real-time monitoring and analytics.

### **Data Lineage**

Real-time event processing and data engineering are complex engineering tasks performed to derived enriched data streams or ADSs. As data evolve through this life cycle, it is important to understand how every byte of data moving though a data platform is derived, and all dependencies are documented to the fullest. Metadata from data in motion and at rest can be used to build data ontologies, which are invaluable information for any data source for a data consumer.

Data Governance and Metadata framework to govern and classify data used is Apache Atlas™ [24]. Apache Atlas™ provides the ability to ingest metadata from data sources, classify, and build lineage information for every dataset.

### **N:1 Data Product Catalog**

The N:1 Data Product Catalog (N:1DPC) contains a collaborative environment for data assets, including data feeds and enriched streams for the Optical Network use case. The data feeds information includes the Avro™ schema information, including producer, consumer, and topic data. Prospective users, such as developers, can view schema information, download samples, and view information about the feed and how it might relate to their use cases. Quality performance measures provide accountability for the producers and show reliability for consumers of the data.

The N:1DPC contains Optical Network data information, including metadata, as well as information for all the fields in the dataset. This includes a description of the field, data types, and any SQL or other language used for derived calculations or engineered features. Lineage information is also available, as trace to source is important to determine data quality and viability.

## **Conclusion**

Charter Communications' Network:1 will enable and produce vast amounts of data to drive human and machine intelligence. As part of this environment, applications will rely on having near real-time uses of streaming and event data that will require a new data governance approach as compared to traditional methods.

We align the near real-time streaming and event data environment within N:1 with lessons from Total Quality Management practices that emphasize the design and consistent production of products, which also must meet their desired customer expectations. Within N:1, these data are the valuable products that are designed and intended for the various humans and machines applications. Due to the nature and complexity of this system, data governance controls must be placed throughout the data value chain and not just downstream once the data is at rest. This shift from traditional data governance practices involves data governance controls being placed at the point of production, upstream in distribution, and at various key downstream touchpoints.

Emphasizing these practices with the Optical Networking example, we have shown various technical data governance components during these phases. This includes ensuring data production has the ability to consistently produce as it is designed. We also emphasize the importance of data governance during data distribution, especially with near real-time consumption applications and various dependent paths the data can flow after this phase. The N:1 Data Distribution Bus governs data topics and the various touchpoints during this phase for the producers and consumers. Meanwhile, schemas that always “travel” with the data are critical to ensure consumers receive data as expected from topic producers. Finally, a collaborative N:1 Data Product Catalog centralizes data assets, including data source feeds, ADSs and other downstream data products or models, conveying the metadata and components for these assets, showing data lineage, and linking data products with decision outcomes.

While some of these governance improvements lead to more upfront work and automation, the collective benefit for N:1 and the entire enterprise far exceed the cost. Just as a consumer should have a clear understanding of a product being purchased, this data governance solution ensures that network engineers and architects, data scientists, developers, and business users can understand and utilize data at various phases according to their business applications.

Overall, the key in the N:1 environment is that data can no longer be thought of as an asset only once it is at rest. It can be consumed and enriched while it is in motion, so it should be treated as an asset once produced, while in motion, once it is at rest, and once it becomes the output of various analytics models. Consequently, data governance must be instrumented and travel with the data to enable effective at scale human and machine decisions for N:1.

## Abbreviations

1NF	First Normal Form
ADS	Analytics Data Set
API	Application Programming Interface
CM MAC	Cable Modem Media Access Control
CMTS	Cable Modem Termination System
CPE	Customer Premise Equipment
DGI	Data Governance Institute
DOCSIS	Data Over Cable Service Interface Specification
DPC	Data Product Catalog (N1:DPC)
EMS	Equipment Management System
IoT	Internet of Things
IP	Internet Protocol
IT	Information Technology
JSON	JavaScript Object Notation
KSQL	Streaming SQL for Apache Kafka
MSO	Multiple System Operators
N:1	Network:1
QoS	Quality of Service
ROI	Return on Investment
SME	Subject Matter Expert
SQL	Structured Query Language
sTQM	Total Quality Management

# Bibliography & References

- [1] Cisco VNI. (2017). “The Zettabyte Era: Trends and Analysis.”  
<https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/vni-hyperconnectivity-wp.html>
- [2] Kimball, Ralph and Ross, Margy. (2011). The Data Warehouse Toolkit: The Complete Guide to Dimensional Modeling. John Wiley & Sons.
- [3] Jacobs, R. F. (2014). Operations and Supply Chain Management. New York: McGraw-Hill Irwin.
- [4] Wende, K. (2007). A Model for Data Governance - Organizing Accountabilities for Data Quality Management. Association for Information Systems.  
<https://aisel.aisnet.org/cgi/viewcontent.cgi?article=1079&context=acis2007>
- [5] Deloitte Consulting Management. The Increasing Importance of Enterprise Data Governance and Management/Case Study.
- [6] Thomas, Gwen. Data Governance Institute Framework. Data Governance Institute.  
[http://www.datagovernance.com/wp-content/uploads/2014/11/dgi\\_framework.pdf](http://www.datagovernance.com/wp-content/uploads/2014/11/dgi_framework.pdf)
- [7] Data Governance Survey Results. Price Waterhouse Coopers. March 2016.  
[https://www.pwc.fr/fr/assets/files/pdf/2016/05/pwc\\_a4\\_data\\_governance\\_results.pdf](https://www.pwc.fr/fr/assets/files/pdf/2016/05/pwc_a4_data_governance_results.pdf)
- [8] Nagle, Tadhg et. A01. (2017) Only 3% of Companies’ Data Meets Basic Quality Standards. Harvard Business Review <https://hbr.org/2017/09/only-3-of-companies-data-meets-basic-quality-standards>
- [9] Forbes Insights. (2016) Breakthrough Business Intelligence. How Stronger Governance Becomes a Force for Enablement.  
[https://images.forbes.com/forbesinsights/qlik\\_bi/BreakthroughBusinessIntelligence.pdf](https://images.forbes.com/forbesinsights/qlik_bi/BreakthroughBusinessIntelligence.pdf)
- [10] Hiskey, Michael. (2017). He Who Rules the Data, Rules The World: A Brief History of Data Governance. CIO Network. <https://www.forbes.com/sites/ciocentral/2017/11/16/he-who-rules-the-data-rules-the-world-a-brief-history-of-data-governance/#689b76fa39b5>
- [11] Fleming, Oliver et al. (2018). Ten Red Flags Signaling Your Analytics Program Will Fail. McKinsey&Company. <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/ten-red-flags-signaling-your-analytics-program-will-fail>
- [12] Gamov, Viktor & Gaur, Vijay. Charter Communications Architecture Review Engagement Report. Confluent. 2018, June.
- [13] Shapira, Gwen. (2015). Yes, Virginia, You Really Do Need a Schema Registry. Confluent.  
<https://www.confluent.io/blog/schema-registry-kafka-stream-processing-yes-virginia-you-really-need-one/>
- [14] Saha Barna & Srivastava, Divesh. Data Quality: The Other Face of Big Data. AT&T Labs-Research. <https://people.cs.umass.edu/~barna/paper/ICDE-Tutorial-DQ.pdf>
- [15] Sallam, Rita. (2017). Magic Quadrant for Business Intelligence and Analytics Platforms.

- [16] Yelp. <http://www.yelp.com>
- [17] Data-Over-Cable Service Interface Specifications. DOCSIS 3.1, CM-SP-PHYv3.1-111-170510. Cable Television Laboratories, Inc. 2017
- [18] Confluent, Inc. [https://www.confluent.io/about/#about\\_confluent](https://www.confluent.io/about/#about_confluent)
- [19] Confluent, Inc. Confluent Schema Registry. <https://docs.confluent.io/current/schema-registry/docs/index.html>
- [20] Apache Avro™. <https://avro.apache.org/docs/current/>
- [21] Confluent Avro™ Kafka Data. <https://www.confluent.io/blog/avro-kafka-data/>
- [22] Snowflake Computing. <https://www.snowflake.com/about/>
- [23] Confluent KSQL. <https://www.confluent.io/product/ksql/>
- [24] Apache Atlas. <https://atlas.apache.org/>

# Using MILP (Mixed Integer Linear Programming) for RF Bandwidth Optimization

A Technical Paper prepared for SCTE•ISBE by

**Tom Holloran**

Principal Operations Research Scientist  
Charter Communications  
14810 Grasslands Dr, Englewood, CO 80112  
720.518.2296  
thomas.j.holloran@charter.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
System and Process Description .....	4
1. Business Requirements .....	4
1.1. Minimization of the number of 6 MHz EIA slots .....	5
1.2. Assignment of voice and data individual content to preferred frequency locations .....	5
1.3. Assignment of voice and data content according to standardized templates.....	5
1.4. Minimization of moves/changes required to use only the minimum number of EIA slots.....	5
1.5. Adherence to engineering constraints .....	5
1.6. System and Process Design .....	6
2. Technical Approach .....	6
2.1. Optimization versus Heuristic Approach.....	6
2.2. Data Driven Approach .....	8
2.3. Finance-infused Cost Metrics.....	8
2.4. Data Repository Creation.....	9
2.5. ROI Tracking.....	9
3. RF Bandwidth Optimization System .....	10
3.1. Modeling Platform.....	10
3.2. Control Tables.....	11
3.3. Optimized RF Channel Map.....	15
3.4. Results Visualization.....	16
3.5. Execution Modes.....	20
4. Business Integration and Automation .....	20
Conclusions.....	21
Abbreviations.....	22
Bibliography & References .....	22

## List of Figures

Title	Page Number
Figure 1 - RF Spectral Bandwidth Mathematical Formulation .....	7
Figure 2 - RF Channel Map Contents.....	8
Figure 3 - System Configuration Overview .....	9
Figure 4 - Analytics Agile Development Lifecycle .....	10
Figure 5 - Mathematical Model Visualization .....	11
Figure 6 - Operational RF Channel Map.....	12
Figure 7 - Bandwidth Demand Control Table.....	13
Figure 8 - Patterns Control Table .....	13
Figure 9 - Preferences Control Table .....	14
Figure 10 - Moves Control Table.....	14
Figure 11 - Levers Control Table.....	15
Figure 12 - Optimization Results Summary Table Example .....	16
Figure 13 - Sankey Diagram of Content Moves Required.....	17
Figure 14 - Detailed Content Re-Assignment Table Sample.....	18
Figure 15 - EIA Slot Visualization Example .....	19
Figure 16 - Data Flow Example.....	20



# Introduction

While DOCSIS software advances and video analog to digital conversions have dramatically increased the amount of video and data that can be delivered across a MHz of spectrum, annual IP data growth rates of 40-60% have made a 6 MHz EIA slot of spectrum an increasingly valuable commodity. Managing MSO product delivery on a fixed amount of RF spectrum has also become very challenging -- to the point that each slot of spectrum requires management like any other resource in the MSO supply chain.

This paper describes a set of algorithms and a supply chain process that can be used to identify the absolute minimum number of EIA slots that are required to meet the demand for any set of existing and future linear and switched video, digital audio, and IP data. Additionally, the algorithms can identify the minimum number of changes that are required to move content from existing spectrum locations into the identified minimum EIA slot allocation, and can maximize the placement preferences for content within defined ranges of spectrum.

The optimization algorithms of this paper are based on integer linear programming. However, they are implemented using a data-driven approach where tables containing engineering rules and location preferences drive all mathematical model generation. The underlying optimization system allows the modeler to treat content and EIA slots as supply chain commodities. Detailed changes to the optimization model can be made without in-depth knowledge or training in mathematical programming optimization techniques.

The algorithms and process contained in this paper are not limited in any way by the hardware and software required for implementation. Most mathematical optimization software packages that have a generalized integer programming capability can be used to implement the algorithms. A variety of software packages exist that can be used to visualize optimized RF Channel Maps. And a variety of software programming languages will facilitate the required functions of pattern generation and solution fitting.

## System and Process Description

### 1. Business Requirements

The business requirements for RF bandwidth optimization include the following:

- Minimization of the number of 6 MHz EIA slots being used for voice and data in each RF Channel Map
- Assignment of voice and data individual content to the preferred frequency locations
- Assignment of voice and data content to frequency ranges to those defined in standardized RF channel map templates
- Minimization of content movement from existing locations in order to accommodate minimum slot utilization and location into standardized/preferred frequency ranges
- Maximization of quality and capacity by adherence to all engineering constraints
- Implementation of a system and process that will reduce the time required to design RF channel map reconfigurations, that can be used by a large number of users, that does not require knowledge of advanced mathematics to use it, and enables use for both tactical (short-term) and strategic (longer term) planning.

### **1.1. Minimization of the number of 6 MHz EIA slots**

Minimizing the number of EIA channels, total frequency, or bandwidth required is the primary goal of spectral bandwidth optimization. By minimizing the bandwidth in use, EIA channels are opened up to accommodate video and/or data growth. Identifying the minimum number of EIA slots determines when plants need to be upgraded, how upgrades can be avoided, what/when architectural changes affecting capacity need to be made, etc. The EIA slot is the highest value asset assumed in optimizing bandwidth.

### **1.2. Assignment of voice and data individual content to preferred frequency locations**

Given a specific minimum number of EIA slots required, a secondary business requirement is placing voice and data content into preferred frequencies, and avoiding certain reserved frequencies. Preferred locations may be designated for quality reasons, or they may just be preferences to standardize locations across plants for ease of maintenance/continuity.

### **1.3. Assignment of voice and data content according to standardized templates**

A business requirement to move toward standardizing the locations of certain types of content makes broad engineering changes or capacity upgrades easier to implement. Standardized template implementation in bandwidth optimization is handled as a preference, a secondary objective to minimizing the amount of bandwidth that a set of content consumes. However, the business requirement for standardization does allow for setting preference costs high enough that they become more important than other location preferences, or even total EIA slot bandwidth.

### **1.4. Minimization of moves/changes required to use only the minimum number of EIA slots**

Another business requirement in optimizing bandwidth is minimizing the amount of reconfiguration and content movement that is necessary to use only the minimum number of EIA slots. There are alternative optimum (alternative configurations that use the minimum number of EIA slots), so this business requirement specifically addresses identifying the set of content assignments that minimize bandwidth use but also result in the minimum disruption to the existing RF channel map.

### **1.5. Adherence to engineering constraints**

Adherence to engineering constraints may be an obvious business requirement. But it is important that they are not overlooked as they definitely effect any bandwidth optimization solution. Linear video content may have attached contractual constraints. Engineering may prefer contiguous blocks of DOCSIS be a certain width or minimum width based on capacity testing. DOCSIS 3.1 OFDM blocks must be contiguous to take advantage of their inherent design. Operations capacity engineers may need open EIA slots in a location that is planned to be expanded in the near future. Video engineers may require SDV and VOD pools to be adjacent to facilitate future changes. A bandwidth optimization system needs to be developed to accommodate these types of engineering constraints, but to the extent possible, must also anticipate that other constraints will be required as architectures and software change over time.

## 1.6. System and Process Design

Business requirements for system and process design can be summarized in the following bullet points for design goals:

- System and process design should reduce the time required to identify needed bandwidth reconfigurations
- System and process design should be capable of accommodating a large number of users
- System and process design should not require advanced math knowledge to operate
- System and process design should accommodate both tactical (short-term) and strategic (longer-term) planning
  - Use by individual head end engineers on individual channel maps, accommodating local requirements for reserved frequencies, use of rolloff frequencies, etc.
  - Use by strategic corporate planners to identify minimum bandwidth upgrades required system wide by certain business planning scenarios, to produce 5-year budgeting projections on service group splits, plant upgrades, etc, to estimate the impact of traffic growth projections and architectural design changes, and other corporate-wide business initiatives

## 2. Technical Approach

The RF bandwidth optimization technical approach is based on the following:

- Use of optimization vs. heuristics
- Flexible and streamlined processes through data-driven design
- Cost metrics derived by Finance corporate
- Development of a data repository
- Ongoing ROI tracking and management

### 2.1. Optimization versus Heuristic Approach

A heuristic approach uses best practices or approximation techniques to identify a good solution from among the set of all alternatives. An optimization approach evaluates all alternatives and identifies the absolute best obtainable solution from among that full set of alternatives. Simply evaluating all possible alternatives (explicitly enumeration) is the most straightforward method of optimization. There are also mathematical optimization techniques like linear programming which can implicitly (versus explicitly) evaluate all alternatives through the use of convergence algorithms. They are able to identify the absolute best obtainable alternative without having to evaluate all alternatives individually.

The optimization approach was chosen over a heuristic approach due to the high value placed on each 6MHz slot of spectrum. We want to know that the absolute minimum bandwidth requirement has been identified. The mixed integer linear programming optimization approach was chosen due to the sheer number of alternatives that must be evaluated to guarantee the minimum bandwidth requirement has been found. If there are 50 patterns (unique sets of content that can be placed on an EIA slot) and 116 EIA slots (750 MHz plant, for example), then the total number of pattern to EIA slot combinations is 50\*116 or 5,800. Each pattern combination is either chosen or not chosen, resulting in a combined set of  $2^{5800}$  sets of assignments that must be evaluated to determine optimality.

The mixed integer linear programming optimization approach uses the mathematical formulation shown in Figure 1. It is based on a cutting stock formulation that has been used in manufacturing for decades. Because of its cutting stock structure, extremely large combinatorial problems can be solved very quickly. In the case of RF bandwidth optimization, the optimization algorithm generally takes less than 10 seconds on a standard-issue laptop to identify an optimal solution from among  $2^{5800}$  possible alternatives.

The interpretation of the mathematical formulation in Figure 1 can be stated simply as:

Minimize the number of required EIA slots while meeting as many assignment preferences as possible in the bandwidth available, and fitting as much content as possible into the bandwidth available.

Minimize:	$\sum_{i=1}^M \sum_{j=1}^N c_{ij} x_{ij} + \sum_{i=1}^M e_i y_i + \sum_{k=1}^K f_k z_k$
Subject to:	$\sum_{i=1}^M \sum_{j=1}^N a_{jk} x_{ij} + z_k \geq D_k \quad \forall k, k = 1, \dots, K$ $\sum_{j=1}^N x_{ij} - 1000 y_i \leq 0 \quad \forall i, i = 1, \dots, M$ $\sum_{j=1}^N x_{ij} \leq 1 \quad \forall i, i = 1, \dots, M$
where:	<p><math>y_i = 0</math> or <math>1</math> (1 if EIA <math>i</math> is used, 0 if EIA <math>i</math> is left open)</p> <p><math>x_{ij} = 0</math> or <math>1</math> (1 if pattern <math>j</math> is assigned to EIA <math>i</math>, 0 otherwise)</p> <p><math>z_k</math> = the number of content type <math>k</math> assignments that cannot be accommodated in the channel map</p> <p><math>c_{ij}</math> = preference cost plus the cost of changing content assignments when pattern <math>j</math> is assigned to EIA <math>i</math>,</p> <p><math>e_i</math> = cost of using EIA slot <math>i</math>,</p> <p><math>f_k</math> = cost of not being able to accommodate a unit of demand for content assignment type <math>k</math>,</p> <p><math>a_{jk}</math> = number of content assignments of type <math>k</math> in pattern <math>j</math>,</p> <p><math>D_k</math> = number of content assignments of type <math>k</math> required in the RF Channel Map</p> <p><math>N</math> = number of individual patterns (sets of content assignments that fill an EIA slot) that can be used</p> <p><math>M</math> = number of EIA's available for possible use</p> <p><math>K</math> = total number of content assignment types (e.g. HD, SD, DA, SDV, VOD, DOCSIS, RESERVED)</p>

**Figure 1 - RF Spectral Bandwidth Mathematical Formulation**

## 2.2. Data Driven Approach

An approach that is data driven simply implies that by setting up tables of differing size and content, a customized mathematical model can be constructed and solved. The contents of the tables drive varying sized of RF channel maps as well as varying sets of engineering constraints and operational preferences. Data driven also implies that a model that is constructed will also always solve. No infeasible solutions are possible. Identification of a shortfall in bandwidth or an inability to accommodate certain content due to lack of allowable assignment patterns is merely identified as part of the optimization results.

Figure 2 contains part of an RF channel map. Part of adopting a data driven approach means that total content bandwidth requirements, total number of EIA slots in the channel map, total moves required to change to a minimum set of EIAs, etc. can and will all be derived from the read of the channel map.

Hub CLLI	ChannelMap	EIAChannel	CurrentAllocation	Bandwidth	Bandwidth Type
TRCYMI21713	Traverse City Region, MI (tc_om01m)	2	HBO 2 HD West	9.5	HD4
TRCYMI21713	Traverse City Region, MI (tc_om01m)	2	HBO Family HD West	9.5	HD4
TRCYMI21713	Traverse City Region, MI (tc_om01m)	2	HBO HD West	9.5	HD4
TRCYMI21713	Traverse City Region, MI (tc_om01m)	2	HBO Signature HD West	9.5	HD4
TRCYMI21713	Traverse City Region, MI (tc_om01m)	11	Charter Mainstreet	2.75	SD
TRCYMI21713	Traverse City Region, MI (tc_om01m)	11	Local Access	2.75	SD
TRCYMI21713	Traverse City Region, MI (tc_om01m)	11	PBS	2.75	SD
TRCYMI21713	Traverse City Region, MI (tc_om01m)	11	PEG Access	2.75	SD
TRCYMI21713	Traverse City Region, MI (tc_om01m)	11	Public Access	2.75	SD
TRCYMI21713	Traverse City Region, MI (tc_om01m)	15	Independent TV Station HD	19.4	HD2
TRCYMI21713	Traverse City Region, MI (tc_om01m)	15	PBS HD	19.4	HD2
TRCYMI21713	Traverse City Region, MI (tc_om01m)	16	CLI analog carrier	38.8	RESERVED
TRCYMI21713	Traverse City Region, MI (tc_om01m)	17	Comedy Central HD East	12.75	HD3
TRCYMI21713	Traverse City Region, MI (tc_om01m)	17	NBCSN HD (NBC Sports Network)	12.75	HD3
TRCYMI21713	Traverse City Region, MI (tc_om01m)	17	TBS HD East (Turner Broadcasting System)	12.75	HD3
TRCYMI21713	Traverse City Region, MI (tc_om01m)	18	MC 70s	0.4	MC
TRCYMI21713	Traverse City Region, MI (tc_om01m)	18	MC 80s	0.4	MC
TRCYMI21713	Traverse City Region, MI (tc_om01m)	18	MC 90s	0.4	MC
TRCYMI21713	Traverse City Region, MI (tc_om01m)	61	Analog Pilot	38.8	ANALOG
TRCYMI21713	Traverse City Region, MI (tc_om01m)	62	VOD Channel	38.8	VOD
TRCYMI21713	Traverse City Region, MI (tc_om01m)	63	VOD Channel	38.8	VOD
TRCYMI21713	Traverse City Region, MI (tc_om01m)	74	AVN	38.8	RESERVED
TRCYMI21713	Traverse City Region, MI (tc_om01m)	75	AVN	38.8	RESERVED
TRCYMI21713	Traverse City Region, MI (tc_om01m)	76	OPEN	38.8	OPEN
TRCYMI21713	Traverse City Region, MI (tc_om01m)	77	Zodiac	38.8	RESERVED
TRCYMI21713	Traverse City Region, MI (tc_om01m)	79	SDV Channel	38.8	SDV
TRCYMI21713	Traverse City Region, MI (tc_om01m)	80	SDV Channel	38.8	SDV
TRCYMI21713	Traverse City Region, MI (tc_om01m)	87	DOCSIS Channel	38.8	DOCSIS3.0
TRCYMI21713	Traverse City Region, MI (tc_om01m)	88	DOCSIS Channel	38.8	DOCSIS3.0
TRCYMI21713	Traverse City Region, MI (tc_om01m)	89	DOCSIS Channel	38.8	DOCSIS3.0
TRCYMI21713	Traverse City Region, MI (tc_om01m)	90	DOCSIS Channel	38.8	DOCSIS3.0
TRCYMI21713	Traverse City Region, MI (tc_om01m)	91	DOCSIS Channel	38.8	DOCSIS3.0
TRCYMI21713	Traverse City Region, MI (tc_om01m)	92	DOCSIS Channel	38.8	DOCSIS3.0
TRCYMI21713	Traverse City Region, MI (tc_om01m)	93	DOCSIS Channel	38.8	DOCSIS3.0
TRCYMI21713	Traverse City Region, MI (tc_om01m)	94	DOCSIS Channel	38.8	DOCSIS3.0

**Figure 2 - RF Channel Map Contents**

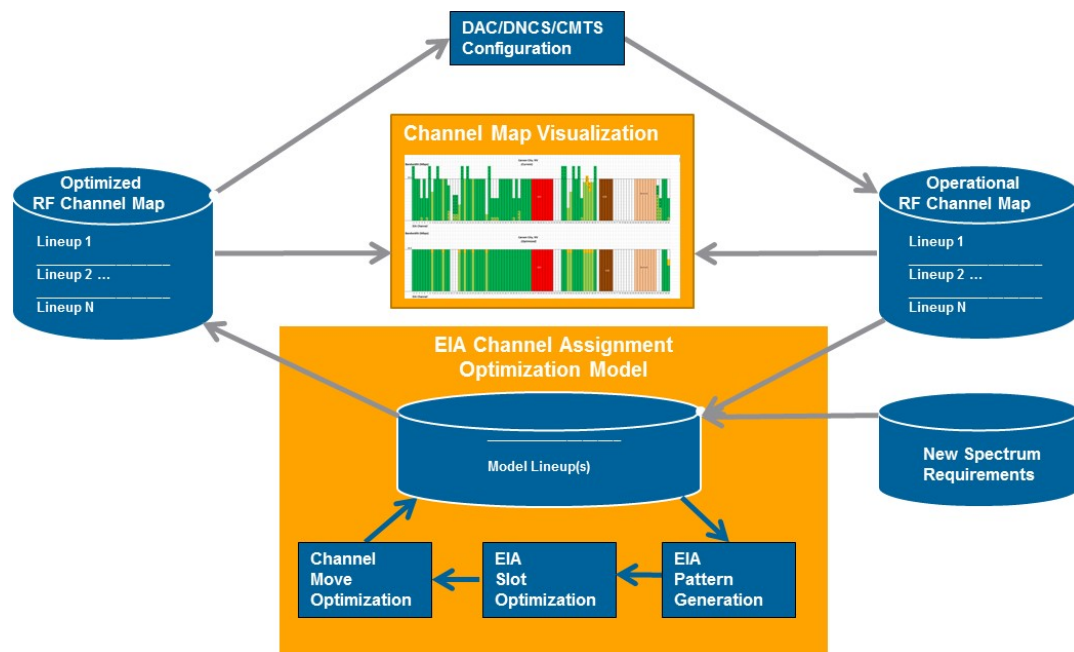
## 2.3. Finance-infused Cost Metrics

A technical approach that uses cost metrics derived from corporate-backed financial reporting is very important to identifying appropriate tradeoffs in bandwidth design. It is also important to accurately track improvements ongoing and justify the use of an optimization approach. Finance metrics leveraged should include CAPEX, OPEX, EBITDA, etc. The key metric in RF bandwidth optimization is the cost of a 6 MHz EIA slot. Reconfiguration FTE and hardware/software costs are also important. Derivation of future costs required for longer-term strategic planning should also be based on hardened corporate financial projections whenever possible.

## 2.4. Data Repository Creation

The technical approach to RF bandwidth optimization should incorporate creation of a data repository. Automation of data creation/read and configuration changes/write is ideal, but creation and maintenance of a data repository of RF channel map configurations and cost in any form is an absolute requirement for bandwidth optimization to be successful.

Figure 3 shows a system overview containing a data repository and its interfaces.

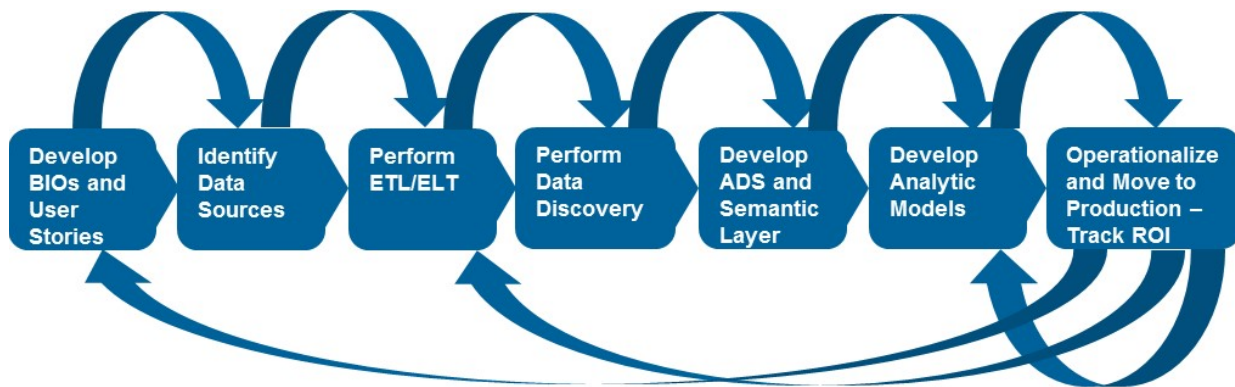


**Figure 3 - System Configuration Overview**

## 2.5. ROI Tracking

ROI tracking of improvements is a key component to justifying the investment in optimization. It is also important to identify potential changes and improvements that can be gained based on costs incurred and/or improved.

Successful creation of an RF bandwidth optimization system is best accomplished using agile development. Figure 4 shows the closed-loop agile analytics process adopted as part of the RF bandwidth optimization technical approach.



**Figure 4 - Analytics Agile Development Lifecycle**

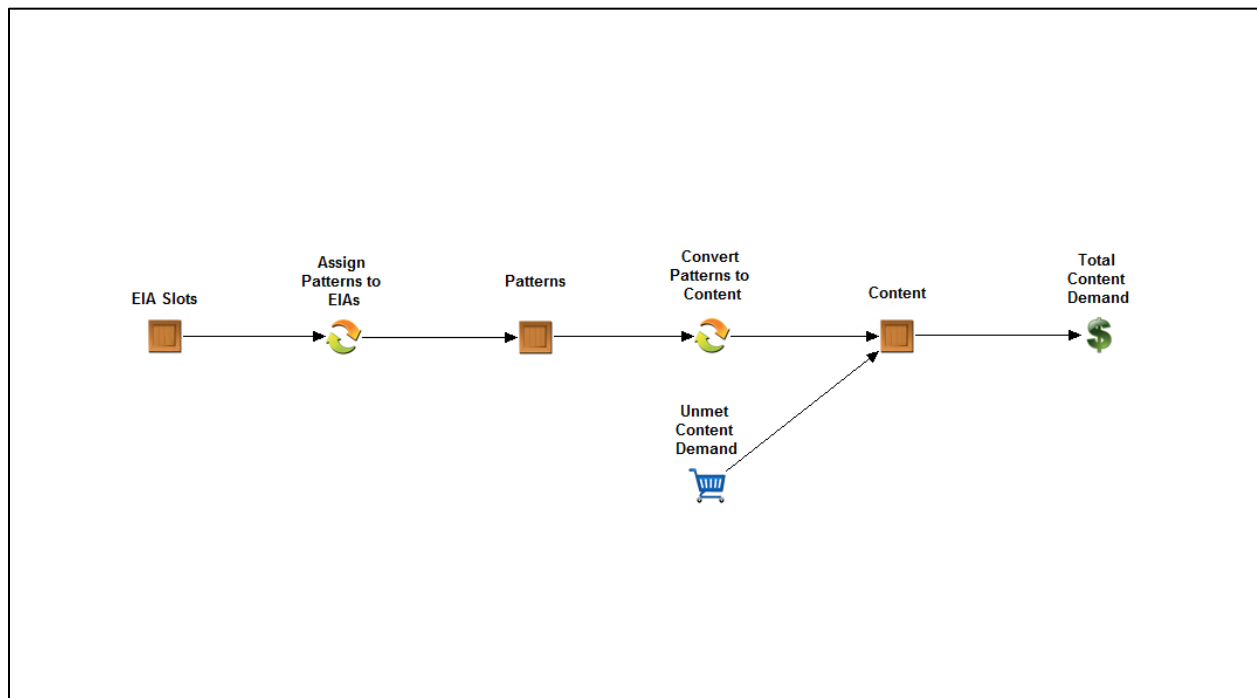
### 3. RF Bandwidth Optimization System

Major components of the RF bandwidth optimization system consist of an optimization modeling platform, a set of driving control tables, bandwidth assignment visualization tools, and the method of model data / model execution management.

#### 3.1. Modeling Platform

The modeling platform consists of a mathematical model generator, and optimizer, and a solution formatter. The model generator and solution formatter can be created using a wide variety of programming languages and/or data blending tools. The optimization module requires commercial mathematical programming software to solve the larger spectral bandwidth problems. A 750 MHz plant can easily produce 3000+ alternative assignments or  $2^{3000}$  combinations to be evaluated in order to identify an optimal set of assignments. Fortunately commercial mathematical programming software packages are capable of solving these problems in less than 1 minute. In the RF Bandwidth Optimization system, model optimization averages only about 10 seconds.

A graphical representation of the optimization model is shown in Figure 5. In supply chain terms an inventory of EIA slots is established, EIA slots are cut into patterns representing all combinations of allowable sets of content that will fit on an EIA, EIA patterns are then cut into individual pieces of content that are matched against the demand required to accommodate a defined RF channel map. The flow is then submitted to the optimizer which determines the minimum number of EIA slots that is required to meet all content demand, and at least cost.



**Figure 5 - Mathematical Model Visualization**

### 3.2. Control Tables

There are six control tables that define the RF bandwidth being modeled along with all of the content demand requirement and engineering/operations preferences for assigning content to bandwidth.



ChannelMap									
ChannelMap	EIAChannel	Type	CurrentAllocation						
Traverse City Region, MI 750	73	SD	In Demand PPV						
Traverse City Region, MI 750	73	SD	Showtime West						
Traverse City Region, MI 750	73	SD	Starz East						
Traverse City Region, MI 750	73	SD	ThrillerMax East						
Traverse City Region, MI 750	74	RESERVED	AVN						
Traverse City Region, MI 750	75	RESERVED	AVN						
Traverse City Region, MI 750	76	OPEN	OPEN						
Traverse City Region, MI 750	77	RESERVED	Zodiac						
Traverse City Region, MI 750	78	SD	C-Span 3						
Traverse City Region, MI 750	78	SD	Discovery Channel East						
Traverse City Region, MI 750	78	SD	ESPN 2						
Traverse City Region, MI 750	78	SD	FX East						
Traverse City Region, MI 750	78	SD	HGTV East (Home & Garden Television)						
Traverse City Region, MI 750	79	SDV	SDV Channel						
Traverse City Region, MI 750	80	SDV	SDV Channel						
Traverse City Region, MI 750	81	SDV	SDV Channel						
Traverse City Region, MI 750	82	SDV	SDV Channel						
Traverse City Region, MI 750	83	SDV	SDV Channel						
Traverse City Region, MI 750	84	SDV	SDV Channel						
Traverse City Region, MI 750	85	SDV	SDV Channel						
Traverse City Region, MI 750	86	SDV	SDV Channel						
Traverse City Region, MI 750	87	DOCSIS 3.0	DOCSIS Channel						
Traverse City Region, MI 750	88	DOCSIS 3.0	DOCSIS Channel						
Traverse City Region, MI 750	89	DOCSIS 3.0	DOCSIS Channel						
Traverse City Region, MI 750	90	DOCSIS 3.0	DOCSIS Channel						
Traverse City Region, MI 750	91	DOCSIS 3.0	DOCSIS Channel						
Traverse City Region, MI 750	92	DOCSIS 3.0	DOCSIS Channel						
Traverse City Region, MI 750	93	DOCSIS 3.0	DOCSIS Channel						
Traverse City Region, MI 750	94	DOCSIS 3.0	DOCSIS Channel						
Traverse City Region, MI 750	95	HD4	Bloomberg TV HD						

**Figure 6 - Operational RF Channel Map**

The Operational RF Channel Map identifies all of the content currently being accommodated on the plant along with its EIA slot location. Total content is identified as the base content requirement to be optimized along with the number of EIA channels (spectrum size) in the channel map to be optimized.

Figure 6 shows a sample portion of an Operational RF Channel Map – containing rows identifying plant name, EIA slot, content type, and specific content name.

Bandwidth Demand			
	Existing Demand	Additional Demand	Total Demand
SD	138	0	138
SD2	0	0	0
HD	0	0	0
HD2	6	0	6
HD3	31	0	31
HD4	121	0	121
DA	44	0	44
DOCSIS 3.0	24	0	24
DOCSIS 3.1	0	16	16
SDV	8	0	8
VOD	6	0	6
SDV-VOD	0	0	0
RESERVED	6	0	6
OPEN	0	0	0

**Figure 7 - Bandwidth Demand Control Table**

A bandwidth demand control table like that shown in Figure 7 is generated based on information in the RF channel map of Figure 6.

The bandwidth demand control table contains an additional demand column where content demand derived from the RF channel map can be modified (added to or subtracted from) to model future demand or account for variations from a base map.

Patterns											
	Contiguous	SD	SD2	HD	HD2	HD3	HD4	DA	D3.0	D3.1	
Pattern01:2-0-0-0-0	1	0	0	0	2		0	0	0	0	
Pattern02:1-1-0-2-0	1	2	0	0	1	1	0	0	0	0	
Pattern03:1-1-0-1-8	1	1	0	0	1	1	0	8	0	0	
Pattern04:1-1-0-0-13	1	0	0	0	1	1	0	13	0	0	
Pattern05:1-0-2-0-0	1	0	0	0	1	0	2	0	0	0	
Pattern06:1-0-1-3-0	1	3	0	0	1	0	1	0	0	0	
Pattern07:1-0-1-2-8	1	2	0	0	1	0	1	8	0	0	
Pattern08:1-0-1-1-14	1	1	0	0	1	0	1	14	0	0	
Pattern09:1-0-1-0-19	1	0	0	0	1	0	1	19	0	0	
Pattern10:1-0-0-6-0	1	6	0	0	1	0	0	0	0	0	

**Figure 8 - Patterns Control Table**

A patterns control table like the sample shown in Figure 8 is also generated off of the RF channel map contents. Every combination of content that can be accommodated on a single EIA slot is identified as a pattern. Prior to executing the optimization, undesired patterns can be removed from the table. The patterns control table Contiguous column is used to identify the number of an individual pattern type that must exist adjacent to each other.

Preferences	EIA002	EIA003	EIA004	EIA005	EIA006	EIA095	EIA096	EIA097	EIA098	EIA099	EIA014	EIA015	EIA016
Pattern01:2-0-0-0-0	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern02:1-1-0-2-0	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern03:1-1-0-1-8	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern04:1-1-0-0-13	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern05:1-0-2-0-0	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern06:1-0-1-3-0	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern07:1-0-1-2-8	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern08:1-0-1-1-14	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern09:1-0-1-0-19	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern10:1-0-0-6-0	0	0	0	0	0	0	0	0	0	0	0	0	1000
Pattern81:DOCSIS3.0	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000
Pattern82:DOCSIS3.1	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000
Pattern83:SDV	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000
Pattern84:VOD	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000
Pattern85:RESERVED	1000	1000	1000	1000	1000	1000	0	1000	1000	1000	1000	1000	0
Pattern86:OPEN	0	0	0	0	0	0	0	0	0	0	0	0	0

**Figure 9 - Preferences Control Table**

The preferences control table is also generated from the base RF channel map of Figure 6. It contains a preference cost for assigning an individual set of content (pattern) to each EIA. The set of patterns is based on the Patterns table of Figure 8. Default preference costs are placed into the table, but can be modified as needed/desired.

This is the control table that is used to bias alternative minimum EIA optimal solutions toward preferred frequency locations for content, or toward established standardization guidelines.

Moves	EIA002	EIA003	EIA004	EIA005	EIA006	EIA095	EIA096	EIA097	EIA098	EIA099	EIA014	EIA015	EIA016
Pattern01:2-0-0-0-0	4	4	4	4	4	4	3	1	2	2	12	3	0
Pattern02:1-1-0-2-0	4	4	4	4	4	4	2	1	2	2	10	3	1
Pattern03:1-1-0-1-8	4	4	4	4	4	4	2	1	2	2	11	3	1
Pattern04:1-1-0-0-13	4	4	4	4	4	4	3	1	2	2	12	3	1
Pattern05:1-0-2-0-0	2	2	2	2	2	2	1	1	0	0	12	1	1
Pattern06:1-0-1-3-0	3	3	3	3	3	3	1	1	1	1	9	2	1
Pattern07:1-0-1-2-8	3	3	3	3	3	3	1	1	1	1	10	2	1
Pattern08:1-0-1-1-14	3	3	3	3	3	3	1	1	1	1	11	2	1
Pattern09:1-0-1-0-19	3	3	3	3	3	3	2	1	1	1	12	2	1
Pattern10:1-0-0-6-0	4	4	4	4	4	4	2	1	2	2	6	3	1

**Figure 10 - Moves Control Table**

A moves control table is also auto-generated based on the RF channel map initial assignments – identifying the number of content assignments that will change if a specific set of content (pattern) is assigned to an individual EIA.

This table is used to bias alternative optimal minimum EIA solutions toward those with the least impact/disruption. The weight that is placed on biasing toward move minimization is set through a single global move cost in the Levers control table of Figure 11.

Levers					
Lever		Value		Description	
				Optimization Options:	
				1=Minimize the number of EIA	
				channels required,	
				2=Minimize the number of EIAs	
				required, the cost of Moves/re-	
				and the Preference	
				costs associated with placing	
Objective		1		individual patterns on individual EIAs	
EIA Channel Value		\$100,000		The value of each EIA channel	
				The unit cost per move associated	
				with moving or re-assigning a	
				currently allocated unit from its	
				assigned EIA channel to another EIA	
Move Unit Cost		\$1,000		channel	
				The minimum number of EIA channels	
				to remain in use. A value greater	
				than the minimum number of EIA	
				relaxes the model	
Minimum EIA Channels		1		objective.	

**Figure 11 - Levers Control Table**

This Levers control table of Figure 11 is used to set optimization modeling objectives and assign costs to EIA slots and individual content moves.

Depending on the level of detail being analyzed in the optimization, the Objective can be set to 1) simply identify the minimum number of EIA slots required to accommodate bandwidth demand or 2) to identify the minimum EIA slot solution that also maximizes content assignment preferences and minimizes changes from existing content assignment locations.

EIA Channel Value is used to indicate potential savings through minimum EIA slot use, and also to weight the relative importance of minimizing the number of EIA slots versus the importance of minimizing disruption and maximizing assignment location preferences.

By manipulating the value of the Move Unit Cost, different levels of emphasis on disruption can be defined.

By manipulating the Minimum EIA Channels value, re-assignment of content can be reduced at the cost of using more than the minimum number of EIA slots. This is specifically accomplished by defining a Minimum EIA Channel count that is greater than the minimum number of EIA slots required to accommodate the content – established in a previous run of the optimization model.

### 3.3. Optimized RF Channel Map

The EIA Channel Assignment Optimization Model system component of Figure 3 generates a mathematical model based on the control tables, optimizes the model using a mixed integer linear programming optimizer, and builds an Optimized RF Channel Map based on the resulting optimized solution.

The Optimized RF Channel Map contains the same content as the original Operational RF Channel Map of Figure 6 (rows identifying plant name, EIA slot, content type, and specific content name) – but reflects the optimized assignments of content to EIA – assignments that minimize EIA slot utilization, maximize preferred location of content, etc.

### 3.4. Results Visualization

A standard summary of optimization model results is shown in Figure 12. It shows that through optimization, an additional 13 EIA slots of spectrum can be freed up. The minimum number of moves required in order to minimize EIA slot use is 35. If freeing an EIA slot is valued at \$100K then the total resulting improvement is shown as \$1.3 million, or \$1.265 million if each content move reduces the benefit by \$1K. The Optimal Patterns table identifies the various EIA slot configurations that would be used to produce the minimum EIA slot configuration. For example, 28 EIA slots would be comprised of Pattern 44 which contains 4 HD4 content providers on a single EIA. The Solution Type Usage table shows that 18 total EIA slots would be open in the optimized RF channel map and 28 additional digital audio channels could be accommodated within the set of patterns chosen by the model.

Scenario Traverse City Region, MI 750						
EIA Channels	EIA Channels In Use	EIA Channels Not Used	Solution EIA Channels In Use	Solution EIA Channels Not Used	Solution EIA Channel Improvement	Solution EIA Channel Cost Improvement
<b>116</b>	<b>111</b>	<b>5</b>	<b>98</b>	<b>18</b>	<b>13</b>	<b>\$1,300,000</b>
EIA Channels Not Changed	EIA Channels Changed	Total Moves	Moves Cost	Solution Net Value		
<b>79</b>	<b>37</b>	<b>35</b>	<b>35,000</b>	<b>\$1,265,000</b>		

Solution Type Usage			
Type	Used Capacity	Unused Capacity	Unmet Demand
SD	138	0	0
SD2	0	0	0
HD	0	0	0
HD2	6	0	0
HD3	31	0	0
HD4	121	0	0
DA	44	28	0
DOCSIS 3.0	24	0	0
DOCSIS 3.1	0	0	0
SDV	8	0	0
VOD	6	0	0
SDV-VOD	0	0	0
RESERVED	5	0	0
OPEN	0	18	0


Optimal Patterns	
Pattern Name	Pattern Count
Pattern44:0-0-4-0-0	28
Pattern81:DOCSIS3.0	24
Pattern86:Open	18
Pattern17:0-3-0-0-0	9
Pattern83:SDV	8
Pattern84:VOD	6
Pattern85:Reserved	5
Pattern67:0-0-0-13-0	4
Pattern56:0-0-1-10-0	3
Pattern01:2-0-0-0-0	2
Pattern68:0-0-0-12-8	2
Pattern05:1-0-2-0-0	1
Pattern10:1-0-0-6-0	1
Pattern18:0-2-1-1-0	1
Pattern25:0-1-2-2-0	1
Pattern34:0-1-0-9-0	1

Figure 12 - Optimization Results Summary Table Example

While a standard summary table and revised RF channel map can fully define optimization system results, further visualization is important when exercising the optimization system iteratively and when attempting to identify the best solution from among alternative optimal solutions. Results visualization is also very important to gaining acceptance of optimized solutions, quantifying benefits, and recognizing the impact that changes will have on moving to an optimized re-allocation.

There are two sets of visualizations that are particularly beneficial – visualization of detailed content moves/re-assignments, and visualization of before/after EIA slot allocations across the full plant spectral bandwidth.

Detailed content changes required to implement an optimized configuration can be visualized at a high-level using a Sankey diagram. A sample Sankey diagram is shown in Figure 13.



**Figure 13 - Sankey Diagram of Content Moves Required**

The Sankey diagram identifies visually each of the changes that would be required (individual lines on the diagram) and the relative amount of bandwidth associated with each change (width of each line).

Low-level detailed content changes can best be visualized in tabular form. A sample portion of a move/re-assignment table for an RF channel map is shown in Figure 14.

From	To	Program Name	Program Type	Bandwidth Lineup
From 0029	To 0019	IndiePlex HD East	HD4	9.5
From 0029	To 0026	movieplex HD	HD4	9.5
From 0029	To 0051	RetroPlex HD	HD4	9.5
From 0030	To 0043	ABC Family HD East	HD4	9.5
From 0030	To 0046	MTV HD East	HD4	9.5
From 0030	To 0055	Nick HD East	HD4	9.5
From 0031	To 0024	CNN HD	HD3	12.8
From 0031	To 0042	Discovery Channel HD East	HD3	12.8
From 0031	To 0095	USA Network HD East	HD3	12.8
From 0032	To 0011	Animal Planet HD	HD3	12.8
From 0032	To 0012	Big Ten Network HD	HD3	12.8
From 0032	To 0056	TLC HD East (The Learning Channel)	HD3	12.8
From 0033	To 0028	History HD East	HD4	9.5
From 0033	To 0038	FX HD East	HD4	9.5
From 0033	To 0055	National Geographic Channel HD East	HD4	9.5
From 0034	To 0040	Bravo HD East	HD4	9.5
From 0034	To 0044	MSNBC HD	HD4	9.5
From 0034	To 0098	truTV HD East	HD4	9.5
From 0035	To 0045	ESPN 2 HD	HD3	12.8
From 0035	To 0053	ESPN HD	HD3	12.8
From 0036	To 0014	AMC HD East	HD4	9.5
From 0036	To 0049	Spike TV HD East	HD4	9.5
From 0036	To 0097	Travel Channel HD East	HD4	9.5
From 0070	To 0013	SundanceTV HD East	HD4	9.5
From 0070	To 0048	Smithsonian Channel HD	HD4	9.5
From 0070	To 0098	Velocity HD	HD4	9.5
From 0071	To 0011	HBO Signature East	SD	2.8
From 0071	To 0011	MoreMax East	SD	2.8
From 0071	To 0024	Showtime 2 East	SD	2.8
From 0071	To 0024	Showtime Extreme East	SD	2.8
From 0071	To 0024	Showtime Showcase East	SD	2.8
From 0071	To 0024	SundanceTV East	SD	2.8
From 0071	To 0024	TeenNick	SD	2.8
From 0071	To 0025	HBO 2 East	SD	2.8


**Figure 14 - Detailed Content Re-Assignment Table Sample**

The table shows each individual content assignment's original location in the first column (From) and its optimized location in the second column (To).

Before/after visualization of full plant spectral bandwidth can be viewed using bar charts. Using bar charts, the bandwidth differences between current RF channel assignments and optimization alternative assignments can be quickly visualized.

A bar chart format shows how the full spectrum is being utilized and the location of content assignments relative to each other.





**Figure 15 - EIA Slot Visualization Example**

Figure 15 shows an example of how EIA bar chart visualization can be used to identify differences in channel map allocations.

The top bar chart of Figure 15 shows the initial content assignments of a channel map which was using 111 of 116 available EIA slots. The different types of content are identified by color. The height of each bar indicates how much of each 6 MHz EIA slot's capacity is being utilized.

The middle bar chart shows a minimum EIA slot assignment which has freed up 13 additional EIAs by reassigning content – using only 98 of the 116 available EIA slots.

The lower EIA bar chart of Figure 15 shows an alternative minimum EIA slot configuration that has been biased toward placing content in alignment with a desired standardized preference template – while still adhering to the minimum number of EIA slots. The preference template is defined by the Preferences Control Table of Figure 9.

Viewing the three bar charts in combination, comparisons can be made to see how much EIA utilization has improved with the EIA slot minimization, how much content has shifted positioning in the spectrum, and how much the optimized assignments were able to adhere to the engineering preferences defined in the control tables.



### 3.5. Execution Modes

The RF bandwidth optimization model is designed to be executed according to the data flow shown in Figure 16.

Data blending software is used to calculate the number of moves that will be required if an individual pattern is placed on an existing EIA slot. It is also used to generate the bandwidth demand by content type from an original Operational RF Channel Map. It is also used to create the mathematical matrix input to the commercial mixed integer programming optimizer.

The MILP (Mixed Integer Linear Programming) commercial optimizer then optimizes the model and produces a set of tables reflecting an optimized RF Channel Map and its associated chosen solution summary results.

Visualization tools then place optimization results into summary results tables, Sankey diagrams, and bar spectral bandwidth bar charts for analysis.

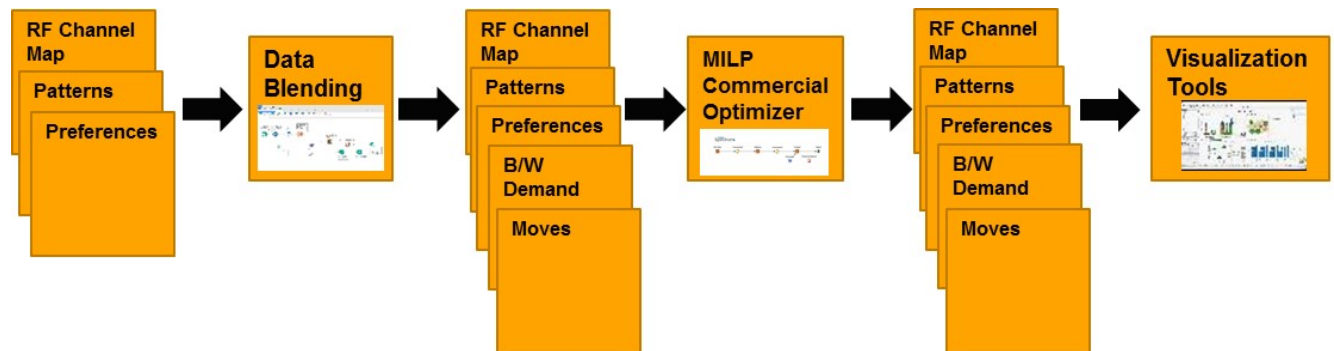


Figure 16 - Data Flow Example

## 4. Business Integration and Automation

In order to effectively integrate the optimization modeling system into the business there are four primary requirements – accurate automated RF channel maps, IP data data traffic demand, video traffic demand, and engineering architectural requirements/direction.

The RF Channel Map (all content assignments in each 6 MHz slot of the spectrum) requires consolidation of allocations to linear video, switched digital video, DOCSIS 3.0 SC-QAM, DOCSIS 3.1 OFDM blocks, and channels reserved for special uses like CLI. Integration of multiple data sources is required. Without automation, the time required to generate and maintain RF Channel Map currency creates a major hurdle that must be overcome to effectively optimize bandwidth.

Accurate prediction of IP data traffic demand is important in order to integrate bandwidth optimization modeling into the business. Peak traffic CMTS utilization can give a good view of current use, but with the movement to more IP video and changing data profiles, it is becoming important to use MAC-level traffic analysis from sources like IPDR to produce accurate forecasts of traffic growth. MAC cable modem mix (e.g. D3.0 vs. D3.1) and individual subscriber segmentation are becoming more and more important in assessing traffic volume requirements on individual service groups.

Accurate prediction of video traffic demand is similarly necessary integrate bandwidth optimization into the business. Compression techniques like MPEG4, increased movement toward switched digital pooling, migration of subscribers to IP video (e.g. cord cutters), and even migration of linear video products to IP video make accurate assessment of bandwidth demand increasingly difficult.

Tying in to the engineering plans for bandwidth architecture changes, new DOCSIS software releases, bandwidth reconfigurations like upstream mid-splits and/or high-splits, etc. is another key need for successful business integration and automation. Synchronization of capacity planning across operations, engineering, video, and data organizations is needed for bandwidth optimization to be effective.

## Conclusions

The three major conclusions that can be drawn regarding use of MILP in optimizing spectral bandwidth are:

1. Data acquisition and data management are the biggest hurdles and most critical success factors.
2. Mixed integer linear programming is a very viable approach to optimizing bandwidth, and
3. There are many benefits to be gained through optimizing spectral bandwidth.

RF Channel Maps containing the contents/allocation of each EIA slot for each plant are necessary for any modeling or optimization. Documentation of preferred standardized configurations and engineering preferences also need to be compiled. Storage in an automated, accessible repository provides high returns in terms of reducing the time required (both modeling and manual manipulation) to analyze bandwidth use.

Mixed integer linear programming optimization for spectral bandwidth optimization requires very little compute time and model scenarios can be turned quickly. This is important when solving for all plants across the corporation as well as iteratively working to find the best optimal configuration associated with an individual channel map. Compute time is generally less than 5 seconds to evaluate as many as  $2^{5000}$  combinations of content assignments to slots. It can also be concluded that MILP can be used without mathematical programming knowledge when designed with a data-driven approach.

Finally, and most importantly, there are many benefits to be had by optimizing spectral bandwidth. With the move to digital, packing program content can open up a lot of bandwidth – avoiding plant upgrades and accommodating unanticipated IP data traffic growth. This is especially important since the value of 6MHz of bandwidth is large and continuing to grow. Benefits in terms of time savings also accrue with optimization modeling. The time required to maintain RF Channel Maps and manipulate configurations to determine the feasibility of upgrades can be reduced. The time required to analyze strategic engineering alternatives can also be greatly reduced.

# Abbreviations

5GL	Fifth Generation Language
CLI	Cumulative Leakage Index
CMTS	Cable Modem Termination System
DOCSIS	Data Over Cable Service Interface Specificationw
EIA	Electronic Industries Association
IPDR	Internet Protocol Data Record
MAC	Media Access Control
MHz	Megahertz
MILP	Mixed Integer Linear Programming
MSO	Multiple-Sytems Operator
OFDM	Orthogonal Frequency Division Multiplexing
RF	Radio Frequency
SC-QAM	Single Carrier Quadrature Amplitude Modulation

# Bibliography & References

*ARRIS Cable Technician Pocket Guide Subscriber Access Networks, Document MX0398 Revision P*; ARRIS Enterprises, Inc., 2014, pp. 2-1,2-37.

CableLabs Data-Over-Cable Service Interface Specifications DOCSIS®3.1, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I18-190422.

<https://specification-search.cablelabs.com/CM-SP-MULPIv3.1>

CableLabs Data-Over-Cable Service Interface Specifications DOCSIS®3.1, CCAP™ Operations Support System Interface Specification, CM-SP-CCAPv3.1-I15-190422.

<https://specification-search.cablelabs.com/CM-SP-MULPIv3.1>

Fredrick S. Hillier and Gerald J. Lieberman, *Introduction to Operations Research*, McGraw-Hill, tenth edition, 2015, ISBN: 0073523453. 1010 Pages.

# **Can Wireless Compete With Wired Access To The Home**

## **A Review of Fixed Wireless Access Technology And Economics**

A Technical Paper prepared for SCTE•ISBE by

**Kashif Shakil**

Customer Solutions Sales Director

Ericsson

6300 Legacy Drive Plano, TX 75024

(972) 679-3737

Kashif.shakil@ericsson.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Reasons for FWA Momentum.....	6
Fixed Wireless Access Network .....	6
Spectrum Options for FWA .....	7
Starting Up and Evolving FWA Network .....	9
3GPP Data Rate Evolution .....	10
FWA with 5G .....	11
Improving 5G mmWave Reach .....	11
CBRS for FWA .....	12
Massive MIMO for FWA.....	12
Summary of Technical Section .....	16
FWA Business Considerations.....	16
FWA economics .....	16
Business Case Example .....	18
FWA business models .....	18
FWA Field Experiences .....	19
Conclusions.....	22
Abbreviations.....	23
Bibliography & References .....	23

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Fixed Wireless Access network.....	7
Figure 2 - LTE throughput ladder .....	10
Figure 3 - Combining mid band with mmWave band .....	12
Figure 4 - Coverage under broadcast beam .....	13
Figure 5 - Beamforming with SU-MIMO.....	14
Figure 6 - Beamforming with MU-MIMO .....	15
Figure 7 - Comparison of 4T4R and 64T64R user throughput and coverage .....	15
Figure 8 - Downlink throughput under different RF conditions.....	19
Figure 9 - Setting for multiple user MIMO testing in ideal radio conditions .....	20
Figure 10 - Setting for MU-MIMO field testing in varying radio conditions .....	21
Figure 11 - FWA with beamforming and MU-MIMO in rural area.....	22

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Broadband access technology comparison .....	5
Table 2 - Spectrum for FWA .....	8
Table 3 - Results of throughput testing with 16 layer MU-MIMO in good radio conditions.....	20
Table 4 - Results of 8 layer MU-MIMO testing in varying radio conditions.....	21

# 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. G.fast promises higher data rates for shorter straight loops. For instance, 600 Mbps for 200m distance. Speed limitation start kicking for longer distances. 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 3.1 FD offers 10 Gbps shared downstream capacity. It is an attractive option but fiber and 5G offerings could disrupt DOCSIS also.
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	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. New operators can disrupt existing DOCSIS and fiber-based offerings. Established operators should study leveraging 3GPP wireless to defend and grow their business.
Proprietary wireless	Wireless is an attractive option for rural broadband. However, operators should consider issues when using proprietary wireless technologies. These technologies are not standards based, they may be limited to one vendor, have smaller ecosystem, an uncertain roadmap & lack of economies of scale. Compared to 3GPP, these technologies may not offer same level of quality.	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 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.

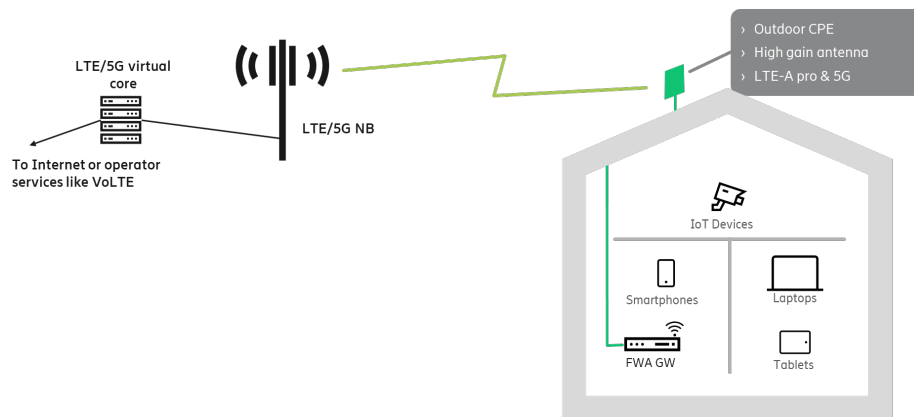
## **Reasons for FWA Momentum**

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 frequencies in 5 GHz band 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.
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.

## **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 non-mobile supported users, 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 may not provide baseband pooling gains due to fixed nature of traffic. A centralized baseband pool may also increase fronthaul transmission costs.

## 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**

<b>Frequency</b>	<b>Benefits</b>	<b>Challenges</b>	<b>Availability</b>
Mid band FDD (PCS, AWS etc.)	Widely available ecosystem Licensed Good propagation	Fully utilized in urban and suburban areas Narrower channel bandwidths	Now
2.5 GHz TDD	~200 MHz of licensed spectrum Best propagation amongst TDD spectrum US ecosystem available today Highest predictability due to licensed spectrum	Cost of acquiring spectrum Majority of spectrum in populated areas is owned	Now
3.5 GHz CBRS	150 MHz of spectrum Global LTE ecosystem Good balance between propagation, power and reuse Interference managed via SAS (Spectrum Access System)	Spectrum demand in urban and suburban areas	Late-2019 as defined by CBRS ecosystem certification timelines
5 GHz	555 MHz of spectrum	Propagation challenges - Maximum of 36 dBm EIRP. Prone to interference due to contention-based access method	Today
5.9 GHz-6.425 GHz	500 MHz of spectrum	Will likely follow unlicensed framework established for 5 GHz	Estimated 2022
24 GHz – 39 GHz	Channels of 100 MHz possible Carrier agg of 400 MHz or more	Significant challenges - propagation	Starting mid-2019
57 GHz – 71 GHz	14 GHz of spectrum Suitable for point-to-point	Significant challenges - propagation and atmospheric absorption	Partially Today; 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
4. Higher peak and average rates can be offered

## 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 on poles 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

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.

## 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.

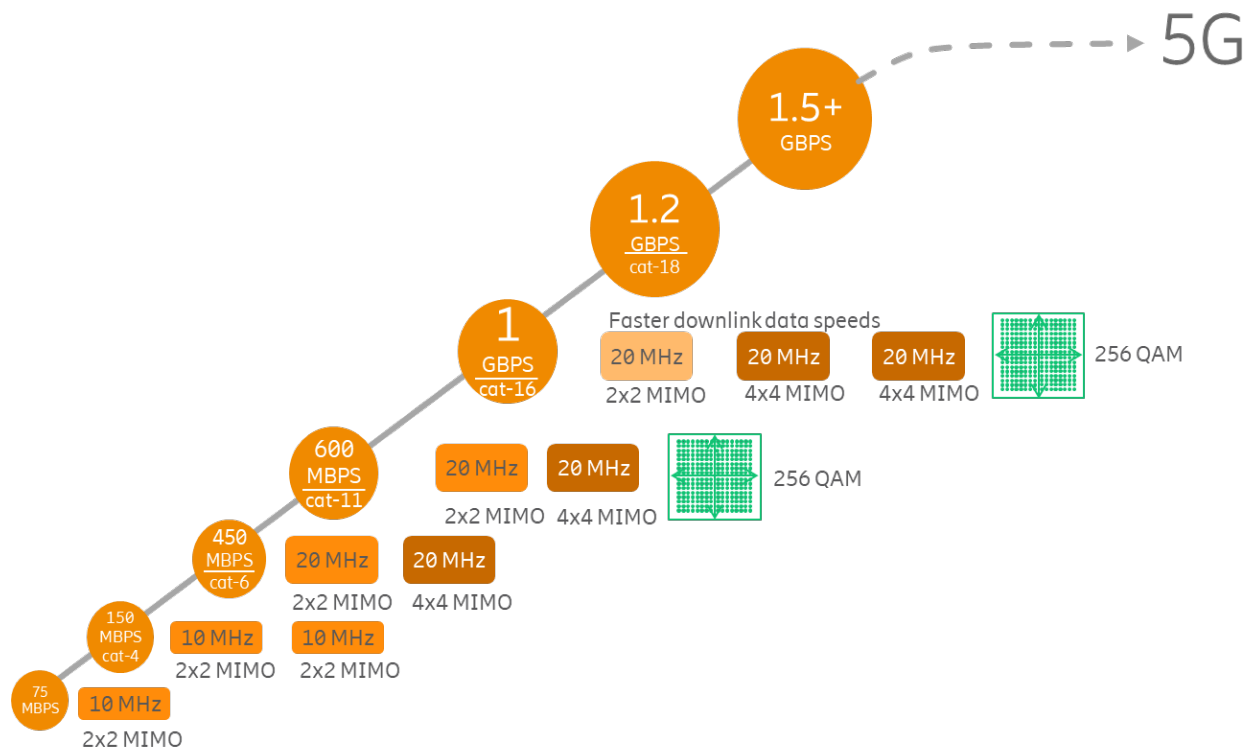
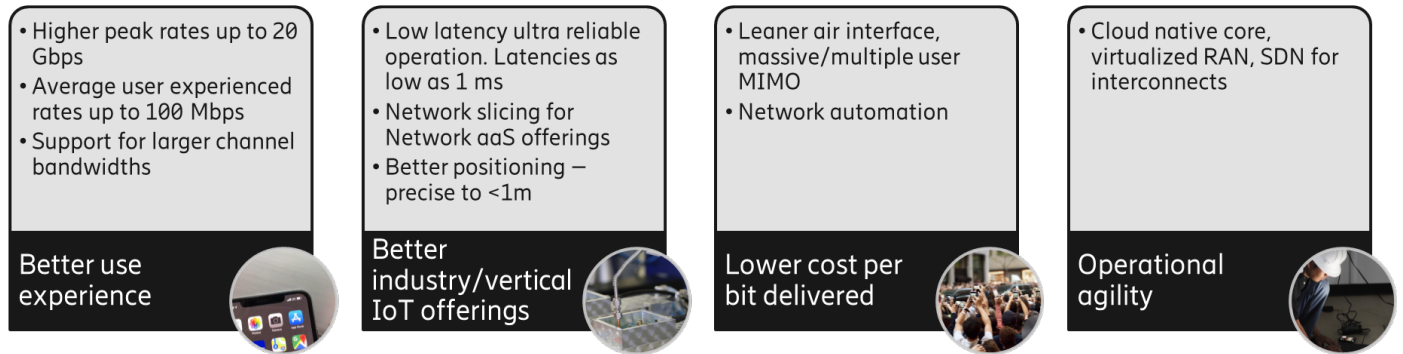


Figure 2 - LTE throughput ladder

# 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

## 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 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 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 3 - Combining mid band with mmWave band**

## 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.

## 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 4 - 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 5 - 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 6 - 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.



*Coverage shown with 3.5Ghz, 3x20 MHz, LoS TDD Frame config 2. Actual results vary based on clutter, terrain and other conditions*

**Figure 7 - Comparison of 4T4R and 64T64R user throughput and coverage**

# 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.

## 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.

## 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 for capacity. 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. However, outdoor CPE may require costly professional installation. Indoor CPEs could be self-installed by the subscriber. On the other hand, self-install prevents operators from guaranteeing performance. Moreover, indoor CPEs may appear as cell edge user to the network, consuming larger share of air interface resources, deteriorating system performance. 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.

Similarly, a higher market share will allow operator to generate more revenue per site. An FWA business case resting on smaller market share or lower subscriber density 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.

# 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 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;3.5 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 market take rates, pricing, and targeted network design metrics</li></ul>	<ul style="list-style-type: none"><li>– Network deployment can begin when 3.5GHz spectrum usage is enabled by FCC's SAS/ESC certification</li><li>– Network deployment using 2.5 GHz licensed spectrum possible today</li></ul>

\* Subscriber growth assumed 10% to 33% in 5 years, 3.5 GHz CBRS spectrum; assumes 60 MHz spectrum

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

## 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.
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.

## 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 8 - 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 9 - Setting for multiple user MIMO testing in ideal radio conditions**

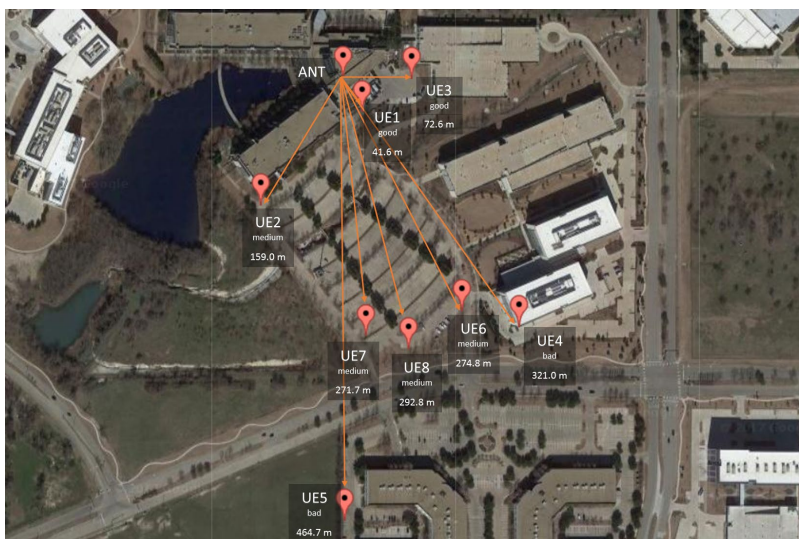
These tests use a single 20 MHz LTE carrier with 64 QAM modulation in mid band TDD spectrum. 16 CPEs are configured to transmit simultaneously. System is set up to use up to 16 layers in both DL UL. Total system throughput of ~740 Mbps was observed in the DL.

**Table 3 - Results of throughput testing with 16 layer MU-MIMO in good radio conditions**

	UE1	UE2	UE3	UE4	UE5	UE6	UE7	UE8
DL Tput(Mbps)	47.16	46.93	46.52	46.72	46.1	46.32	46.31	45.32
UL Tput(Mbps)	8.72	8.91	8.78	8.92	8.89	8.9	8.89	8.92
	UE9	UE10	UE11	UE12	UE13	UE14	UE15	UE16
DL Tput(Mbps)	45.98	45.83	45.64	45.9	47.3	47.13	45.98	46.9
UL Tput(Mbps)	8.9	8.87	8.93	8.86	8.9	8.88	8.92	8.93

In the next field trial, our objective is to compare performance of massive MIMO and multiple user MIMO against traditional LTE in 2T2R, 2x2 single user MIMO. UE placement is more realistic: with 2 CPEs in good radio conditions, 4 in medium radio conditions and 2 in poor coverage. Massive MIMO system is configured only for 64 QAM and 8 layers. 2x2 MIMO system is configured for LTE transmission mode 3 and 256 QAM. We use a single 20 MHz TDD LTE carrier in single band and our spectrum band is in the mid bands.





**Figure 10 - Setting for MU-MIMO field testing in varying radio conditions**

Results show a total system throughput of 178 Mbps for massive/MU-MIMO MIMO system vs. 84 Mbps 2x2 SU MIMO system. This is a 2x improvement in system capacity, as well as considerable improvement in individual UE throughput. With 16 layers, we could expect a 4x improvement in system throughput. Enabling 256 QAM for massive MIMO system would result in higher system capacity gains.

Note: System capacity gains from massive MIMO are tied to layer utilization. Users in over-lapping beams cannot be MU-MIMO users, limiting gains in more practical situations.


**Table 4 - Results of 8 layer MU-MIMO testing in varying radio conditions**

		MU-MIMO DL Throughput	SU MIMO DL Throughput
		<i>Mbps</i>	<i>Mbps</i>
●	UE1	43.83	16.93
●	UE2	11.78	9.25
●	UE3	44.73	15.30
●	UE4	4.69	4.92
●	UE5	17.19	7.03
●	UE6	22.65	9.19
●	UE7	17.80	13.24
●	UE8	15.13	8.76
	SUM	177.78	84.63

● Good RF conditions  
 ● Medium RF conditions  
 ● Poor RF conditions



Finally, we show coverage and range of a massive MIMO system in rural area. We use single 20 MHz carrier in mid band with outdoor CPEs. We also use MU-MIMO. High single user throughput performance is achieved at more than 10 km away from the cell site. Total system capacity of 200 Mbps was achieved for 8 MIMO layers.



**Figure 11 - FWA with beamforming and MU-MIMO in rural area**

## 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>

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)

# **Machine Learning Applications in Cable TV Advertising – Usage and Challenges**

A Technical Paper prepared for SCTE/ISBE by:

Srilal M Weerasinghe PhD  
Principal Engineer  
Charter Communications  
8560 Upland Drive, Englewood, CO 80112-7138  
720-699-5079  
[srilal.weera@charter.com](mailto:srilal.weera@charter.com)

Robbie Mills III  
Media Content Manager  
Charter Communications  
316 East Morehead St, Charlotte, NC 28202  
704-973-7461  
[Robbie.Mills@charter.com](mailto:Robbie.Mills@charter.com)

Vipul Patel, Charter Communications

Basil Badawiyeh, Charter Communications

Mike Terada, Charter Communications

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
1. TV Advertising – Quality Control at Ad Ingest .....	3
2. Machine Learning in Carrier-Class Video Applications – Challenges.....	4
2.1. Technical Challenges .....	4
2.2. Machine Learning Models for Image and Video Classification .....	5
2.3. Performance considerations .....	5
2.4. Limitations of Current Machine Learning Tools.....	5
3. Lab Evaluation .....	6
3.1. Image Analysis .....	6
3.2. Video Analysis.....	6
3.3. Types of Errors.....	6
3.3.1. False Positives.....	6
3.3.2. False Negatives .....	7
3.3.3. Machine Learning Tool Performance.....	8
4. Proposed Solution.....	8
4.1. Steps Summary.....	8
4.2. Solution Details .....	11
Conclusions.....	11
Abbreviations.....	12
Bibliography & References .....	12

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – False Positive - Fireworks .....	7
Figure 2 – False Negative - Alcoholic Beverage .....	7
Figure 3 – JSON file of audio script of the parsed ad.....	8
Figure 4 – Summary of Steps of the Proposed Solution .....	9
Figure 5 – Proposed Solution for Mitigating ML Results .....	10

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Machine Learning Identification of Activities for Ads* .....	6
Table 2 – Machine Learning Detection and Error Mitigation.....	7

# Introduction

Use of machine learning (ML) for image and video analyses would often include face recognition, personalization and recommendations. An emerging trend is the application of AI technology for TV advertising. In this paper, we present the unique challenges in applying machine learning to carrier-class video advertising. We focus the discussion on a specific use case that is common to all ad supported TV services.

The selected use case is Ad Ingest Quality Control (QC). In the United States, TV commercials are subjected to various rules and regulations. Ads containing specific content (e.g. Alcohol, firearms) are barred from airing during certain TV programs. Identifying these categories may pose a challenge, as off-the-shelf machine learning products are more oriented towards facial recognition. That is to be expected perhaps, as the video ML products were primarily intended for surveillance and sports applications. However, our research indicates that by judiciously combing metadata from multiple data streams, machine learning analysis results can be improved.

The intent of the paper is to outline the results and recommendations of a proof-of-concept study that will be helpful to the carrier-class video services community.

## Content

### 1. TV Advertising – Quality Control at Ad Ingest

Multi-channel video programming distribution/ distributor (MVPD) is a highly regulated industry in the US. The term covers not only traditional cable companies, but any entity that provides TV service to consumers via fiber, coax, satellite, DSL and wireless. With the advent of internet-based TV service (also known as OTT), the moniker is modified as V-MVPD (virtual MVPD). In all cases, the content distributors could be responsible for the displayed video content, including advertisements [1]. This places the onus on the content distributor (also known as service provider/network operator), to prevent the ‘non-compliant’ content from reaching the TV audience.

In the context of the present discussion, there is a distinction between movie content and ads. While movies/episodes are originated from mainstream studios (and are properly vetted), the TV ads could originate from a multitude of sources. Therefore it is necessary to identify any non-compliant ads prior to airing at the Ad Ingest Quality Control (QC). Today, this is done manually by trained individuals. They examine tens of thousands of ads a month and quarantine the failed ones. The challenge is to automate that process with an AI/ML engine embedded into the workflow.

First we examine the basis for non-compliance of ads. When a TV commercial is deemed non-compliant, the restriction usually stems from one of the three categories below.

#### a) Regulatory Compliance

The Regulatory constraints are primarily stipulated by FCC [1] but could also be under the purview of FTC and FDA [2] [3]. Listed below are some examples of regulatory requirements overseen by federal agencies. See the references cited above for full requirements.

- Ads related to alcohol, tobacco, firearms, gambling etc. must meet federal guidelines.
- A political ad is required to display a statement from the sponsor for at least 4 seconds.

- Truth-in-advertising – An ad may be deemed deceptive for misleading/missing information.
- Ads promoting certain lotteries, cigarettes or smokeless tobacco products are not allowed.
- Ads must comply with loudness mitigation requirements of CALM Act.

#### b) Contractual Compliance

Contractual constraints are imposed by content providers such as ESPN. An example would be the restriction on alcohol ads during ESPN Little League World Series program. For a complete list of applicable restrictions, see reference [4].

#### c) Business/Operational Compliance

These are generally operational guidelines and best practices established by the enterprise. Being sensitive to audience needs as well as delivering quality content could enhance a company credibility. One example is ‘frequency capping’ or limiting the display of the same ad multiple times.

## **2. Machine Learning in Carrier-Class Video Applications – Challenges**

Identifying the above categories programmatically poses a challenge to ML tools, as off-the-shelf products are more oriented towards facial recognition. A familiar ML application is creating a ‘bounding box’ around a face and tracking it through a video-clip. Such applications are useful in sports and surveillance, however they are not directly applicable to MVPD market. The latter requires comprehensive ML analyses of multiple streams (video, audio and textual metadata).

In common usage, Machine Learning video products do a multi-pass analysis (each pass to identify faces, common objects, celebrities etc.). The results are presented as content descriptor metadata (labels). An accompanying ‘confidence level’ indicates the accuracy of prediction. Per our lab testing, off-the-shelf ML tools didn’t meet our needs right out of the box. It may be because the video content/Ads detection is still a nascent technology. Adapting such products for carrier-class video applications requires a certain amount of post-processing. Else, the results could be tainted with false positives or the tool may fail to identify content adequately (false negatives).

### **2.1. Technical Challenges**

To train a neural network, a good selection of examples and counter-examples is needed. Else the machine learning model would be susceptible to ‘overfitting’. That is, the model will fit the existing data well, but would fail when it encounters a new instance of the target data. While this is not an issue with common objects (e.g. cars) due to the abundance of examples, it is a challenge for objects with ambiguous signatures (such as fireworks or alcohol). Distinguishing ‘fireworks’ from similar signatures (‘bright lights in a dark background’), is not an easy task. Similarly, an image classifier may find it hard to differentiate ‘beer’ from a similarly colored liquid in a bottle (e.g. olive oil).

The need for proper counter-examples becomes more acute as we move from image analysis to video activity identification. This is discussed in detail in the ‘Issues Noted in Our Testing’ section below.

Next we present a short overview of applicable deep learning algorithms.

## **2.2. Machine Learning Models for Image and Video Classification**

General multi-perceptron based neural networks (ANN) are not able to meet carrier-class video classification requirements. Training time and accuracy would be hard to achieve. Convolutional neural networks (CNN) is the Deep learning based technology used for image classification. Most products use ‘transfer learning’ model; first training the model on a large public dataset such as ImageNet or Inception and then fine tuning it to meet the specific requirements. While image analysis has only spatial dependence, video analysis involves the temporal component.

For time series analysis, recurrent neural networks (RNN) deep learning model is the standard technique, due to its ability to store events happened in the past. However, it is well known that RNN, with many hidden layers, suffers from the vanishing gradient problem. This issue also manifests as the exploding gradient problem. (A simpler interpretation is that Tangent of the angles being very close to 0 or 90 degrees, respectively). The root cause is the exceedingly small derivatives of the ‘loss function’ (or error), during back propagation. A solution is to disregard certain intermediate steps to avoid extreme values of the gradients. A popular model for handling such sequence data is the Long Short Term Memory (LSTM) algorithm. LSTM discards certain data (via the ‘forget gate’) to reset the cell state thus keep the values getting extreme.

In the field of deep learning, new algorithms are routinely being developed (Fast R-CNN, Faster R-CNN etc.). These are mainly for improving the speed of analysis, as updating millions of parameters (weights and biases) associated with hidden states takes a lot of time.

## **2.3. Performance considerations**

In our testing, the processing time as measured was not close to real-time. One reason could be the ML engines operate in multi-pass mode. This is necessary because at Ad-Ingest quality control, the ML engine works as a gate-keeper. On the other hand, if the intent is to find a single signature (e.g. either guns or alcohol), a single pass would be sufficient.

We have tested machine learning models in appliance mode as well as in the cloud. The cloud-based implementation is preferred if the data also resides on the same cloud. The appliances would be GPU-based (as opposed to CPU), due to the large number of cores which facilitates parallel computing. We tested with NVidia GTX and also plan to benchmark with NVidia DGX (with thousand TFLOPs of computing speed),

## **2.4. Limitations of Current Machine Learning Tools**

To improve the detection accuracy, Machine Learning tools tend to use increasingly sophisticated algorithms. However, the algorithmic approach alone did not seem to produce expected results. Obtaining optimal results within a reasonable time is a challenge. Searching each video frame for a multitude of categories (alcohol, gambling, drugs, violence, trademarks, copyrighted content, explicit content, political content etc.) is time consuming. It could also be irrelevant (i.e. searching for all manners of firearms or medications would be wasteful, in the case of a beer ad).

To improve the results, we propose adding a software engine to the workflow to perform additional analyses.

### 3. Lab Evaluation

Our findings are presented below in a vendor agnostic manner.

#### 3.1. Image Analysis

Content descriptors (Labels) need to be sufficiently descriptive for effective contextual analysis, i.e. instead of generic labels such as ‘person/human’, the ML tool needs to identify whether a person is young/old, male/female, mood etc.

#### 3.2. Video Analysis

Activity identification is a challenge for current ML tools. This is a burgeoning field of research at premier AI/ML research institutions [5]. For the MVPD space, ‘activity identification’ would open up new applications. E.g. identifying a car chase from a video (as opposed to cars in a still image) would offer new ad opportunities. Table-1 below depicts sample activities that are relevant to contextual advertising.

**Table 1 – Machine Learning Identification of Activities for Ads\***

Dominant Activity	Suggested Ad Types
Cooking	Kitchen Appliances & Utensils, Cooking Classes
Car chase	New Cars, Auto Repairs, Auto Insurance
Shopping	Retail Stores
Eating	Food, Restaurants
Dancing	Clothing , Personal Care, Alcohol
Drinking	Alcohol
Social gathering	Clothing, Jewelry
Kids playing	Toys, Food and Drinks, Medicines, Clothing
Sports activities	Sports Related Products
Anxiety, Arguing	Pain Medications, Lawyers

(\*examples only)

#### 3.3. Types of Errors

##### 3.3.1. False Positives

In this example, the tool misidentifies the bright light in the dark background as ‘fireworks’ (with a high confidence level).





**Figure 1 – False Positive - Fireworks**

**Table 2 – Machine Learning Detection and Error Mitigation**


Detected Category	Initial Confidence Level	New Confidence Level
Fireworks has been detected from 00:00:02 to 00:00:03	90%	< 30%

In the “Proposed Solution” section below we present a methodology to mitigate this issue.

### **3.3.2. False Negatives**

In this example, the tool fails to identify the alcoholic beverages in the image analysis. However, the term ‘Cocktails’ is noted in the audio transcript as depicted in the JSON file (Figure 3).

In the “Proposed Solution” section below we present a methodology to mitigate the false negative impact.



**Figure 2 – False Negative - Alcoholic Beverage**

```

{
  "id": 4,
  "text": "You can enjoy our hot tub cocktails and R Florida.",
  "confidence": 0.9069,
  "language": "en-US",
  "instances": [
    {
      "Start": "0:00:16.74",
      "End": "0:00:19.82",
    }
  ]
}

```

**Figure 3 – JSON file of audio script of the parsed ad**

The JSON file in Figure 3 indicates the word ‘cocktails’ as parsed from the audio transcript. This data is available even though the image analysis failed to recognize alcoholic beverage in the video.

### **3.3.3. Machine Learning Tool Performance**

Another issue with some ML products is the excessive time taken for video analysis. In our studies, a 30-second ad would take 2-3 minutes for a multi-pass analysis. This can be improved substantially with faster GPU processors.

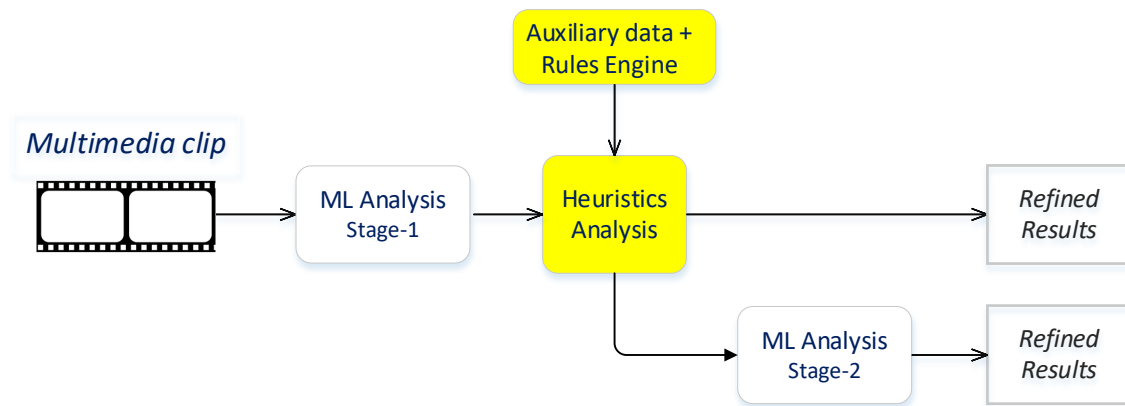
## **4. Proposed Solution**

Current Machine learning products treat metadata of each stream separately; e.g. video/image analysis is separate from audio or text analysis, albeit each may use neural networks based classification algorithms. We believe that interrelating the video, audio and text data could enhance the accuracy of predictions. For example, a gambling ad for a casino may have telltale signs on video-audio-text streams. These accompanying signatures (supplementary/auxiliary data on multiple streams) are utilized by the software decision module introduced below.

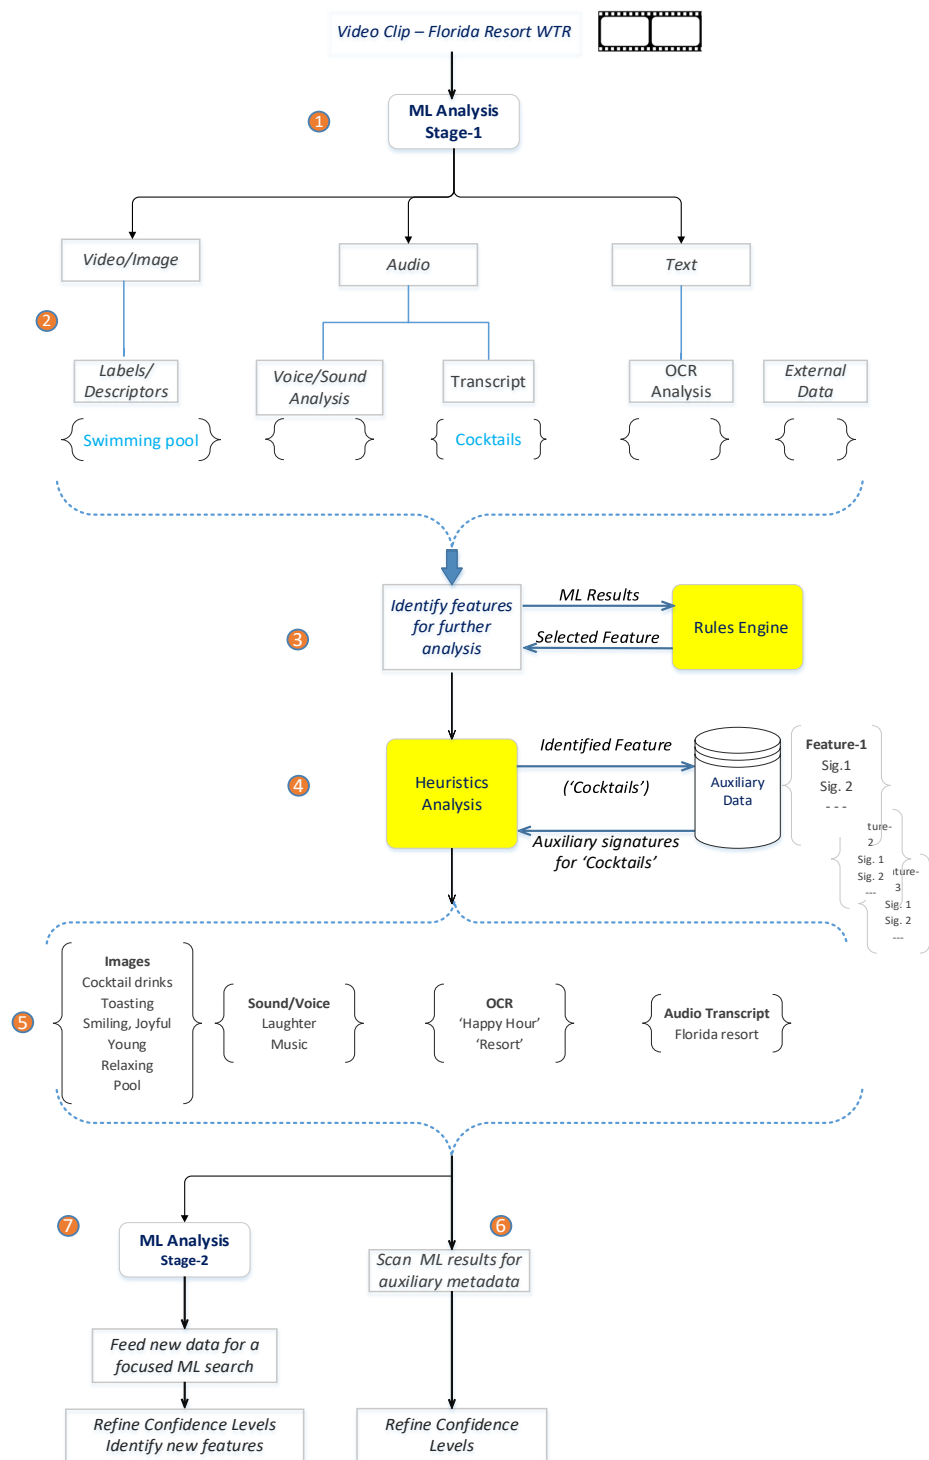
The proposed solution consists of multiple stages. A modified workflow is introduced with an embedded decision module to accommodate heuristic analysis. ‘**Heuristic**’ in the present context would mean an educated guess based on supplementary data. It is not a rigorous deterministic algorithm, but yields results in a reasonable time. Note that a signature-based deterministic approach is not guaranteed to work in the selected use cases. For example, some beer ads do not use the term ‘beer’ in the audio stream. In such cases, auxiliary signs in other streams (images of joy, relax, young people, bottles/cans, OCR data) could be strong clues.

### **4.1. Steps Summary**

1. First pass is a general ML analysis to derive ‘content descriptors’.
2. Next, a heuristic analysis is performed using auxiliary data to assess the initial results.
3. The ‘confidence levels’ are reassessed and revised based on rules set.
4. Option to perform a more refined 2<sup>nd</sup> stage ML analysis for content classification.



**Figure 4 – Summary of Steps of the Proposed Solution**



**Figure 5 – Proposed Solution for Mitigating ML Results**

## 4.2. Solution Details

(Numbering below corresponds to Figure 5)

1. Ad creatives/video content are fed into the ML engine (these could be 30-second Ads, long-form ads, TV episodes, movies or other multimedia content)
2. ML Engine conducts machine learning based analysis. The results are the identified content descriptors and confidence levels.
3. Identify features for further analysis. The criteria is based on rules built previously, such as a list of keywords.
4. Heuristic Analysis – Retrieve ‘auxiliary data’ for the feature identified above. These signify plausible signatures that may appear in other streams for a given feature.
5. This step depicts sample auxiliary signatures for the term ‘Cocktails’. These are pre-populated in the database.
6. In this step, previous ML results (from the first pass) are fed into the software module programmatically. It searches for the presence of auxiliary data in other streams. Based on the analysis, ‘confidence level’ is adjusted. (i.e. If the metadata terms ‘drinks’, ‘toasting’ appear in the video analysis, the confidence level for ‘Cocktails’ is increased. Conversely, if there are no supporting auxiliary data, the confidence level is lowered.
7. Optionally, a second stage ML analysis is supported for a more refined search. Using Auxiliary data for the classification algorithm would enable a focused and accurate search.

Using the above process, the false positive/negative impacts are mitigated. Column 3 of Table-2 shows the results of applying heuristic analysis. In the case of fireworks, if there are no auxiliary signs on other streams (such as ‘noise’), then the Confidence Level is lowered by a factor. The Rules Engine contains the pre-set value of the multiplier (e.g. 0.75)

Note that in the case of real ‘fireworks’, an image frame taken a second later would have the lights diminished. That heuristic signature could be used to differentiate fireworks from other lights and reduce false positives.

In the same fashion, if the auxiliary signatures are present in other streams (as in Figure 5), then the original confidence level is multiplied by a factor (e.g. 1.25) which would increase the final confidence level value.

## Conclusions

Based on our testing, visual analysis alone is not sufficient to make meaningful recommendations for carrier-class video (unlike surveillance or sports use cases). A multi-stream analysis of Video, Audio and Text (OCR) streams would provide a better contextual interpretation.

Machine learning applications to carrier-class video services is still a nascent field. We outlined some of the unique challenges. A multi-stream heuristic method was proposed to complement the current algorithmic approach. The MVPD space is a fertile field for AI/ML applications, and much work still needs to be done.

# Abbreviations

AI/ML	Artificial Intelligence/Machine Learning
DL	Deep Learning
ANN	Artificial Neural Network
CNN	Convolutional Neural Network
RNN	Recurrent Neural Network
LSTM	Long Short Term Memory
R-CNN	Region based Convolutional Neural Network
SSD	Single Shot Detector
TFLOP	Trillion floating-point operations per second
JSON	Java Script Object Notation

## Bibliography & References

- [1] FCC Guidelines for Ads - <https://www.fcc.gov/consumers/guides/complaints-about-broadcast-advertising>
- [2] FTC Guidelines for Ads - <https://www.ftc.gov/tips-advice/business-center/guidance/ftcs-endorsement-guides-what-people-are-asking>
- [3] FDA Guidelines for Ads - <https://www.fda.gov/media/82590/download>
- [4] ESPN Advertising Guidelines - [http://www.espn.com/adspecs/guidelines/en/ESPN\\_AdStandardsGuidelines.pdf](http://www.espn.com/adspecs/guidelines/en/ESPN_AdStandardsGuidelines.pdf)
- [5] MIT AI Lab research - <http://moments.csail.mit.edu/explore.html>

# **An Analysis of How to Deploy Low Power WAN IoT Using HFC and Fiber Network Infrastructure**

A Technical Paper prepared for SCTE•ISBE by

**Patricio Sebastian Latini**  
Principal Technologist  
CASA Systems  
100 Old River Rd. – Andover, MA  
+1 (305) 504-9250  
patricio.latini@casa-systems.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
Content .....	4
1. IoT background .....	4
2. Low Power Area Networks (LPWAN).....	6
3. LoRaWAN.....	7
3.1. Link Budget .....	10
3.2. Effects of Aloha in LoRaWAN .....	13
3.3. Signal Propagation Models .....	15
3.4. Frequency Plan .....	16
3.5. Scenario Modeling.....	16
4. HFC Networks .....	20
4.1. HFC Node + 0 Networks .....	20
4.2. HFC Node+1/2 Networks .....	20
5. Network Density and Throughput .....	21
6. Powering and Backhaul Connectivity.....	23
6.1. Powering .....	23
6.2. Backhaul Connectivity .....	23
7. LTE Based IOT .....	24
Conclusion .....	25
Abbreviations.....	26
Bibliography & References .....	27

## List of Figures

Title	Page Number
Figure 1 - The dimensions of communications in IoT. Retrieved from (ITU, 2012) .....	5
Figure 2 - Creation of environments and intelligent spaces. Adapted from (Smith, 2012) .....	6
Figure 3 - Projection of connected Devices. Retrieved from (Ericsson, 2016) .....	7
Figure 4 - Spreading of Symbols in LoRaWAN.....	8
Figure 5 - LoRaWAN package structure. Retrieved from (Semtech, 2013) .....	10
Figure 6 - Channel Throughput vs Channel Load for LoRaWAN.....	14
Figure 7 - Pathloss vs Distance for Dense Urban scenario.....	17
Figure 8 - Pathloss vs Distance for Sub-Urban scenario.....	19
Figure 9 - Node+0 sample node distribution.....	20
Figure 10 - Node+1/2 sample node distribution.....	21
Figure 11 - LoRAWAN Architecture. Retrieved from TheThingsNetwork.....	24

## List of Tables

Title	Page Number
Table 1 - Code Rates for LoRaWAN. Adapted from (Noreen, Bounceur, & Clavier, 2017) .....	8
Table 2 - Spreading Factor vs Chip Length .....	12



Table 3 - Spreading Factor vs. Link Sensitivity and Budget .....	12
Table 4 - Channel Bitrate vs Channel Bandwidth, Spreading Factor and Coding Rate.....	13
Table 5. Channel Bitrate vs Spreading Factor and Coding Rate for Aloha MAC.....	14
Table 6 - Total gateway Bitrate vs Spreading Factor and Coding Rate for Aloha MAC in an 8-channel gateway. ....	15
Table 7 - Spectrum bands for LoRaWAN in the USA.....	16
Table 8 - Public Applications Communication Parameters.....	22
Table 9 - Residential Applications Communication Parameters .....	22
Table 10 - Device density for cities .....	22
Table 11 - Devices and Bitrate per gateway .....	23
Table 12 - LTE IOT Protocols .....	24

# Introduction

In the last few years the telecommunications world has been focused on developing and deploying specific, viable IoT infrastructure as well as IoT-based business and use cases. However, most of the deployments are done in an OTT (Over-The-Top) manner over existing internet connections. This is in large part due to the network agnostic nature of many consumer IoT applications. IoT applications directly connect to their IoT providers through open internet connections, giving multiple subscriber operators (MSOs) the traffic load, but cutting them out of any control of the quality of the service and also its revenues.

This paper analyzes the main alternatives of Low Power WAN (LPWAN) IoT native protocols such as LoRaWAN running over unlicensed spectrum compared to mobile based protocols such as LTE-M and NB-IOT using a typical MSO infrastructure as their support.

A detailed analysis of RF footprint for both alternatives is presented by using the existing infrastructure of the MSOs to support physical mounting, powering and network backhauling by either using in-home or out-of-home alternatives. This analysis focuses on the positioning of LoRaWAN access points in key positions of the hybrid fiber coax (HFC)/fiber network in order to effectively serve remote sensors in different types of scenarios of device densities, such as dense urban and suburban.

A backhauling analysis is presented showing the key elements to properly support the most important key performance indicators (KPIs) of IoT networks such as packet loss, latency and bandwidth over DOCSIS® transport.

Lastly a security and network transport layer model is presented in order to properly support thousands of remote access points/small cells without the necessity of dedicated managed transport networks. The use of internet security (IPSEC) protocol is analyzed together with the requirements for the tunnel termination requirements for supporting the mentioned topology.

The resulting conclusions will allow the cable operators to better understand how the different LP-WAN protocols behave at the RF level in certain configurations that are well aligned and resource efficient with current HFC-fiber infrastructure deployments in MSOs. This understanding may help MSOs better plan for LP-WAN network deployments.

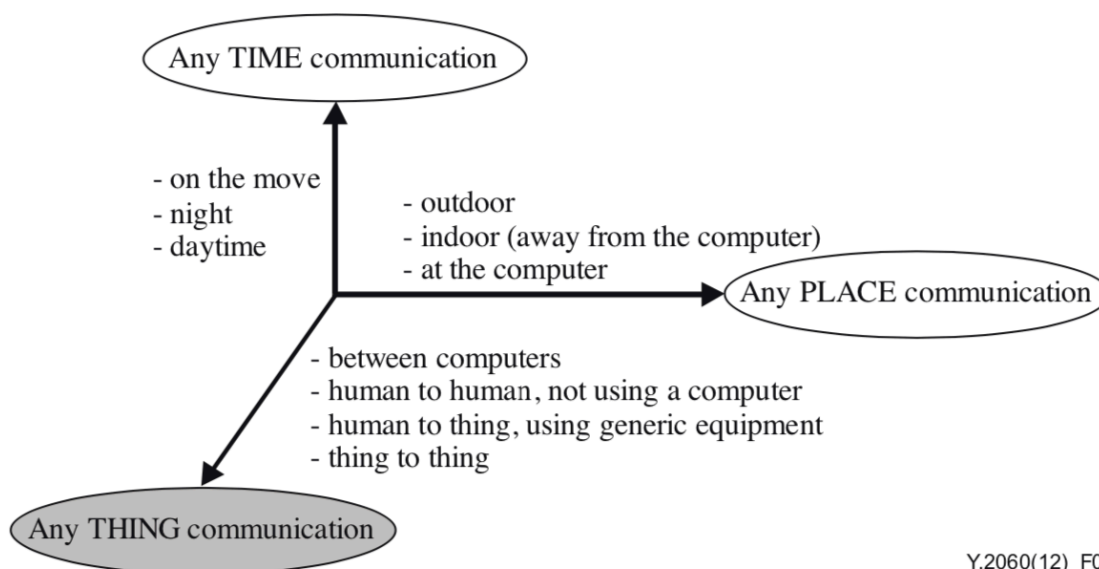
## Content

### 1. IoT background

The term "Internet of Things" (Internet of Things or IoT) was first used by Kevin Ashton in 1999 in the context of a presentation on how to improve the efficiency of a company's supply chain systems of provision of goods through the use of radiofrequency markers (RFID) instead of bar codes and how to achieve an automated data collection through the use of a network and mainly the elimination of the human factor in said data collection (Ashton, 2009 ). This is particularly important given that the term "Internet of Things" brings with it the concept that "things start to use the network in such a way that people do not need to" as Neil Gershenfeld (1999) mentioned

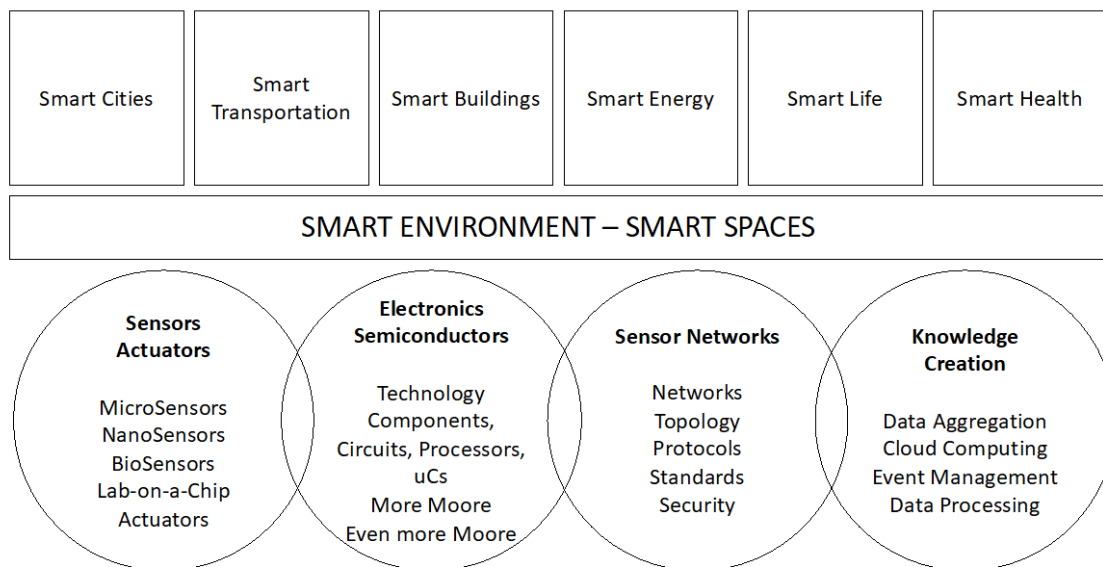
and relates the automation of the communication of things, now tied to the growth of the internet as a network of global interconnection (Gershenfeld, 1999).

In 2012, one of the broadest definitions of the Internet of Things was made by the International Telecommunications Union (ITU). Which defines it as "the global infrastructure for the information society that facilitates the provision of advanced services through the interconnection of objects (physical and virtual) thanks to the interoperability of present and future information and communication technologies" (ITU, 2012). Likewise, each of these things or objects is an object of the physical world (physical objects) or the world of information (virtual objects) that can be identified and integrated into communication networks. A very important factor also defined by the ITU is the ubiquity of communication. As seen in Figure 1, the Internet of Things adds a new dimension, that of communication between objects, and not just between computers or people (ITU, 2012).



**Figure 1 - The dimensions of communications in IoT. Retrieved from (ITU, 2012)**

A goal of vital importance for the development of the internet of things was the creation of intelligent environments and spaces with "own life" things (see eg smart transport, products, cities, rural areas, health) (Smith, 2012). Said intelligent spaces or environments require a set of technologies functioning in a coordinated manner, which ranges from: the sensors, going through local processing, the interconnection of the sensors, and reaching the use of mass data processing in the cloud, as shown in Figure 2.



**Figure 2 - Creation of environments and intelligent spaces. Adapted from (Smith, 2012)**


A list of different Internet of Things applications based on different sensors and in different market segments is shown in Table 1. The list includes 8 different verticals in which the most striking scenarios are mentioned (Libelium, 2018).

To exchange data between applications, devices and things, there are several communication standards such as Bluetooth, WiFi and several mobile standards based on GSM (2G / 3G / 4G). In general these standards have achieved, with technological advances, significant increases in transmission of data rates and therefore the ability to transmit images or videos in real time. Most of the IoT use cases mentioned above do not require very high data transmission rates, but they do require very low energy consumption, since the sensors in general are located in remote areas with access challenges and very limited space for the placement of batteries.

## 2. Low Power Area Networks (LPWAN)

As a result, the goal of Low Power Wide Area Networks (LPWAN) has become a central issue in the IoT. LPWAN is a broad term where there is a variety of technologies used to connect sensors and controllers to the Internet without the use of WiFi or traditional cellular networks. Two primary standards have emerged for LPWAN networks: those based on cell phones, for example, NB-IoT or LTE-M; and those developed natively for IoT use cases such as LoRaWAN and SigFox. The predominant design considerations are low power consumption (up to more than 10 years of autonomy), strong penetration, the connection of a large number of sensors and very low bandwidth devices (Hassan, 2018).

From 2016 through at least 2022, IoT devices are expected to increase in number at a compound annual rate of 21 percent, driven by the new use cases mentioned above (Ericsson, 2016), therefore the need for spectrum will be marked to serve nearly 29 billion devices of which at least 2100 million will be of the LPWAN type, Figure 3.



**Figure 3 - Projection of connected Devices. Retrieved from (Ericsson, 2016)**

### 3. LoRaWAN

Compared with other modulation techniques, the spread spectrum technique used in LoRaWAN ensures a greater range of links, as well as better immunity to interference. LoRaWAN uses a 125 kHz to transmit the signal but also allows the use of scalable bandwidth between 125 kHz, 250 kHz or 500 kHz (Reynders, Meert, & Pollin, 2016). The use of a wider band makes LoRaWAN resistant to noise, Doppler effects, long-term variations of oscillators and fading. However, the use of a narrowband signal in a much wider band makes the spectrum less efficiently used, unless a generation of perfectly orthogonal signals is achieved between the different transmitters (Noreen, Bounceur, & Clavier, 2017).

The transmitter generates signals by varying its frequency over time and keeping the phase constant between adjacent symbols. The signal transmitted is a signal similar to noise that is resistant to multipath fading and Doppler, and is robust against interference. The receiver can decode even a very attenuated signal of 20 dB below the noise level (Semtech, 2013).

The error correction technique used in LoRaWAN to further increase the sensitivity of the receiver is the FEC (Forward Error Correction) type, particularly through the use of a Hamming code of adjustable length (Europe Patent No. 13154071.8, 2013). The code rate (CR) defines the amount of FEC and LoRaWAN offers CR values between 1 and 4. LoRaWAN uses code rates, Coding Rate =  $4 / (4 + CR)$  or 4/5, 4/6, 4/7 and 4/8. If the code rate is denoted as  $k = n$ , where  $k$  represents useful information and the coder generates  $n$  number of output bits, then  $n - k$  will be the redundant bits. Redundancy allows the receiver to detect and correct errors in the message at the cost of decreasing the effective data rate as evidenced in Table 1.

**Table 1 - Code Rates for LoRaWAN. Adapted from (Noreen, Bounceur, & Clavier, 2017)**

CR	1	2	3	4
Coding Rate	4/5	4/6	4/7	4/8
Efficiency	0.8	0.666	0.571	0.5

In LoRaWAN you can choose a variable spreading factor (SF) as a function of the received signal-to-noise ratio (SNR). This spreading factor adapts the length of the symbol and at the same time also specifies the number of bits per symbol. Therefore, changing the spreading factor gives a variable bit rate between 366 bps for the highest propagation factor (SF = 12) and 48 kbps for the lowest propagation factor (SF = 6) as shown in Equation 1

$$R_b = \frac{BW}{2^{SF}} * SF \text{ [bits/seg]}$$

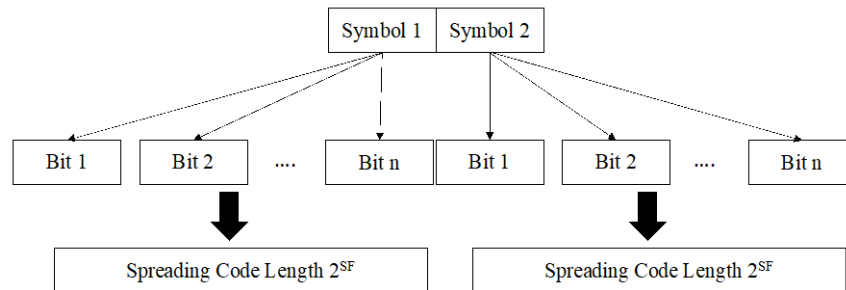
Equation 1. LoRaWAN bit rate

With Coding

$$R_b = \frac{\frac{4}{\frac{4+CR}{2^{SF}}}}{BW} * SF \text{ [bits/seg]}$$

Equation 2. LoRaWAN bit rate with coding

Although the choice of a higher spreading factor increases the bit rate, higher spreading factor reduces the maximum range of the transmission and the same occurs in the opposite direction. Each symbol is scattered with scatter code of  $2^{SF}$  chips in length. In the transmitter, the scatter code is subdivided into codes of length  $2^{SF} / SF$ . Then, each bit of the symbol is scattered using the sub code. Therefore,  $2^{SF}$  chips are needed to propagate a symbol. This same spreading code is also known in the receiver. The replacement of a symbol by multiple information chips means that the spreading factor has a direct influence on the effective data rate (Noreen, Bounceur, & Clavier, 2017), as shown in Figure 4.



**Figure 4 - Spreading of Symbols in LoRaWAN**

The choice of a longer length of spread code improves the transmission distance, but at the cost of a lower bit rate given by the increase in time in the air. This follows from the application of the Shannon-Hartley theorem that establishes the maximum possible information rate for a channel with noise at a given bandwidth as seen in Equation 3 (Proakis, 2000)

$$C = B * \log_2 \left( 1 + \frac{S}{N} \right) \text{ [bits/seg]}$$

Equation 3. Shannon's theorem

Where:

C = Channel capacity

B = Channel bandwidth

S = Average signal power in the receiver

N = Average noise power in the receiver

S / N = Signal to noise ratio in the receiver

If the logarithmic ratio of base 2 to natural base is converted and also assumed that for a spread spectrum application the small signal to noise ratio and the signal power will be much lower than the noise, then  $S / N \ll 1$  and Equation 2 can be rewritten as Equation 4.

$$\frac{C}{B} = 1.43 * \frac{S}{N}$$

Equation 4. Simplified Shannon's Theorem for Spread Spectrum

Therefore, it follows that in order to transmit error-free information to a given signal-to-noise ratio, it is only necessary to increase the bandwidth of the channel to transmit more information. (Semtech Corporation, 2015)

Now, given that the floor noise also called Johnson-Nyquist thermal noise, is determined by the channel bandwidth according to Equation 5 (Sam Lee & Miller, 1998)

$$N_{Floor} = 10 * \log_{10}(k_B * T * B * 1000) \text{ [dBm]}$$

Equation 5. Johnson-Nyquist noise

Where:

$k_B$  = Boltzman constant ( $1.38 * 10^{-13} \text{ m}^2\text{kg/s}^2 \text{ K}$ )

T = Temperature [K]

B = Channel Bandwidth [Hz]

1000 = Conversion of watt to milli watt

Assuming  $T = 293 \text{ K}$  and simplifying  $N_{Floor} = -174 + 10 * \log_{10}(B) \text{ dBm}$ , Equation 6 is obtained, where it is shown that the noise floor depends on the bandwidth of the channel used.

$$N_{Floor} = -174 + 10 * \log_{10}(B) \text{ [dBm]}$$

Equation 6. Johnson-Nyquist noise at room temperature

If you want to obtain the minimum detectable signal or sensitivity in a receiver, you should consider this noise figure, and the minimum signal to noise ratio required for the modulation to be used as shown in Equation 7.

$$S = -174 + 10 * \log_{10}(B) + NF + SNR [dBm]$$

Equation 7. Sensitivity of a radio receiver

### 3.1. Link Budget

If you wish to establish the link budget, this is determined by Equation 7

$$\text{Link Budget} = \text{Minimum Detectable Signal} - \text{Maximum Transmit Power [db]}$$

Equation 7. Link calculation

The maximum range is given by the link budget considering the signal attenuation in free space (FSPL) according to Equation 8. (Sam Lee & Miller, 1998).

$$FSPL = 20 * \log_{10}(d) + 20 * \log_{10}(f) + 20 * \log_{10}\left(\frac{4\pi}{c}\right) [dB]$$

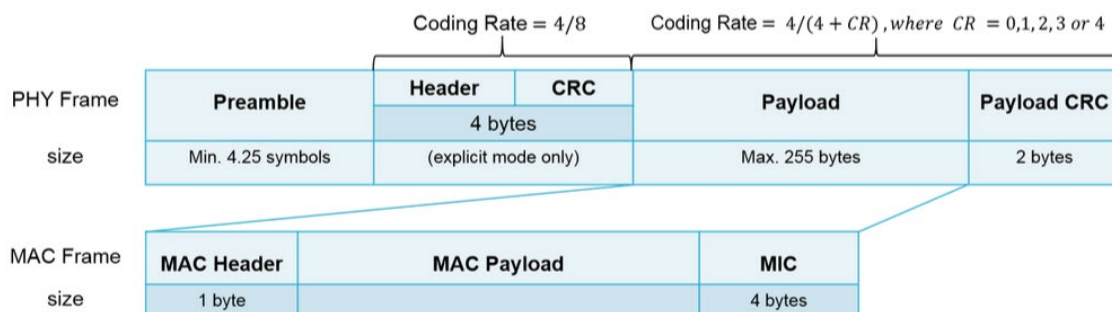
Equation 8. Attenuation of free space

Considering f=900 MHz

$$FSPL = 20 * \log_{10}(d) + 31.53 [dB]$$

Equation 9. Attenuation of free space at 900 Mhz.

The LoRaWAN protocol establishes a packet format that can be seen in Figure 5, which shows that the higher the error correction rate (CR), the longer the packet will be for a given payload.



**Figure 5 - LoRaWAN package structure. Retrieved from (Semtech, 2013)**



The air time a LoRaWAN package is given by Equation 10.

$$T_{Packet} = T_{Preamble} + T_{Payload} [seg]$$

Equation 10. LoRaWAN Air Time

The preamble time LoRaWAN package is given by Equation 11.

$$T_{Preamble} = (n_{Preamble} + 4.25) * T_s [seg]$$

Equation 11. LoRaWAN Preamble Time

The time of a LoRaWAN symbol is given by Equation 12 and is related to the symbol rate.

$$T_s = \frac{1}{R_s} [seg]$$

Equation 12. LoRaWAN Symbol Time

The LoRaWAN symbol rate is given by Equation 13 and is related to the channel bandwidth (BW) and the spreading factor (SF).

$$R_s = \frac{BW}{2^{SF}}$$

Equation 13. LoRaWAN Symbol Rate

The total time of the LoRaWAN payload is given by Equation 14 and is related to the channel bandwidth (BW) and the spreading factor (SF).

$$T_{Payload} = PL_{Symb} * T_s$$

Equation 14. LoRaWAN Payload Time

The total time of a LoRaWAN package is given by Equation 15 that derives from Equation 14, Equation 11, Equation 12 and Equation 13 and is related to the channel bandwidth (BW) and the spreading factor (SF).

$$T_{Packet} = (n_{Preamble} + 4.25 + PL_{Symb}) * \frac{2^{SF}}{BW}$$

Equation 15. LoRaWAN packet time

Considering the previous points, it can be summarized for a LoRaWAN system that the air time of a packet is multiplied exponentially for higher spreading factor values as shown in Table 2.

**Table 2 - Spreading Factor vs Chip Length**

SF	$2^{SF}$
6	64
7	128
8	256
9	512
10	1024
11	2048
12	4096

Using Table 3 the link budget is displayed for different spreading factors based on the gain of 2.5 dB for each factor step. For the elaboration of said use table a typical receiver noise figure of 6 dB, as specified in the manuals of the receiver and a maximum transmission power of 14 dBm, as specified for LoRaWAN 2.0. (Semtech, 2013)

**Table 3 - Spreading Factor vs. Link Sensitivity and Budget**

SF	chips/symbol	SNR Limit [dB]	Noise Figure [dB]	BW [Hz]	Sensitivity [dBm]	TX Power [dBm]	Link Budget [dB]
6	64	-5	6	125000	-122.03	14.00	-136.03
7	128	-7.5	6	125000	-124.53	14.00	-138.53
8	256	-10	6	125000	-127.03	14.00	-141.03
9	512	-12.5	6	125000	-129.53	14.00	-143.53
10	1024	-15	6	125000	-132.03	14.00	-146.03
11	2048	-17.5	6	125000	-134.53	14.00	-148.53
12	4096	-20	6	125000	-137.03	14.00	-151.03

Considering Equation 9, this puts the maximum theoretical reach of LoRaWAN for SF=12 at 900 Km. The current world record using standard equipment is 702 km (The Things Network, 2017).

In Table 4, it is clear, that the impact of using a larger spreading is significant on the total bitrate of the channel.

**Table 4 - Channel Bitrate vs Channel Bandwidth, Spreading Factor and Coding Rate**

		SF						
BW	500000	6	7	8	9	10	11	12
CR	0	46875	27344	15625	8789	4883	2686	1465
	1	37500	21875	12500	7031	3906	2148	1172
	2	31250	18229	10417	5859	3255	1790	977
	3	26786	15625	8929	5022	2790	1535	837
	4	23438	13672	7813	4395	2441	1343	732
BW	250000	6	7	8	9	10	11	12
CR	0	23438	13672	7813	4395	2441	1343	732
	1	18750	10938	6250	3516	1953	1074	586
	2	15625	9115	5208	2930	1628	895	488
	3	13393	7813	4464	2511	1395	767	419
	4	11719	6836	3906	2197	1221	671	366
BW	125000	6	7	8	9	10	11	12
CR	0	11719	6836	3906	2197	1221	671	366
	1	9375	5469	3125	1758	977	537	293
	2	7813	4557	2604	1465	814	448	244
	3	6696	3906	2232	1256	698	384	209
	4	5859	3418	1953	1099	610	336	183

### 3.2. Effects of Aloha in LoRaWAN


The MAC of LoRaWAN is based on Pure ALOHA. If we define S as the average number of packets generated per interval; the traffic source  $\lambda$  consists of a high number of users who form an independent poisson source with an aggregate packet rate of X packets/s, the packet time width is supposedly fixed with a period of T seconds. It can be considered that each user generates packets infrequently. S can also be expressed as the channel throughput rate. A node delays the transmission of a previously collided packet with a random time. Therefore, the total traffic is not only new packets but repetition of retransmission of collided packets

$$\begin{aligned}
 S &= \lambda T \\
 G &\geq S \\
 G(n) &= \lambda(n)T \\
 S &= G(n)P_{Suc} = \lambda(n)T * e^{-\lambda(n)2T}
 \end{aligned}$$

Equation 15. LoRaWAN packet time

In Pure ALOHA, a successful transmission happens if the channel is free during the time period 2T (vulnerability period). The probability that there are no transmissions in the 2T period is

Psuc. The total channel traffic could be expressed as presented in Equation 15. With these constraints, it is possible to show that the maximum channel throughput is 18% (Kleinrock, 1975) as shown in figure 5. (Polloneli, 2019)



**Figure 6 - Channel Thoughtput vs Channel Load for LoRaWAN**

The total maximum adjusted bandwidth per channel considering Aloha MAC for 125 KHz Channels is shown in table 5.

**Table 5. Channel Bitrate vs Spreading Factor and Coding Rate for Aloha MAC**

0.18	125000	6	7	8	9	10	11	12
CR	0	2109	1230	703	396	220	121	66
	1	1688	984	563	316	176	97	53
	2	1406	820	469	264	146	81	44
	3	1205	703	402	226	126	69	38
	4	1055	615	352	198	110	60	33

Considering an 8 channel (125 KHz.) gateway, the total bandwidth is found in Table 6.

**Table 6 - Total gateway Bitrate vs Spreading Factor and Coding Rate for Aloha MAC in an 8-channel gateway.**

8	125000	6	7	8	9	10	11	12
CR	0	16875	9844	5625	3164	1758	967	527
	1	13500	7875	4500	2531	1406	773	422
	2	11250	6563	3750	2109	1172	645	352
	3	9643	5625	3214	1808	1004	552	301
	4	8438	4922	2813	1582	879	483	264

### 3.3. Signal Propagation Models

There are a few propagation models for signals in the UHF frequency range, and the best known is the Okumura-Hata model. Hata's model gives the value of pathloss at given distance between a base station and a mobile user. It considers several factors such as frequency, antenna heights and others (Hata, 1980).

$$L_P = 69.55 + 26.26 \log(f) - 13.82 \log(h_B) - CH + [44.9 - 6.55 \log(h_B)] \log(d)$$

For a medium small city

$$CH = (1.1 \log(f) - 0.7)h_M - (1.56 \log(f) - 0.8)$$

$L_P$  = Path loss in urban areas. Unit: decibel (dB)

$h_B$  = Height of base station above ground (meters)

$h_M$  = Height of mobile station above ground (meters)

$f$  = Transmit frequency (MHz)

CH = Antenna height correction factor

d = Distance from gateway to device (km)

This model works very well for most IoT applications however its accuracy is limited by a minimum base station height of 30m (98 feet). In this analysis, the base station will be on the fiber node location, which is typically between 8-12m (26 to 39 feet) in height, so this model is not directly applicable.

A more accurate model for low height antennas is described in (Vilardi, 2012) and better applies for the scenario studied in this paper. This model considers urban and suburban models as well as outdoor/indoor commercial concrete walls and residential wood frame and/or brick wall applications.

This model has been done on a 900 MHz band, so perfectly applies for this study.

$$L_P = 20 \log(h_B) + 20 \log(h_M) - 43.36 \log(d) - A - B - C$$

$L_P$  = Path loss in urban areas. Unit: decibel (dB)  
 $h_B$  = Height of base station above ground (meters)  
 $h_M$  = Height of mobile station above ground (meters)  
 $d$  = Distance from gateway to device (m)  
 $A$  = Constant – Area type - 24.3 dB in suburban and 29.3 dB in urban  
 $B$  = Constant – Building Attenuation – 0 dB Outdoor, 17.7 dB for Commercial Building (concrete) and 5.4 dB for suburban house (wood and brick frame)  
 $C$  = Constant – Shadowing Affect - 0 dB Outdoor, 9.3 dB for Commercial Building and 6.4 dB, for residential

### 3.4. Frequency Plan

The spectrum in the USA for LoRaWAN has 64 uplink channels available (125 kHz each) (channels 0-63) starting at 902.3 MHz which increment every 200 kHz up to 914.9 MHz  
 There are 8 overlapping uplink channels (500 kHz each) (channels 64-71) from 903 MHz which increment every 1.6 MHz up to 914.2 MHz.  
 For gateway to node communication, there are 8 downlink channels (500 kHz each) (channels 0-7) from 923.3 MHz which increment every 600 kHz up to 927.5 MHz.

**Table 7 - Spectrum bands for LoRaWAN in the USA**

Uplink sub-bands	Frequency range (MHz)	Channels
Sub-Band 1	902.3 - 903.7	0-7
Sub-Band 2	903.9 - 905.3	8-15
Sub-Band 3	905.5 - 906.9	16-23
Sub-Band 4	907.1 - 908.5	24-31
Sub-Band 5	908.7 - 910.1	32-39
Sub-Band 6	910.3 - 911.7	40-47
Sub-Band 7	911.9 - 913.3	48-55
Sub-Band 8	915.5 - 914.9	56-63
Downlink sub-bands	Frequency range (MHz)	Channels
Downlink sub-band	903 - 914.2	64-71

### 3.5. Scenario Modeling

A) Dense Urban with Concrete Buildings and Indoor sensors

For a typical LoRaWAN network with a gateway running on a Fiber node on the strand let's assume  $h_B = 9m = 30\ ft$  and  $h_M = 1.5m = 4\ ft$   
 Model boundary conditions

$$h_B = 9m$$


$$\begin{aligned}
h_M &= 1.5m \\
A &= 29.3dB \\
B &= 17.7dB \\
C &= 9.3dB
\end{aligned}$$

$$L_P = 20 \log(9) + 20 \log(1.5) - 43.36 \log(d) - 29.3 - 17.7 - 9.3$$

$$L_P = 19.08 + 3.52 - 43.36 \log(d) - 29.3 - 17.7 - 9.3$$

$$L_P = -43.36 \log(d) - 33.7$$

Plotting the pathloss as a function of the distance as in Figure 7, the thresholds for operation in the most resilient and faster spreading factors are marked, SF=12 (red) and SF=6 (green).



**Figure 7 - Pathloss vs Distance for Dense Urban scenario**

The maximum distance for different spreading factors can be calculated based on its link budget as show in table 4.

$$d = 10^{\left(\frac{L_B + 33.7}{-43.36}\right)}$$

For SF=12,  $L_B = -151.08$  dB

$$d = 10^{\left(\frac{-151.08 + 33.7}{-43.36}\right)}$$

$$d = 509.5m$$

For SF=6,  $L_B = -136.03$  dB

$$d = 10^{\left(\frac{-136.03+33.7}{-43.36}\right)}$$

$$d = 229.1m$$

The conclusion is that in a dense urban scenario with concrete buildings a LoRaWAN network can operate with the least efficient transmission mode up to 509.5m (1671 feet) and the limit for the most efficient profile is 229.1m (751 feet).

#### B) Sparse Suburban with Houses and Indoor sensors

For a typical LoRaWAN network with a gateway running on a fiber node on the strand, assume  $h_B = 9m = 30\ ft$  and  $h_M = 1.5m = 4\ ft$ .

Model boundary conditions

$$h_B = 9m$$

$$h_M = 1.5m$$

$$A = 24.3dB$$

$$B = 5.4dB$$


$$C = 6.4dB$$

$$L_p = 19.08 + 3.52 - 43.36 \log(d) - 24.3 - 5.4 - 6.4$$

$$L_p = -43.36 \log(d) - 13.5$$

Plotting the pathloss as a function of the distance as in Figure 8, the thresholds for operation in the most resilient and faster spreading factors are marked, SF=12 (red) and SF=6 (green).





**Figure 8 - Pathloss vs Distance for Sub-Urban scenario**

$$d = 10^{\left(\frac{L_B + 13.5}{-43.36}\right)}$$

For SF=12,  $L_B = -151.08$  dB

$$d = 10^{\left(\frac{-151.08 + 13.5}{-43.36}\right)}$$

$$d = 1489.25m$$

For SF=6,  $L_B = -136.03$  dB

$$d = 10^{\left(\frac{-136.03 + 13.5}{-43.36}\right)}$$


$$d = 669.7m$$

The conclusion is in a dense urban scenario with concrete buildings a LoRaWAN network can operate with the least efficient transmission mode up to 1489.25m (0.89 miles) and the limit for the most efficient profile is 669.7m (0.42 miles).

## 4. HFC Networks

### 4.1. HFC Node + 0 Networks

HFC networks with 0 amplifiers in cascade typically have short coaxial cable runs that expand in average no more than 200 meters (600 feet) from the node location and even less than 150 meters in dense cities like the example in Figure 9.




**Figure 9 - Node+0 sample node distribution**

That scenario is really aligned if the LoRaWAN gateway is collocated in the fiber node with the LoRaWAN configuration presented in section 3.3 where a maximum distance of 229 m (751 feet) is required to serve stations using maximum capacity SF=6.

In this scenario not all fiber nodes need to have a LoRa gateway because the distance exceeds signal reach.

### 4.2. HFC Node+1/2 Networks

HFC Networks with 1 amplifier in cascade typically have medium coaxial cable runs that expand in average 500 meters (1500 feet) from the node location in not so dense cities with houses as illustrated in Figure 10.



**Figure 10 - Node+1/2 sample node distribution**

This scenario is really aligned if the LoRaWAN gateway is collocated in the fiber node with the LoRaWAN configuration presented in section 3.4; where a maximum distance of 669 m (751 feet) is required to serve stations using maximum capacity SF=6.

In this scenario all fiber nodes need to have a LoRa gateway, as the distance is well aligned with the wireless reach.

## **5. Network Density and Throughput**

Considering the different IoT applications, it is very important to be able to model the traffic patterns and device densities according to the area density. In this section we will consider public usage and some residential use cases where data traffic is not intensive. The device density per suburban and urban areas was taken from (Huang, 2011).

Table 8 Shows the average data rates and reporting periods for public applications.

**Table 8 - Public Applications Communication Parameters**

	Reporting Period [s]	Average Transaction Rate [1/s]	Average Message Size [bytes]	Data Rate [bps]
Credit Machine in Grocery	120	0.00833	24	1.6000
Credit Machine in Shop	1800	0.00056	24	0.1067
Roadway Signs	30	0.03333	1	0.2667
Traffic Lights	60	0.01667	1	0.1333
Traffic Sensors	60	0.01667	1	0.1333

Table 9 shows the average data rates and reporting periods for residential applications.

**Table 9 - Residential Applications Communication Parameters**

	Reporting Period [s]	Average Transaction Rate [1/s]	Average Message Size [bytes]	Devices per Home	Data Rate [bps]
Smart Meters	9100	0.00011	20	3	0.053
Home Security System	600	0.00167	20	1	0.267
PHEV	4200	0.00024	12	2	0.046
				Total per Home	0.365

Table 10 shows the average density of devices in different areas.

**Table 10 - Device density for cities**

	density per Square Meter					
	Homes	Grocery Stores	Restaurants	Road Signs	Traffic Lights	Traffic Sensors
Urban (NYC)	0.00384400	0.00020947	0.00220000	0.00031647	0.00001503	0.00001503
Suburban (Washington)	0.00147922	0.00002312	0.00034988	0.00009433	0.00001144	0.00001144

Table 11 shows the average device quantity per gateway on different areas considering the radius where it is most likely that those devices will use SF=7, and its corresponding average traffic in bits per second.

**Table 11 - Devices and Bitrate per gateway**

			devices per gateway						
	Cell Radius [m]	Cell Area [sq. m]	Homes	Grocery Stores	Restaurants	Road Signs	Traffic Lights	Traffic Sensors	Total Devices
Urban (NYC)	200	125664	483.05	26.32	276.46	39.77	1.89	1.89	829.38
Suburban (Washington)	500	785398	1161.78	18.16	274.80	74.08	8.99	8.99	1546.79
			bps per gateway						
Urban (NYC)	200	125664	176.38	42.12	29.49	10.61	0.25	0.25	259.09
Suburban (Washington)	500	785398	424.20	29.06	29.31	19.76	1.20	1.20	504.72

This table shows that the required bandwidth is well below the maximum bandwidth per serving area according to section 3.3.

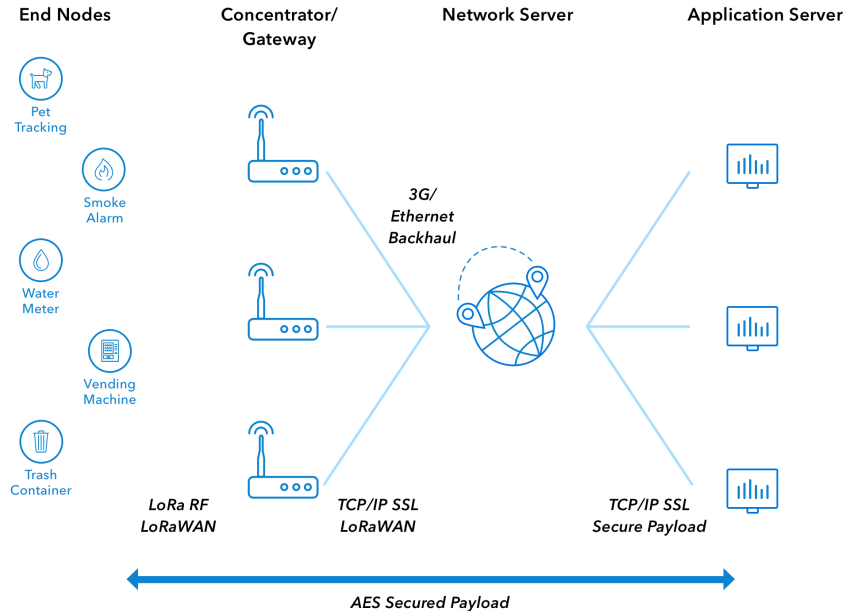
## 6. Powering and Backhaul Connectivity

### 6.1. Powering

Most standalone LoRaWAN 8 channel gateways on the market have a power consumption below 5W, even with 3G/LTE backhaul and embedded GPS. Using an embedded cable modem or ethernet connection to the fiber node instead of wireless backhaul can average a power consumption of 5W. This is well below the total power consumption of a traditional HFC fiber node, which ranges from 70 to 120W and 140 to 160W for a distributed access enabled node. Embedding the LoRaWAN gateway in the fiber node provides significant benefits on the powering area, as the gateway can be powered with the fiber node regulated power supply.

### 6.2. Backhaul Connectivity

In any LoRaWAN networks (Figure 11) a backhaul connection is required to connect the gateway devices to the application servers. In an HFC network model this backhaul can be provided by an embedded cable modem or direct ethernet connection in the fiber node. This approach provides a big benefit for the reliability of the backhaul, requiring negligible resources from the DOCSIS® network. As analyzed in the previous sections, an 8-channel gateway can require a maximum of 13.5 kbps of bandwidth. Compared to the hundreds of megabits available for DOCSIS® modems, this is totally negligible. Given this, a real-time high priority delivery unsolicited grant service (UGS) service flow can be used on that modem in order to ensure immediate delivery of the information.



**Figure 11 - LoRAWAN Architecture. Retrieved from TheThingsNetwork.**


## 7. LTE Based IOT

It is really important to understand that LPWAN networks and LTE based IOT are not mutually exclusive options of IOT networks. As seen before, LoRaWAN provides a really good alternative for low bandwidth sensor applications like smart meters or city infrastructure support. On the other hand, LTE based solutions provide higher bandwidth capabilities and an evolution to support low power end devices as shown in table 12.

**Table 12 - LTE IOT Protocols**

	LTE Cat 0	LTE Cat M1	LTE Cat NB1
3GPP Release	Release 12	Release 13	Release 13
Downlink Peak Rate	1 Mbit/s	1 Mbit/s	250 Kbit/s
Uplink Peak Rate	1 Mbit/s	1 Mbit/s	250 Kbit/s (multi-tone)
			20 Kbit/s (single-tone)
Number of Antennas	1	1	1
Duplex Mode	Full or Half Duplex	Full or Half Duplex	Half Duplex
Device Receive Bandwidth	1.4 – 20 MHz	1.4 MHz	180 kHz
Receiver Chains	1 (SISO)	1 (SISO)	1 (SISO)
Device Transmit Power	23 dBm	20 / 23 dBm	20 / 23 dBm
Power	+	-	--


Even if this is not the main focus of this paper, there are several options of HFC backhauled LTE picocells which can provide access infrastructure for LTE IOT devices, in the three below pictures typical coverage for an LTE Picocell of 2W of power is shown.



Cell Diameter: 366m  
 Cell Radius: 183m  
 Cell Area (m): 50,691m<sup>2</sup>  
 Cell Area (ft): 545,635 sq.-ft



Cell Diameter: 312m  
 Cell Radius: 156m  
 Cell Area (m): 22,304m<sup>2</sup>  
 Cell Area (ft): 240,076 sq.-ft



Cell Diameter: 346m  
 Cell Radius: 173m  
 Cell Area (m): 17,773m<sup>2</sup>  
 Cell Area (ft): 190,873 sq.-ft

Comparing the coverage of these LTE picocells with LoRaWAN gateways it can be seen that the coverage is really similar and in line with the size of an HFC fiber node.

## Conclusion

This paper described the state of the art of current low power area networks for IoT, then analyzed the key parameters related to link budget calculation a bandwidth capacity for LoRaWAN networks serving urban and suburban areas.

Next, the best coverage zone for these areas was calculated and correlated that with different HFC network designs, followed by a sensor density and required bandwidth for that optimal coverage zone. Power and Backhauling was analyzed using the fiber node as connection point.

Lastly, a brief comparison between LTE based IoT and LoRaWAN was shown, focusing on their differences and how they can be both deployed on HFC Networks.

As a final conclusion: HFC networks provide the right supporting infrastructure to add support for IoT networks without significant effort. Wireless coverage zones are well aligned with fiber node locations and provide required power and IP connectivity. Adding IoT services can provide MSOs with an extra revenue stream without a significant investment and allow them to provide appropriate services.

# Abbreviations

AP	access point
bps	bits per second
dB	decibel
DOCSIS	data over cable service interface specification
FEC	forward error correction
GPS	global positioning system
HFC	hybrid fiber-coax
Hz	hertz
IoT	internet of things
IP	internet protocol
ISBE	International Society of Broadband Experts
ITU	International Telecommunications Organization
IPSEC	internet protocol security
KPI	key performance indicator
LPWAN	low power wide area network
LoRaWAN	long range wide area network
LTE	long term evolution
MAC	media access control
MSO	multi service operator
OTT	Over the top
RF	radiofrequency
SCTE	Society of Cable Telecommunications Engineers
SNR	signal to noise ratio
UGS	unsolicited grant service
WAN	wide area network
WiFi	wireless fidelity



## Bibliography & References

- Adelantado, F., Vilajosana, X., Tuset-Peiro, P., Martinez, B., Melia-Segui, J., & Watteyne, T. (2017). Understanding the Limits of LoRaWAN. *IEEE Communications Magazine*, 34 - 40.
- Ashton, K. (2009, Enero 22). *That 'Internet of Things' thing: in the real world, things matter more than ideas*. Retrieved Junio 30, 2018, from <http://www.rfidjournal.com/articles/view?4986>
- Bankov, D., Khorov, E., & Lyakhov, A. (2016). On the Limits of LoRaWAN Channel Access. *2016 International Conference on Engineering and Telecommunication* (pp. 10-14). IEEE.
- Ericsson. (2016). *Internet of Things forecast*. Retrieved from <https://www.ericsson.com/en/mobility-report/internet-of-things-forecast>
- Gershenfeld, N. (1999). *When Things Start to Think*. Henry Holt and Co.
- Haas, Z. J., & Deng, J. (2003). On optimizing the backoff interval for random access schemes. *IEEE TRANSACTIONS ON COMMUNICATIONS, VOL. 51*, 2081-2090.
- Hassan, Q. (2018). *Internet of Things A to Z: Technologies and Applications*. Wiley-IEEE Press.
- ITU, T. S. (2012). *Overview of the Internet of things*. International Telecommunication Union.
- Kranz, M. (2018, Enero 11). *6 ways the Internet of Things is improving our lives*. Retrieved from World Economic Forum: <https://www.weforum.org/agenda/2018/01/6-ways-the-internet-of-things-is-improving-our-lives/>
- Libelium. (2018). *50 Sensor Applications for a Smarter World*. (Libelium) Retrieved June 30, 2018, from [http://www.libelium.com/resources/top\\_50\\_iot\\_sensor\\_applications\\_ranking/](http://www.libelium.com/resources/top_50_iot_sensor_applications_ranking/)
- LoRa Alliance. (2017). LoRaWAN Specification 1.1. *LoRaWAN Specification 1.1*. San Ramon, CA, USA: LoRa Alliance. Retrieved from [https://loralliance.org/sites/default/files/2018-05/lorawan1\\_0\\_2-20161012\\_1398\\_1.pdf](https://loralliance.org/sites/default/files/2018-05/lorawan1_0_2-20161012_1398_1.pdf)
- Nikolic, M., Chandramouli, B., & Goldstein, J. (2017). Enabling Signal Processing over Data Streams. *SIGMOD'17*. Chicago, Illinois, USA.
- Noreen, U., Bounceur, A., & Clavier, L. (2017). A study of LoRa low power and wide area network technology. *2017 International Conference on Advanced Technologies for Signal and Image Processing*. Fez, Morocco: IEEE.
- Petajajarvi, J., Mikhaylov, K., Roivainen, A., Hanninen, T., & Pettissalo, M. (2015). On the coverage of LPWANs: range evaluation and channel attenuation model for LoRa

- technology. *2015 14th International Conference on ITS Telecommunications (ITST)*. Copenhagen, Dinamarca: IEEE.
- Presidencia de la Nacion. (2018). Resolucion 581/2018. *Boletin Oficial de la Republica Argentina*.
- Proakis, J. (2000). *Digital Communications, 4 Ed.* McGraw Hill.
- Reynders, B., Meert, W., & Pollin, S. (2016). Range and Coexistence Analysis of Long Range Unlicensed Communication. *23rd International Conference on Telecommunications (ICT)*. Thessaloniki, Greece: IEEE.
- Sam Lee, J., & Miller, L. E. (1998). *CDMA Systems Engineering Handbook*. London: Artech Hoise.
- Sampieri, D. R., Collado, D. C., & Lucio, D. M. (2010). *Metodología de la investigación*. Mexico: McGraw Hill Education.
- Seller, O. B., & Sornin, N. (2013). *Europe Patent No. 13154071.8*.
- Semtech. (2013). *SX1272/3/6/7/8: LoRa Modem Designer's Guide*. Retrieved from [https://www.semtech.com/uploads/documents/LoraDesignGuide\\_STD.pdf](https://www.semtech.com/uploads/documents/LoraDesignGuide_STD.pdf)
- Semtech Corporation. (2015, May). *LoRa™ Modulation Basics*. Retrieved from Semtech Corporation: <https://www.semtech.com/uploads/documents/an1200.22.pdf>
- Smith, I. G. (2012). *The Internet of Things 2012 - New Horizons*. Halifax, UK: Platinum.
- Valenta, V., Masalek, R., Baudoin, G., Villegas, M., Suarez, M., & Robert, F. (2010). Survey on Spectrum Utilization in Europe: Measurements, Analyses and Observations. *5th International ICST Conference on Cognitive Radio Oriented Wireless Networks and Communications*.
- Wixted, A. J., Kinnaird, P., Larijani, H., Tait, A., Ahmadinia, A., & Strachan, N. (2016). Evaluation of LoRa and LoRaWAN for wireless sensor networks. *2016 IEEE SENSORS*. Orlando, FL, USA: IEEE.

# **Can Future Networks Survive Without Artificial Intelligence?**

## **A Comprehensive Overview of Frameworks and Applications of AI**

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

**Emilia Gibellini**

Scientist

Telecom Argentina S.A.

egibellini@teco.com.ar

**Carlos Germán Carreño Romano**

Scientist

Telecom Argentina S.A.

caromano@teco.com.ar

**Gabriel Carro**

VP Engineer and R&D

Telecom Argentina S.A.

gcarro@teco.com.ar

# Table of Contents

Title	Page Number
Table of Contents .....	2
Abstract.....	4
Content .....	4
1 Introduction.....	4
1.1 Motivations and Overview.....	4
1.2 Artificial Intelligence.....	5
1.2.1 A brief history.....	5
1.3 Data Analytics .....	6
2 AI and Data Analytics in Communications, Networks and Services.....	7
2.1 Definitions .....	8
2.2 Applications.....	9
2.2.1 Traffic classification.....	11
2.2.2 Spectrum use.....	12
2.2.3 Proactive Network Maintenance .....	12
2.2.4 Capacity planning .....	12
2.2.5 Beamforming .....	13
2.2.6 Massive MIMO .....	13
2.2.7 Encoding.....	13
2.2.8 NFV/SDN.....	14
2.2.9 Wireless and Mobile Networks .....	14
2.2.10 Network Slicing .....	15
2.2.11 Mobile Edge Computing .....	15
2.2.12 Optimizing Customer Experience on Video using Machine Learning.....	16
3 Technical and Theoretical Framework.....	16
3.1 Machine Learning.....	16
3.1.1 Artificial Neural Networks .....	17
3.1.2 Time Series models .....	17
3.1.3 A/B Testing .....	17
3.1.4 Clustering .....	18
3.1.5 K-Nearest Neighbors.....	18
3.1.6 Natural Language Processing .....	18
3.1.7 Deep Learning techniques.....	18
3.2 Non-Machine Learning Tools.....	20
3.2.1 Locality-Sensitive Hashing .....	20
3.2.2 Collaborative filters.....	20
3.2.3 Process Mining .....	20
3.2.4 Digital Twins .....	21
3.3 The Knowledge Plane.....	21
4 Strategy .....	23
Conclusions.....	25
Abbreviations.....	26
Bibliography & References .....	27

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - The evolution of Artificial Intelligence. ....	6
Figure 2 - Worldwide amount of data created per year in zettabytes.....	7
Figure 3 - Map of AI .....	9
Figure 4 - Potential annual impact of AI technologies in global industries by 2025. ....	10
Figure 5 - Artificial Neural Network conceptual diagram .....	17
Figure 6 - The four planes in network architecture.....	21
Figure 7 - Plane of Knowledge in a SDN.....	22
Figure 8 - Detailed Plane of Knowledge in a SDN .....	23
Figure 9 - Telecom Argentina knowledge-defined network .....	24

# Abstract

In the last two years, Artificial Intelligence (AI) has become a trending topic in the main conferences and international exhibitions. One of the first questions that arises is to refer to the operators and vendors when they use the term AI. These days it really seems that everything gets better with AI, then another important point is hype vs reality. Certainly, there has been a lot of activity in trying to use AI to improve cable network operation and customer experience. On the other hand, future networks will generate a volume of data with exponential growth, and this poses a great challenge. We present in this technical paper a framework of AI technologies, how operators can extract the value from data and gain new insights and efficiencies using artificial intelligence, where to start (the right points) in our networks to apply AI and what human skills will be needed. Finally, we present our biggest challenge: our AI and Data Strategy that will drive our transformation from an MSO and MNO to a Digital Service Provider (DSP).

# Content

## 1 Introduction

In our networks and services, AI has the potential to change, the way we operate, and to become the foundation of the transformation that leads to the fourth industrial revolution. But this requires hard work, a long-term commitment, and a deep cultural change.

Telecom Argentina is currently in a process of integration of its MSO and MNO networks. To accompany the fourth industrial revolution, we need to become a Digital Service Provider (DSP). Telecom is going to a Multiservice Convergent Network and Multi devices approach, where the Client / User can consume their own and third party services (platforms) from any device and connected to any of our access networks.

A digital service is defined by The European Union (EU) Agency for Cybersecurity, more specifically by the Network and Information Security (NIS) directive 2015/1535 as “any service normally provided for remuneration, at a distance, by electronic means and at the individual request of a recipient of services”. According to the NIS directive, DSPs are limited to only three types of services: cloud, online market places and search engines [1].

The term Digital Service Provider applies to any company that distributes media online. In the case of telcos, it is an organization that has moved on from offering core, traditional telecom services, to providing mobile broadband access, services, content and apps, all sold directly from the device [2].

The DSP is not merely a dumb pipe offering shared access to a common utility; it is an online, real-time business that deals with countless transactions every day, managing high volumes of data traffic and multiple devices per user, and often multiple users per account. The mobile and fixed landscape has changed dramatically and CSP's are fine-tuning their businesses, and their network infrastructure, to cater for the digital needs of the data-hungry customer.

### 1.1 Motivations and Overview

The term AI has recently become a buzzword and entered in a sort of semantic satiety similar to what happened a few years ago with the term Big Data. But what exactly is AI? What is it used for? Who invests in it? Can future networks survive without it?

All Operations Support Systems (OSS) in our current and future networks generate a huge amount of data. How can we generate a culture of innovation driven by data? What is the value of the data?

The final aim of all of these efforts is to be able to offer our services in an adequate way for our next generations of clients. They are nowadays putting the requirements in the market and driving the evolution of technologies. Ultimately our clients do not buy technologies, they buy services.

All these services will be enabled by technologies such as 10G, 5G, SDN / NFV, Holographics Displays etcetera. Operating, managing and provisioning future services with automation processes becomes essential.

Operating, managing, provisioning future services with automation processes must be essential. If we want a complete automation, we will need AI technology and data analytics tools.

This technical paper proposes a comprehensive overview of frameworks and applications of AI to network's design, management and operation. We introduce a distributed cognitive system that permeates the network, that is called the knowledge plane. This paper is organized as follows. Following this section we define AI and Data Analytics. In section 2, we expose some applications of AI and Data Analytics to communications, networks and services. Section 3 exposes the technical and theoretical framework, we briefly describe some ML and non-ML algorithms that have found useful and how we are applying them, and we introduce the concept of the knowledge plane. The last section, which is section 4, outlines the key challenges in our path.

## **1.2 Artificial Intelligence**

In this work we define our scope of AI, which is much more related to the current Machine Learning (ML) framework, we propose some technical and theoretical issues of AI, and an overview of applications to networks. In the next sections we hope to answer some of the starter questions.

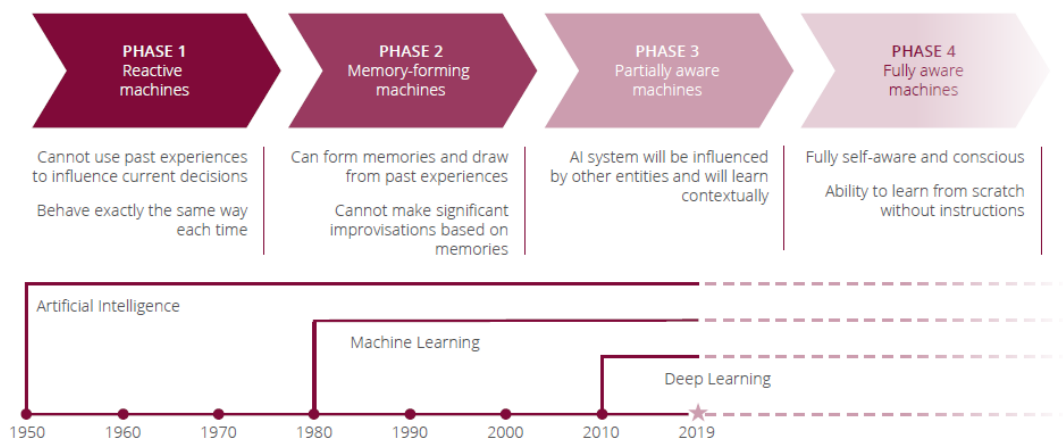
The concept of AI is not new. It was already established in the 1950's as machine intelligence that is capable to process, analyze, and react to input and changing situations by itself. We can say that today AI is a general term that describes multiple technologies, none of which fit completely to the original definition of AI.

According to our chosen bibliography, intelligence is the capacity to do the right thing at the right time, in a context where doing nothing (or making no change in behavior) would be worse. Intelligence then requires the capacity to perceive contexts for action, the capacity to act, and the capacity to associate contexts to actions. AI, by convention, describes (typically digital) artifacts that demonstrate any of these capacities [3].

### **1.2.1 A brief history**

In the 1940's and 1950's, scientists in the fields of Mathematics, Engineering and Computer Science explored the possibilities of artificial brains and tried to define the intelligence of a machine. In 1950, Alan Turing presented a test known today as the Turing test, which defined the concept of Machine Intelligence [4]. John McCarthy is credited with the first definition of AI as *the science and engineering of making intelligent machines*, during a workshop at Dartmouth College in 1956 [5]; Russell and Norvig present a classification of the available definitions of AI in two senses: one based on the function expected to be performed (comparing processes/reasoning of the machine versus the outcome/behavior that it exhibits) and the other about the metrics used for assessing the success of AI (human performance versus an ideal standard of rationality) [6], and in their seminal 1995 book, they give it an arresting description, *agents that*

receive percepts from the environment and take actions that affect that environment. In Figure 1, we present the evolution and stages of AI.



**Figure 1 - The evolution of Artificial Intelligence.**  
Source: Statista, Nvidia.

In 1959 Machine Learning (ML) arises, Arthur Samuel defines it as *the field of study that gives computers the ability to learn without being explicitly programmed* [7]. So, the word learning in this sense is used by analogy with the learning process in animals rather than in humans.

As we see the term AI is, in broad sense, the human intelligence process replicated by a machine, which is an auto programmable entity. People think in robots, in the Turing test, a person asking a robot and being answered by an entity which could be a human itself, something like a chatbot but more realistic. Critics go deeply through the philosophical plane.

In our opinion, this meaning will not be possible at least in the next few years except in movies, so we prefer to write about AI in a narrow sense. It is interpreted as a subset of programming techniques based on statistics, sometimes referred to as ML. In order to be more specific, our approach comprehends the collection of data, polling of the network, generation of useful models, information and algorithms based on business rules. From a statistical outlook, it is an expansion of the classical statistical methods which includes the conception supervised, unsupervised and reinforcement methods. From a computer science point of view, AI differs from classical programming which basically is rules + data, in the idea of rules learned from the data.


In recent years, the progress in AI has been driven mainly by the generation and availability of huge amounts of data and augmented computing power.

### 1.3 Data Analytics

In the last two years, 90% of the world's total data load has been generated, yet only 1% has been processed [8]. With data volumes set to increase exponentially (Figure 2) over the next decade, significant challenges are waiting on our journey.



Companies who can extract the value from data and gain new insights and efficiencies using AI and Data Analytics could boost the fourth industrial revolution [9]. *Data is going to introduce social and economic changes that we see perhaps once or twice in a century*, said ex-CEO of Intel, Brian Krzanich, in his keynote speech at CES 2018.



**Figure 2 - Worldwide amount of data created per year in zettabytes.**  
Source: Statista.

One of the main differences between Data Analytics and ML is that the former extracts value from data by applying exploratory techniques while the latter implies a certain level of inference.

## 2 AI and Data Analytics in Communications, Networks and Services

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
- ITU-T FG ML5G (Est. in 11/2017) 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) (Est. in 2/2017) Defining a cognitive network management architecture based on AI methods and context-aware policies; five deliverables have already been released

At the TM Forum [10] presentations and discussion panels, possible uses and applications in the telecommunications business were presented. Among the advantages of the use of ML, we can mention:

- Fast and automatic analysis of large volumes of data that are becoming more and more complex.
- Getting faster and more accurate results that allow to make reliable and repeatable decisions.
- Focus on behavioral analysis to detect and predict possible anomalous events at an early stage.
- Automate real-time analysis in the orchestration of end-to-end services in a virtualized world.
- Identification and mitigation of security threats in services through predictive analytics and ML to detect attacks that escape traditional preventative static defenses.

- Prediction of churn. Unlike traditional strategies, ML allows a multi-class classification of our clients, for example to predict whether they belong to a low, medium or high-risk class.
- Support for automation and management of network orchestration and traceability of end-to-end transactions across the network and OSS/BSS environment.

For the cable industry in particular, [11] provides an overview of ML algorithms, and how their potential applications could be applied:

- Software Defined Networks (SDN) Routing
- Profile Management on DOCSIS 3.1 cable modems [12]
- Proactive Network Maintenance (PNM): for DOCSIS
- HFC's Network Health KPI

Some applications are being implemented in:


- Internet Traffic Characterization
- Network Traffic Engineering
- Wi-Fi Proactive Network Maintenance (PNM)

## 2.1 Definitions

We have taken some definitions from those work groups and we started to build our own AI and Data Analytics technologies application framework in order to apply it to the Telecom Argentina's networks and services. We will mention some of the definitions from the framework.

- **Data Analytics:** monitoring data to look for patterns and anomalies (without applying intelligence) and applying those patterns towards effective decision making.
- **Artificial Intelligence:** the development of computer systems capable of performing tasks that normally require human intelligence; this includes visual perception, speech recognition, decision-making, and translation between languages.
- **Machine learning:** a type of AI that gives machines the ability to learn automatically and improve from experience without being explicitly programmed.
- **Deep learning:** takes machine learning further by processing information in layers, where the result or output from one layer becomes the input for the next.
- **Automation:** within MSOs and MNOs, this means automation of processes that were previously carried out by people; AI is an enabling technology that may (or may not) help with the process of automation.
- **Cognitive computing:** like AI, cognitive computing is based on the ability of machines to sense, reason, act and adapt based on learned experience, but whereas AI acts on its analysis to complete a task, cognitive computing provides the information to help a person decide.
- **SON (Self-Organizing Networks):** a technology for automating the planning, configuration, management, optimization and healing of mobile radio access networks; it was developed by 3GPP and is sometimes conflated with AI.
- **Explainable AI:** explainability 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) [13].

Conceptually speaking, it is clear from Figure 3 that machine learning is a subset of AI, which includes deep learning algorithms.

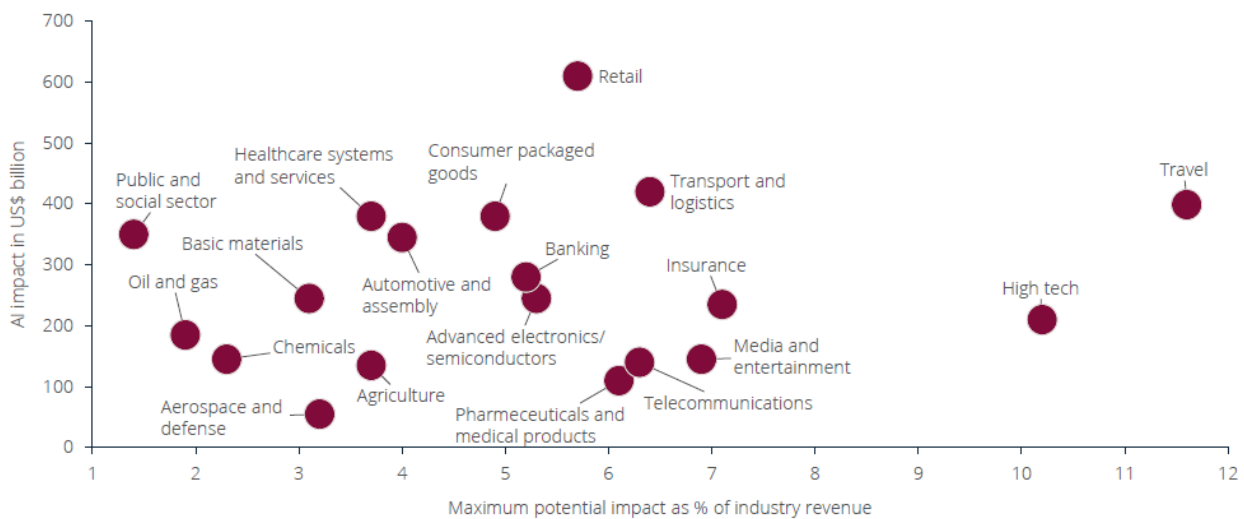


**Figure 3 - Map of AI**

Regarding this last point, and according to [14], [15] 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 models can be explainable, not always they are auditable. The discussion about audit AI is still open [16].

## **2.2 Applications**

Several applications from different fields of knowledge are being hyped during last few years. As shown in Figure 4, some of them include the transport and logistics industry, healthcare systems, agriculture, semiconductors industry, among others. In this document we will focus on network and service-related applications.



**Figure 4 - Potential annual impact of AI technologies in global industries by 2025.**

**Source: McKinsey Global Institute, Statista.**

From a global or top-level perspective, we will mention and describe some of the current and next few years applications of AI applied to networks and services, which we consider of potential interest in these industries. Some of them are related to current deployments and some of them are being tested in lab.

- **Management**

Cable Television industry, Telco and TechCo in general have some technical concerns at the management level. Some of them are the specific technology election for traffic classification and prioritization, the licensing of the spectrum, the outside plant, the attractiveness of the website, the commercial and technological strategy.

Dealing with the mass of users, devices and applications implies certain requirements for the network in order to service data. AI must be one of the strategic management tools for different purposes. Depending on the scenario, it could be used for reducing costs, customizing the user preferences, launching dynamic offers to the market, etcetera.

- **Maintenance**

Maintenance efforts, including corrective, predictive and proactive maintenance are very important. Nowadays Proactive Network Maintenance (PNM) work groups show a good picture about the need of ML algorithms for the classification of network impairments. Power control schemes could be improved also with tools of monitoring and analytics that combine not only hardware and systems raw data but also user experience information coming from, for example, social networks.

In the telco industry there is a trend of monitoring the user experience through specific teams associated with the focus on the service, and not only on the network. Despite the effort to change the vision, this is still considered as a way of corrective maintenance because it reacts when something is wrong. But there are also new ideas, for example digital twins, which could be used not only for troubleshooting but also for simulation of network changes. This approach currently involves tools for testing the effectiveness of a change at low costs through A/B testing. We will talk about these technologies in the rest of the document.

- **Engineering and architecture**

In 5G, big data and analytics are leveraged to extract massive patterns, especially at the physical (PHY) and medium access control (MAC) layers and enable self-organizing operations. Artificial neural networks (ANN) can be used to redefine communication networks, solving a number of non-trivial design problems at runtime and across layers for cognitive link adaptation, resource scheduling, signal classification, and carrier sensing/collision detection, among others [17].

For fixed networks there are several step by step changes within the SDN/NFV roadmap, which implies and make possible the application of algorithms for network deployment at the orchestration level. Machine Learning techniques are interesting in the administration of cloud platforms, particularly when dynamic scaling of capacity is a requirement within the applications. We will cover some notions about technical concerns of these technologies.

The design and optimization of wireless communication systems are becoming more and more challenging, due to the extreme key performance indicators (KPIs) for user experience, efficiency, performance and complex network environments. AI, which can exploit the increasingly massive datasets available from wireless systems, can be used to solve complex and previously intractable problems.

Many problems in wireless communication systems, such as decision making, resource optimization, and network management, can be cast in a form that is suitable to be solved by AI techniques. Based on this observation, it is essential to study when and how to apply AI technology to improve the performance of future wireless communication systems.

### **2.2.1 Traffic classification**

Characterization of network traffic has been proved to be useful to many applications, such as network engineering and management, quality of service (QoS) assurance, distribution of tiered services, automated re-allocation of network resources, security, and even censorship. It provides data on subscriber's usage, which enables the application of data mining to support business goals. Furthermore, governments expect the CSPs to provide the facilities for lawful interception, for which traffic classification is key

The classic approach, which consists on the observation of the port number, is not effective since some applications use well-known port numbers to disguise their traffic or use unpredictable port numbers. Data packet inspection (DPI) techniques base their strategy on the search of keywords in the data packets' headers, which can be expensive in terms of computational effort and infeasible in case the data is encrypted. Moreover, governments may impose privacy regulations constraining the ability of third parties to lawfully inspect payloads at all. In consequence, a new generation of methods, based on machine learning (ML), have been proposed [18].

For the application of ML algorithms, the analysis unit is the flow, which consists of a succession of IP packets that have the same protocol type, source address and destination address. Most approaches consist of classifying the flows according to their statistical properties (distribution of flow duration, flow idle time, packet inter-arrival time and packet lengths). The main challenge to the implementation of such machine learning solutions so far has been the lack of ground truth, highly non-stationary data distribution and algorithms that underperform. There still lies an opportunity for developing highly accurate classifiers.

### **2.2.2 Spectrum use**

AI can be determinant towards deciding on transmission schemes, access methods, carrier frequency and bandwidth, channel modeling and transmission power.

Concerning HFC networks, one of the most important MAC layer functions introduced by the DOCSIS 3.1 specification is profile management. For each cable modem (CM) a profile is set, meaning a set of subcarriers and modulation orders are configured, based on the channel signal-to-noise ratio (SNR). This aims to increase spectral efficiency. The profiles are defined by the CMTS, based on the relationships existing between SNR and modulation. There remains an opportunity for vendors and CSPs to work on AI algorithms and improve profile management.

Next generations of wireless networks depend heavily on channel state information (CSI), which is an estimate of the capacity of a communication link, used to optimize signal transmission. Current methods for the estimation of CSI entail considerable air interface resource overhead. It was discovered that exploiting linear correlations of CSI among, for example, co-located antennas, different time instances, and different frequency subcarriers can alleviate this problem. There exists potential in exploiting nonlinear CSI structures [19].

### **2.2.3 Proactive Network Maintenance**

Degradation of the physical network can be assessed through proactive network maintenance (PNM) practices, which are developed from the analysis of data collected from the Customer Premise Equipment (CPE). This allows for the detection of several network impairments, such as impedance mismatches, internal noise, attenuation, signal ingress and egress, suckouts, among others.

PNM has been around for at least 10 years and has reached maturity. However, new capabilities continue to be developed. An example of this is the expansion of diagnostic measurements provided by DOCSIS 3.1 specification.

Many improvements in the diagnosis of the network are the result of applying machine learning techniques to the CPE data, in conjunction with other data sources. Additionally, the scope of the analysis can be expanded by including other data sources, to assess network reliability and predict the probability of impairment, which, as shown in [20] can be used to reprioritize tasks and so avoid the appearance of further damages.

### **2.2.4 Capacity planning**

At Telecom Argentina, AI is determinant for capacity planning. In the case of the HFC network, we use time series to forecast traffic and resource utilization at total, per regional hub or local site, per service group and per subscriber levels. This leads to a data-driven design, and optimization of resources. Similarly, in the case of wireless networks, the same kind of longitudinal analysis is made at cell site, cell or antenna and subscriber level. Tuning network parameters enables us to cope with subscribers' requirements.

Regarding wireless networks, ML and other state-of-the-art techniques and technologies -such as digital twins- are key to prevent coverage holes. Information from sites and cities infrastructure can be combined and analyzed in order to improve coverage. The need for automated detection of holes is particularly relevant for the case of 5G millimeter-wave signals (mmWave), which are subject to random blockage.

Moreover, in 5G, the link capacity between users and base stations (BS) can be much higher compared to sub-6 GHz wireless systems. Meanwhile, due to the high cost of infrastructure upgrade, it would be difficult

to drastically enhance the capacity of backhaul links between mmWave BS and the core network. As a result, the data rate provided by backhaul may not be enough to support all mmWave links; hence, the backhaul connection becomes the new bottleneck. BS-UE link is characterized by high data rate and unstable connection, while the backhaul link is characterized by relatively limited data rate and stable connection. To balance this mismatch and enhance the system performance, efficient backhaul resource allocation to each user is necessary. Such adaptive control cannot be implemented by traditional resource allocation schemes due to the varying system dynamics [21].

### **2.2.5 Beamforming**

Beamforming technology aims to provide faster, stronger signals with longer range. In contrast to omnidirectional signal transmission, it consists of the broadcasting of signals in specific angles, in the form of beams. Provided that a higher number of active beams increases the system resource utilization, it becomes necessary to leverage system utilization through an optimal definition of the beams to be used. In this sense, AI could provide a solution.

User equipment (UE) measures the beam state information (BSI), which is based on measurements of beam reference signal (BRS), comprising of parameters such as beam index (BI) and beam reference signal received power (BRSRP). Given a set of potential beams, the best is the one with the highest signal strength a.k.a. RSRP. An algorithm could be trained to automatically find the best beam, considering the RSRP and/or the BI [17].

### **2.2.6 Massive MIMO**

Multiple-input multiple-output (MIMO) is a wireless technology that leverages link capacity, by exploiting multipath propagation. Normally, transmitted signals bounce off walls, ceilings, and other objects, reaching the receiving antenna multiple times at different angles and slightly different times. MIMO technology uses multiple, smart transmitters and receivers with an added spatial dimension, increasing performance and range [22]. MIMO antennas usually have two transmitting and two receiving elements, which double the capacity of a basic antenna. Massive MIMO goes further, using multiple elements simultaneously.

The weights for antenna elements for a massive MIMO 5G cell site are critical for maximizing the beamforming effect. AI can be used to identify dynamic change and forecast the user distribution (based on historical data), dynamically optimize the weights of antenna elements, perform adaptive optimization of weights for specific use cases with unique user-distribution and improve the coverage in a multi-cell scenario considering the inter-site interference between multiple 5G massive MIMO cell sites [23].

Additionally, the implementation of massive MIMO implies the installation of many antennas in the base station antenna array, which requires many power amplifiers (PA). The primary problem with PAs is the existing trade-off between their capacity of nonlinearity tracking, predistortion, and impairment correction, and their electrical efficiency. Highest PA efficiency is achieved when constantly feeding the PA at the limit of its highest-power linear region. This is not a feasible solution for high peak-to-average power ratio (PAPR) signals and is not realistic for 5G BSs. Instead, signal-processing-based solutions (ML solutions) are used to provide a better cost-performance trade-off.

### **2.2.7 Encoding**

In [24], the enhanced structure of a Deep Neural Network (DNN) based encoder and decoder was developed for orthogonal frequency-division multiplexing (OFDM) systems and a variation of the autoencoder structure was applied to build a codec for sparse code multiple access (SCMA) systems. There exists a

theoretical possibility of using DNN-based communications systems, which can adapt their operation according to the surrounding environment, including, dynamic channel, mobility, power, and so on. The operation of DNN-based communications systems in practice, especially with respect to their system architectures and the implementation of hardware capable of real-time operation, has not yet been verified.

From the point of view of computational cost, building an encoder and a decoder with a DNN may be challenging because it requires a large number of arithmetic computations (i.e., multiplications and additions) for matrix operations in DNN. Therefore, in order to accelerate operations in the DNN-based encoder and decoder, there is a clear need for dedicated and pipelined digital circuit designs, in which the processing time can be reduced to support the high data rate.

### **2.2.8 NFV/SDN**

The following notes are a synthesis of the survey in [25]. NFV allows customers to transfer the networking functions from vendor-specific and proprietary hardware appliances to software hosted on COTS platforms [26].

Thinking of today, we have several kinds of VMs running on the same type of servers within the datacenter, and we still have routers and switches in different specific boxes. In the NFV context, a new generation of servers with network functionalities designed at the hardware level including also the application layer will be capable of running not only services inside VMs but also network functions such as firewall, routing and switching. This has a tremendous advantage in scaling.

SDN is a new paradigm that was designed to overcome the difficulty in developing and testing new solutions and protocols in production environments, where the underlying code running in business switches and routers are proprietary and closed [27].

The main feature of the SDN paradigm is the separation of the control and data planes. Centralization of the control plane in conjunction with the availability of open APIs, making easy the process of creating and deploying new network configuration and management. This is a useful and powerful abstraction that has as much implications as existing applications and protocols. Maybe one of the main challenges is not technical but strategic: how can two big companies which compete for the same market be ready to sit down, discuss and write code about common APIs?

Both concepts have not only new technical paradigms which are required by the community, they also have consequences at commercial strategy level of big companies. And both concepts are the first steps in order to apply AI models and algorithms in the broad sense inside the networks.

### **2.2.9 Wireless and Mobile Networks**

Improved data rate, low latency and increased capacity for consistent QoS/QoE are the main drivers of 5G networks. The final standard will be published by the ITU in mid-2020, which is also referenced as International Mobile Telecommunications (IMT)-2020. The 5G Tech Forum was created with the participation of the major vendors such as Verizon, Cisco, Ericsson, Nokia and Apple in order to develop early 5G specifications. It is very important to follow what these people say, and how it can relate to the AI environment, thinking not only on applications but also in network algorithms running on the mobile device.

5G and the Cable Industry agree to provide features that support different types of vertical businesses such as IoT, Automotive, Health care, VR&AR, IPTV or Media & Entertainment, and for those applications the



requirements are very different: while automotive industry will need very low latency, IPTV is also currently requiring more bandwidth, and so on.

The access networks are currently being transformed into a new convergent scenario: the converged access transport network (CATN), which blurs the boundaries between access and transport networks driven mainly by the need for bandwidth on existing networks. Radio Access Network (RAN) also must evolve supporting the coexistence of different radio access technologies such as LTE and Wi-Fi. The challenges about data processing are being covered by techniques such as Mobile Edge Computing (MEC).

The use of NFV and SDN technologies will play a significant role in 5G networks, since they allow the network programmability and the fast delivery of new services, enabling network slicing and MEC implementation and orchestration. So, imagine they are running, how we can monitor or operate this level of infrastructure? It is evident that the AI platforms have nowadays a list of requirements to be guaranteed in the next few years.

### **2.2.10 Network Slicing**

Network Slicing refers to the partitioning of a certain physical infrastructure, composed of both network and computational resources, into multiple logical networks, called network slices. This approach has several advantages over the traditional physical networks, such as customization of logical networks according to service requirements, on-demand provisioning to scale resources up or down as conditions change and network resource isolation for improved security and reliability.

The strategy to keep in mind here with the AI approach is to be bendable: more flexible in terms of deployments and in case system fails, we can accept a partial degradation of the services, but we don't want a cut. So different ML algorithms to monitor, maintain the operators reported and make decisions in real time will be needed. Maybe this is a very important point in terms of security and audition of the AI approach. How can a system be audited if a wrong decision is taken by AI?

### **2.2.11 Mobile Edge Computing**

MEC or Multi-access Edge computing has been a trend in mobile networks. The MEC architecture has been standardized by ETSI since 2016, providing IT and Cloud Computing capabilities within RAN. For this, a set of computer and storage resources are deployed at the edges of a mobile operator's network to assist the core data center in supporting computing and communication. It focuses on delivering the services closest to the user to meet certain critical application requirements that are not supported only by Cloud Computing, such as high bandwidth, low latency and jitter, context awareness, and mobility support.

There are studies that deal with Distributed NFV and multiple VIMs (Virtualized Infrastructure Managers) which goes in the complementary approach of the central control plane proposed by SDN. But there are interesting works with these D-NFV; for example, it includes a Virtual Network Life Cycle Manager (VNLM). This element implements a multi-objective resource scheduling algorithm that uses a genetic algorithm to provide near-optimal placement of VNFs over different data centers.

When we consider the use of multiple VIMs without distributed NFV, we usually have a scenario where there are two secondary VIMs, one to manage a data center for virtualization purposes and another to manage a transport network for end-to-end network services provisioning.

### **2.2.12 Optimizing Customer Experience on Video using Machine Learning**

Telecom has its own IPTV and Streaming platform, called FLOW, in order to deliver the best entertainment service to its subscribers, to increase market penetration and to gain competitive advantage. This deployment is based on unmanaged (second screens) and also on managed devices (set top boxes).

To implement our ABR streaming services we need to use different types of video encodings. Encoding is a multi-layer matter of concern. While in DOCSIS we have encoding at the MAC layer, we do not have to forget that the success of video streaming (QoE) depends on the adequate trade-offs between encoding and available bandwidth. In the industry, objective and subjective models for measuring video quality are well known. Recently, objective models emerged using ML algorithms which are trained using databases with subjective evaluations, to combine a variety of classical metrics. Classical metrics are much simpler to implement and at a lower cost but, they produce worse results that do not always fit the human perception. On the contrary, the metrics based on AI produce results very close to the subjective opinion of the customers, but they are more complex to implement and provide development opportunities.

Within the objective metrics based on AI, there are two methods that produce similar results: Video Multimethod Assessment Fusion (VMAF) and Video Quality Model with Variable Frame Delay (VQM-VFD). VMAF is an open-source method proposed by Netflix in 2016 and VQM-VFD was standardized by ITU in 2003. Thus, it's very important for us to develop VMAF as a tool for optimizing FLOW customer experience and to equalize video quality with the other existing video platforms.

## **3 Technical and Theoretical Framework**

In this section we comment on the roadmap of AI for Network Operations in Telecom Argentina. We mention some ML algorithms, as well as some others that are non-ML. Both types shape an ecosystem of techniques and technologies that coexist to make the development of data-driven solutions possible. Far from thinking these methods have to be part of a standard, we mention the algorithms for which we have found interesting applications.


### **3.1 Machine Learning**

The following are well-known algorithms, so rather than focusing on the methodology itself, we emphasize on the practical problems, how they are posed in terms of variables, how these methods provide a solution and the practical implications. Some of the applications are thoroughly documented in publications that we, as members of the STEM team, have submitted to conferences and industry events.

- Artificial Neural Networks (ANN)
- Time Series models
- A/B Testing
- Clustering
- K-Nearest Neighbors (KNN)
- Context Aware Recommendation
- Natural Language Processing (NLP)
- Deep Learning techniques

### 3.1.1 Artificial Neural Networks

Maybe one of the most extended approaches in the application of Artificial Neural Networks (ANN) is the convolutional architecture (CNN), which is originally designed for image recognition and extended for natural language processing, speech processing, and computer vision. The main concept behind these structures consists of the application of a series of transformations to a set of independent variables, in order to predict a response or objective variable (Figure 5). As part of Telecom Argentina STEM team, we have used them to classify the outside plant investments of the optical nodes in the HFC access [28]. Additionally, Marketing teams applied them to predict churn.



**Figure 5 - Artificial Neural Network conceptual diagram**

Other architectures were explored for a variety of purposes. For anomaly detection and forecasting of time series, such as CPU load or traffic growth respectively, the Long-Short Term Memory (LSTM) architecture has been tested effective [29].

### 3.1.2 Time Series models

Time series are often used to forecast the global traffic growth in the industry. Although it seems there is a consensus in the industry about the use of linear regression as the forecast method even when the hypothesis of independency between observations is not met, advanced forecast methods such as ARIMA, exponential smoothing or Kalman filter are, in our experience, more accurate and they have been proved to be useful to explain auto-dependence, stochastic trends and seasonality [30].

### 3.1.3 A/B Testing

A/B Testing is often a rigorous way to determine the results of a change in productive environments. It is an empirical approach in which an experiment is conducted considering two groups of subjects: one is the control and the other, the test. By definition, the groups are mutually exclusive. A certain change or alteration is introduced to the test subjects, so its consequences are measured within the test group and compared to control (which was not subject to changes). The results are driven by the actual data, ending with subjectivities. An application example is testing two different posters for a given TV channel, which are evaluated to determine the the best promoter. A/B Testing will be one of the fundamental keys in the operations toolbox of the future networks, because it allows to test network changes in production environment at lower costs than current migrations.

### **3.1.4 Clustering**

Clustering techniques aim to assign analysis units to classes, according to the appearance of similar characteristics among them. In other words, for a dataset where we have observations and features, cluster analysis is used to find groups such that the observations within the same group are close in the given features space. There exist various definitions of distance, being applied most frequently the Euclidean and Manhattan distances or other correlation-based statistics such as Pearson, Kendall and Spearman correlation coefficients.

There are three main types of algorithms: hierarchical, density-based and non-hierarchical. Hierarchical clustering are useful to find groups and sub-groups within them. Density-based methods do not assume that every observation belongs in a cluster, instead they classify some of the cases as clusters and some others, as 'noise'. Non-hierarchical methods are particularly useful when a priori we can assume that clusters have similar size. In our experience, the latter have been useful to find subscribers with common trends in geographical location and service tier acquired. Another application we found for the same type of methods is to classify cable modem signals, according to the appearance of impairments. This is helping in the development of further capacities for PNM [31].

### **3.1.5 K-Nearest Neighbors**

K-Nearest Neighbors (KNN) is a supervised algorithm that can be used for classification as well as for regression problems. Provided a certain objective variable and set of  $p$  features, which determine a  $p$ -dimensional feature space, we know the location of the points from a training sample and their response values. When a new observation comes along, and we want to predict its response value, the strategy is to look at the response of the  $K$  closest observations in the  $p$ -dimensional feature space, and use them to infer on the new one. The method presents the advantage of being non-parametric and robust. In counterpart, it is computationally expensive to search for the nearest neighbors for all observations in the sample. At Telecom Argentina, this is used to recommend VoD contents to subscribers, based on what other similar users who have watched in the past.

### **3.1.6 Natural Language Processing**

The processing of natural language, also known as NLP, is a branch of computer science associated to the traditional AI definition, which aims to give computers the capacity to understand the meaning of human language. There are efforts directed towards analyzing language according to formal definitions and relationships, executing grammatical, semantical and syntactical analysis. On the other hand, the statistical approach, which consists on extracting words statistics, has been much more successful. One of the reasons for this is when an analysis tries to go deeper by considering relationships between entities and implications to the real world, it loses potential to generalize to other application areas.

Sentiment analysis is one of the most common applications of NLP based on social networks comments in order to evaluate the social experience of the services. Apart from that case, we made some work analyzing the descriptions and comments of the people who monitor the network (like the NOC), to check if the failures or disruptions of the network are being managed uniformly and how they appeared repeatedly along time.

### **3.1.7 Deep Learning techniques**

These notes are based on the article [32]. Deep learning uses multiple layers to represent the abstractions of data to build computational models. Deep learning, which has its roots from conventional neural

networks, significantly outperforms its predecessors. It utilizes graph technologies with transformations among neurons to develop many-layered learning models.

Feature engineering focuses on building features from raw data and is often very domain specific and requires significant human effort. Some very well-known features proposed to compare are Histogram of Oriented Gradients (HOG), Scale Invariant Feature Transform (SIFT), and Bag of Words (BoW).

Deep learning algorithms perform feature extraction in an automated way, which allows researchers to extract discriminative features with minimal domain knowledge and human effort.

Maybe the most prominent current features about Deep Learning techniques are the possibilities given by the capacity of parallel and distributed computing techniques which makes viable lot of implementations.

Threshold logic is the combination of algorithms, mathematics and architectures which allows to emulate the process of thinking of humans, but not to learn in the human sense. The perceptron is the first device within the context of cognition systems, and we all know about how lots of perceptron interconnected conforms a neural network.

The backpropagation algorithm uses the errors in training deep learning models and is one of the first milestones in the creation of neural network architectures. Then the Recurrent Neural Networks (RNN), Recursive Neural Networks (RvNN) such as the Long-Short Term Memory RvNN (LSTM), Deep Neural Networks (DNN), Deep Belief Networks (DBF), Restricted Boltzmann Machines (RBMs) and others are some well-studied architectures.

In each problem we face with Artificial Neural Networks (ANN), we deal with architectures, hyperparameters, optimization algorithms, loss functions, and most important, the nature of the data domain. The mature of this technology is in direct relation with the performance achieved using the pipeline process of the GPUs for processing data vectors instead of finishing in pixels, and since that milestone, several frameworks such as TensorFlow, Torch, Theano, MXNet and others were developed and currently have a big community of developers.

### **3.1.7.1 Deep Learning in Distributed Systems**

There are two main approaches to train models in a distributed system: data parallelism and model parallelism. In the former the model is replicated to all the computational nodes and each model is trained with the assigned subset of data; in the latter all the data is processed with one model where each node is responsible for the partial estimation of the parameters in the model.

With parameter averaging we have for example N slave nodes and one master node, and at time t the weight on the master node is  $W_t$ , then:

$$W_{t+1} = \frac{1}{N} \sum_{i=1}^N W_{t+1,i}$$

is the weight at time t+1. Very often, the objectives of a ML algorithm are optimized using the update:

$$w \leftarrow w - \alpha \sum_{i=1}^n g(w; x_i, y_i)$$

where  $w$  is a vector of dimension  $d$  and the data has a length of  $n$ ,  $x_i$  and  $y_i$  are the dimensions or hyperparameters of the model and  $\alpha$  is the weight.

These simple notations are placed here to mention performance trade-offs between memory and computing cost and new compromises driven by the applications.

If we have MEC the ML models must be solved within the capacity of the edge, so only local data parallelism is admitted. But if we have a centralized computation cluster, then we can apply global models and data parallelism for the whole network, but at the price of latency for applications running far away from the datacenter.

So, the implementation of AI solutions on future networks will have to address this compromise too. MEC is a powerful framework to focus on the models and applications very near the user. This will also require the need for monitors and fail recoveries schemes when, for example, some application is compromising the capacity of the local cluster.

## **3.2 Non-Machine Learning Tools**

As we proceeded with traditional ML algorithms, next we describe some other tools that are not ML, since their implementation does not require tuning parameters, optimization of functions or training models. However, they complement the ML implementations, in some cases they provide contextual information and in some others, they are useful to execute some calculations with reasonable timing.

### **3.2.1 Locality-Sensitive Hashing**

Locality-sensitive hashing (LSH) facilitates the identification of similar observations in a high volume of data. The general idea is to hash data points into buckets so that data points near each other are located in the same buckets with high probability, while data points far from each other are likely to be in different buckets. While the traditional process of looking for pairs of equals has a quadratic computational cost,  $O(n^2)$ , this alternative looks for pairs of similar items without the cost of examining each possible pair, hence a much lower cost. It is particularly useful to detect near-duplicate documents, webpages and other types of files, for large-scale image search and audio/video fingerprinting (A/V fingerprinting) [33]. Fingerprinting enables the identification of characteristics from multimedia, so we expect to use this algorithm to extract features from contents in Telecom Argentina's Flow platform.

### **3.2.2 Collaborative filters**

Collaborative filtering is used for content recommendation. It is based on previous subscriber behavior, and it evaluates similarities between subscribers in terms of the viewed content. An index score for each content is defined with the measurements of each user and then is multiplied by the similarities distance between users. The similarity analysis is usually combined with the results of a clustering algorithm to improve recommendations [34]. In the case of Telecom Argentina, collaborative filtering is used as a complement to KNN clusters to manage recommendations of VoD content.

### **3.2.3 Process Mining**

When we execute an analysis, we intend to answer the questions what happened, why it happened, what is likely to happen next and what is the best that could happen next. The analysis of processes allows us to understand if a certain process model that is assumed is actually being executed, as well as discover new process models underlying the data. Sequences and relationships between possible events are represented with Petri nets, where nodes represent the events, directed arcs indicate the pre-conditions and post-

conditions, and they can be enriched with information about time between events or frequency of appearance [35].

In a digital ecosystem, where every device produces event data (in the form of data logs), we expect that process mining (PM) will provide an alternative to understand subscriber behavior and compare the expected usage patterns to actual ones.

### 3.2.4 Digital Twins


Digital twins are realistic representations of real-world entities, which consist of a combination of mathematical and computational methods, and software services that provide real time synchronization with the real object or process they represent. Their purpose is to judge, analyze, predict and optimize the real entity. They are usually supported by 3D, video, augmented reality and/or virtual reality representations.

The twins are used in manufacturing to reduce errors, reduce planning time, plan for changes, hence reduce cost of changes, and increase the planning maturity [36]. Another field of application is smart cities, as it is the case of the digital twin created for the city of Atlanta [37], to plan for resource allocation, provide security, maximize services, facilitate human activities and prevent disruption while continuously adapting. We expect that the same approach can contribute to improve network capacity planning, engineering and architecture.

## 3.3 The Knowledge Plane

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* [38].

The knowledge plane (Figure 6) 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.

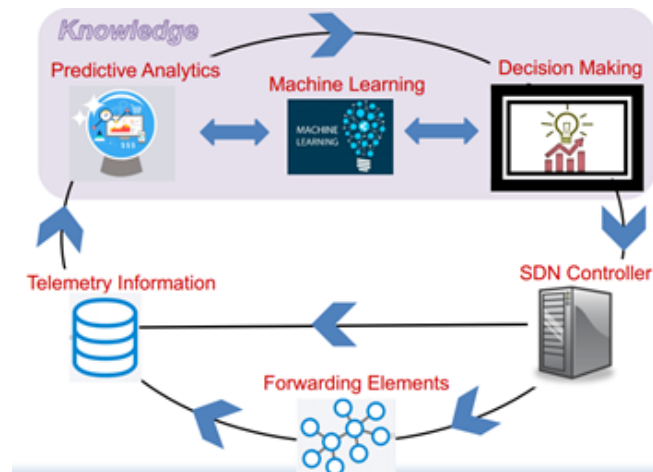


**Figure 6 - The four planes in network architecture.**  
Source: TM Forum.

Several different KP based approaches have been proposed [39]. But it is not until the development of NFV and SDN that such proposal once again takes hold in communities such as the TM Forum, IETF and the Industry.

In [40] 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. The KDN paradigm is also taken by the TM Forum as a proposal to specify future architectures [41].

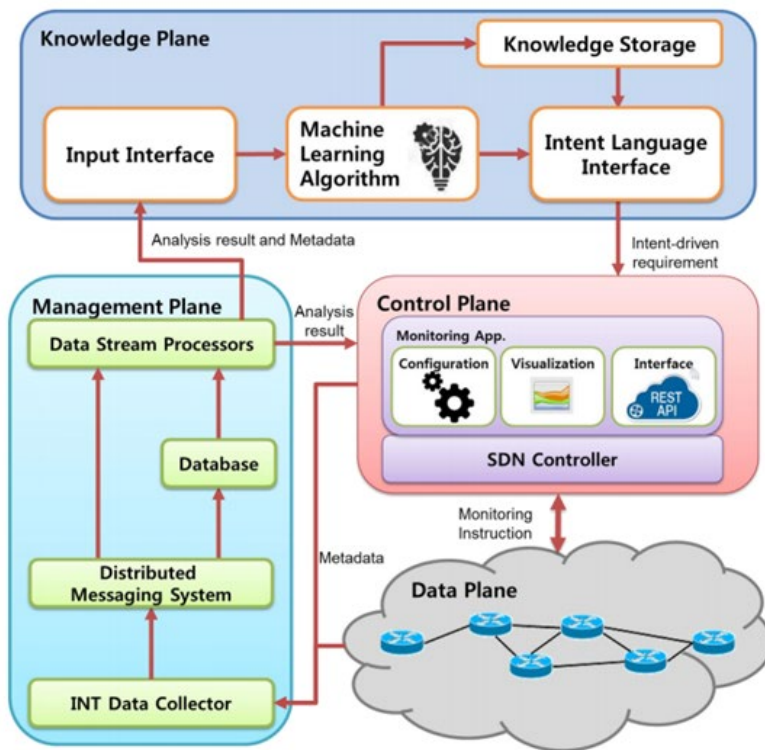
An example of the plane of knowledge in an SDN is presented in Figure 7 [42].



**Figure 7 - Plane of Knowledge in a SDN**

Finally, Figure 8 shows an architecture with more detail presented in APNOMS [43].



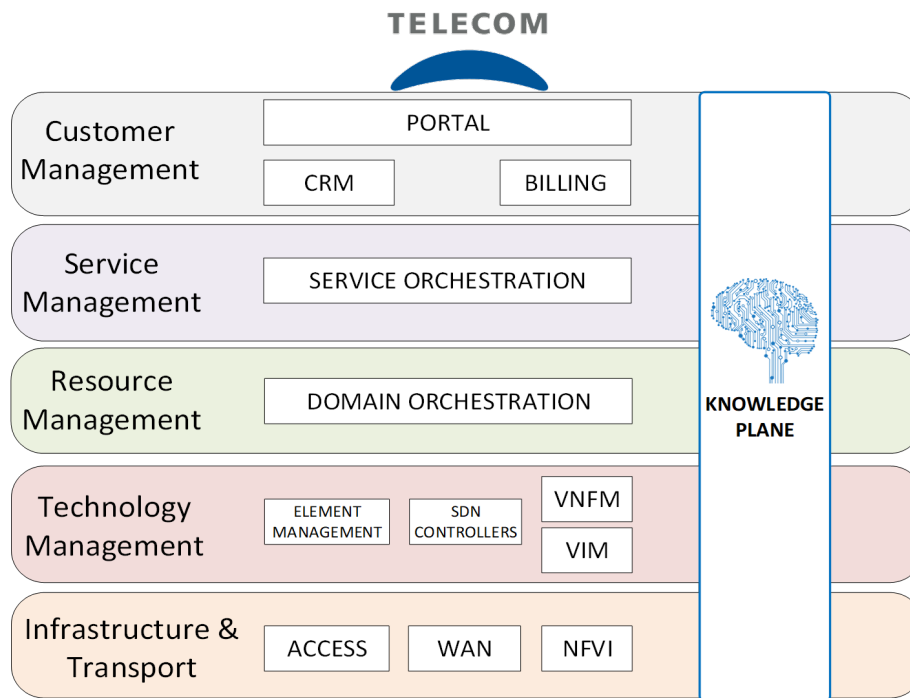


## 4 Strategy

More than a decade has passed since the emergence of a paradigm of autonomous computing in the world of telecommunications. Back then, there was a gap between that paradigm and the capacities of the networks. However, a path has been taken in recent years with the adoption of cloud computing, NFV and SDN.

These technological advances have made available a more agile infrastructure, computing capacity and storage as resources more abundant than ever. Motivated by this evolution, together with the ever-growing need to improve the management and administration of networks and services, we present this first approach to the plane of KP.

We define a reference architecture (Figure 9) with the purpose of automating the services from end to end, with a holistic view of the network and towards a CATN. It is based on RFC 8309, MEF-55, eTOM model (TM Forum) and other initiatives of other CSPs and vendors.



**Figure 9 - Telecom Argentina knowledge-defined network**

We detail some of the main topics around AI and networks. The first steps Telecom Argentina is giving towards the deployment of AI for networks and services are:

- Operations use cases of Deep Learning

The main benefits of the deployment of AI in the context of networks will be achieved by the operations teams. This is often the main starter point to think about a deployment of ML. One of the well-known use cases is anomaly detection.

- NFV/SDN orchestration layer

We consider SDN orchestration layer as the basis for all the AI running applications. But this does not mean we will wait for it before starting with AI.

- Extension of AI applications through integration platforms

Today we are used to having several OSS/BSS systems that are not always interconnected within a common platform. In this sense, integration platforms are key technologies which enable the combination of different data domains and kinds of algorithms through common APIs. Some relevant work about integration platforms with native parallelism has been made by the people of The Apache Software Foundation, particularly with Spark.

- Highly specialized human resources

The structure of the organization is another key concept which cannot be excluded by the AI framework. The knowledge layer is a concept in the SDN paradigm with relevant meaning in the ability of the specialized teams to deeply understand about topics of networks, machine learning and business.

- Distributed Data Lakes and centralized Data Ocean

There are a lot of data domains available in the current deployed networks, but this does not mean all the data domains have to be in a centralized datacenter. There are centralized and distributed approaches for the administration and mining of the data, and a lot of effort to build a data ocean. But as the capillarity of the networks is still growing and merging with the mobile access, the MEC could evolve as a network of distributed data lakes.

- TM Forum framework for application and analytics layers

Hybrid layer models considering the enterprise organizations have been designed consider mainly TMForum for the application and analytics layers and ETSI/IETF for the Physical and MAC layers.

- Open source and microservices-based platforms

Open source is useful at the development and test stages, and sometimes with the adequate support is the right tool for production environments. Vendors are taking that direction to make better products and to guarantee a dynamic evolution and technical support.

- Data Governance

Data governance is still developing the interactions and processes within the organization. The main capabilities are to manage data quality, security and compliance, storage, presentation and distribution of the data within and outside the enterprise. It also has to meet compliance requirements dictated by the organization policies and external regulatory entities [44].

## Conclusions

The Applied Artificial Intelligence STEM Group will focus on the application of AI to decision-making process and auto-remediation to help Telecom Argentina's network keep pace with the growth in network size, traffic volume and service complexity, as well as define new approaches to network operations and customer assurance to support the accelerated deployment of new over-the-top services and collaborate in the digital transformation process, in which Telecom Argentina goes from being an MSO & MNO to a Digital Service Provider.

This work is the first step in the Telecom Argentina's roadmap for AI applied to networks and services. It is a long, challenging journey, however, we are confident that by taking one step at the time, little but concrete, we will accomplish our goals.

The first stage we are working on is to define this framework regarding to the deployments of AI systems in the whole network. More than a year ago we understood that there will be no future networks without AI and we decided to make our first proof of concepts based on our network and services data. We already know we have to continue the direction of this decision but this does not mean we have to stop testing models and algorithms in our labs. Some methodologies like Agile or DevOps are very useful in order to find fast results.

The trainings are some of the most important skills the knowledge layer must develop. Customized trainings based on local data and exposed by local professionals are a good practice.

Working groups such as CableLabs/SCTE, TM Forum and Telecom Infra Project are collaborative environments the industry may maintain, and the company will continue to participate.

Regarding the vendor relations, in this stage we are identifying how they deploy the knowledge layer and how is promoted. We already define a convergent scenario for the merging of the access networks called Converged Access Transport Network (CATN), so the vendor's offered knowledge layer has to match with our requirements. This is a proposal to invite you to follow our path.

## Abbreviations

ABR	adaptive bit rate
AI	Artificial Intelligence
ANN	Artificial Neural Network
AR	augmented reality
ARIMA	autoregressive integrated moving average
Avg	average
B2B	business to business
B2C	business to consumer
BS	base station
BW	bandwidth
CAGR	Compound Annual Growth Rate
CM	Cable Modem
CMTS	Cable Modem Termination System
CVA	Cablevisión S.A.
DNN	Deep Neural Network
DOCSIS	Data Over Cable Service Interface Specification
DSP	Digital Service Provider
EDA	Exploratory Data Analysis
GHz	Giga Hertz
HFC	Hybrid Fiber Coaxial
HHP	household passed
IoT	Internet of things
IPTV	Internet Protocol Television
ISBE	International Society of Broadband Experts
Kbps	kilobits per second
Km	kilometers
KPI	key performance indicator
LTE	long term evolution
Mbps	megabits per second
MEC	multi-access edge computing
MIMO	multiple input, multiple output
ML	Machine Learning
mMIMO	Massive MIMO
MNO	Mobile Network Operator
MSO	Multiple Service Operator
NFV	network functions virtualization
NR	new radio

OPS	operations per second
O-RAN	Open Radio Access Network
OSS/BSS	Operation Support System/Business Support System
PCA	Principal Components Analysis
PNM	Proactive Network Maintenance
QAM	Quadrature Amplitude Modulation
QoE	quality of experience
QoS	quality of service
RAN	radio access network
ReLU	Rectified Linear Unit
SCTE	Society of Cable Telecommunications Engineers
SD	standard deviation
SDN	software-defined network
SG	service group
SoC	system on a chip
SON	self-organizing network
STEM	science, technology, engineering and mathematics
Subs	subscribers
TCO	total cost of ownership
UE	user equipment
URLL	ultra-reliable low-latency
URLLC	ultra-reliable low-latency communications
VR	virtual reality

## Bibliography & References

- [1] «Digital Service Providers (DSP),» National Cyber Security Centre (NCSC), 2019. [En línea]. Available: <https://www.ncsc.gov.ie/dsp/>.
- [2] J. Kyriakakis, «Defining the Digital Service Provider,» LightReading, 29 September 2014. [En línea]. Available: <https://www.lightreading.com/business-employment/business-transformation/defining-the-digital-service-provider/a/d-id/711114>.
- [3] J. Bryson y A. Winfield, «Standardizing Ethical Design for Artificial Intelligence and Autonomous Systems,» *IEEE Computer Society*, vol. 4, n° 2, pp. 10-13, 2018.
- [4] A. M. Turing, «Computing Machinery and Intelligence,» *Mind*, vol. 59, n° 236, pp. 433-460, 1950.
- [5] J. McCarthy, «What Is Artificial Intelligence? - Stanford University,» 12 November 2007. [En línea]. Available: <http://www-formal.stanford.edu/jmc/whatisai/node1.html>.
- [6] S. Russell y P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd Ed., Prentice Hall Series in Artificial Intelligence, 2010.

- [7] A. L. Samuel, «Some Studies in Machine Learning Using the Game of Checkers,» *IBM Journal of Research and Development*, vol. 3, nº 3, pp. 210-229, 1959.
- [8] B. Marr, «How Much Data Do We Create Every Day? The Mind-Blowing Stats Everyone Should Read,» *Forbes*, 21 May 2018. [En línea]. Available: <https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/#2aff0a3960ba>.
- [9] K. Schwab, *The Fourth Industrial Revolution*, Penguin Random House, 2017.
- [10] C. Righetti, *Personal notes on TM Forum Digital Transformation*, Nice, France, 2018.
- [11] K. Sundaresan, N. Metts, G. White y A. Cabellos-Aparicio, «Applications of Machine Learning in Cable Access Networks,» de *SPRING TECHNICAL FORUM, CableLabs SCTE NCTA 2016 Spring Technical Forum Proceedings*.
- [12] G. White y K. Sundaresan, «DOCSIS 3.1 Profile Management Application and Algorithms,» de *SCTE NCTA 2016 Spring Technical Forum Proceedings*.
- [13] H. Hagras, «Toward Human-Understandable, Explainable AI,» *IEEE Computer*, vol. 51, nº 9, pp. 28-36, 2018.
- [14] B. S. I. (BSI), British Standard BS 8611:2016, Robots and robotic devices. Guide to the ethical design and application of robots and robotic systems., London, UK, 2016.
- [15] U. Gasser y V. A. F. Almeida, «A Layered Model For AI Governance,» *IEEE Internet Computing*, vol. 21, pp. 58-62, 2017.
- [16] T. Forum, «IG1184 Service Management Standards for AI R18.5.1,» 2019.
- [17] M. Yao, M. Sohul, V. Marojevic y J. H. Reed, «Artificial Intelligence Defined 5G Radio Access Networks,» *IEEE Communications*, vol. 57, nº 3, pp. 14-21, 2019.
- [18] T. T. T. Nguyen y G. Armitage, «A Survey of Techniques for Internet Traffic Classification Using Machine Learning,» *IEEE Communications Surveys & Tutorials*, vol. 10, nº 4, pp. 56-76, 2008.
- [19] Z. Jiang, S. Chen, A. F. Molisch, R. Vannithamby, S. Zhou y Z. Niu, «Exploiting Wireless Channel State Information Structures Beyond Linear Correlations: A Deep Learning Approach,» *IEEE Communications Magazine*, vol. 57, nº 3, pp. 28-34, 2019.
- [20] L. Wolcott, M. O'Dell, P. Kuykendall, V. Gopal, J. Woodrich y N. Pinckernell, «A PNM System Using Artificial Intelligence, HFC Network Impairment, Atmospheric and Weather Data to Predict HFC Network Degradation and Avert Customer Impact,» de *SCTE/ISBE 2018 Fall Technical Forum*, Atlanta, GA, 2018.
- [21] M. Feng y S. Mao, «Dealing With Limited Backhaul Capacity in Milimeter-Wave Systems: A Deep Reinforcement Learning Approach,» *IEEE Communications Magazine*, vol. 57, nº 3, pp. 50-55, 2019.

- [22] Intel, «Learn about Multiple-Input Multiple-Output,» 25 March 2019. [En línea]. Available: <https://www.intel.com/content/www/us/en/support/articles/000005714/network-and-i-o/wireless-networking.html>.
- [23] O. Dharmadhikari, «Leveraging Machine Learning and Artificial Intelligence for 5G,» Informed blog by CableLabs, 18 June 2019. [En línea]. Available: <https://www.cablelabs.com/leveraging-machine-learning-and-artificial-intelligence-for-5g>.
- [24] M. Kim, W. Lee, J. Yoon y O. Jo, «Toward the Realization of Encoder and Decoder Using Deep Neural Networks,» *IEEE Communications Magazine*, vol. 57, n° 5, pp. 57-63, 2019.
- [25] M. S. Bonfim, K. L. Dias y S. F. L. Fernandes, «Integrated NFV/SDN Architectures: A Systematic Literature Review,» *ACM Computing Surveys*, vol. 51, n° 6, 2019.
- [26] European Telecommunications Standards Institute (ETSI), «Network Functions Virtualisation - An Introduction, Benefits, Enablers, Challenges and Call For Action,» de *SDN and OpenFlow World Congress*, Darmstadt, Germany, 2012.
- [27] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker y J. Turner, «OpenFlow: Enabling Innovation in Campus Networks,» *ACM SIGCOMM Computer Communication Review*, vol. 38, n° 2, pp. 69-74, 2008.
- [28] C. Righetti, E. Gibellini, F. De Arca, C. G. Carreño Romano, M. Fiorenzo, G. Carro y F. R. Ochoa, «Network Capacity and Machine Learning,» de *SCTE•ISBE Cable-Tec Expo 2017*, Denver, CO, 2017.
- [29] C. G. Carreño Romano y N. Clivio, «Sizing Techniques Applied to Network Capacity Planning,» de *IEEE Biennial Congress of Argentina (ARGENCON)*, Tucumán, Argentina, 2018.
- [30] R. Chrobok, «Theory and Application of Advanced Traffic Forecast Methods,» Duisburg-Essen, Germany, 2005.
- [31] Gibellini, E.; Righetti, C., «Unsupervised Learning For Detection Of Leakage From The HFC Network,» de *ITU Kaleidoscope 2018: Machine Learning for a 5G Future*, Santa Fe, Argentina, 2018.
- [32] S. Pouyanfar, S. Sadiq, Y. Yan, H. Tian, Y. Tao, M. Presa Reyes, M. L. C. S. C. Shyu y S. S. Iyengar, «A Survey on Deep Learning: Algorithms, Techniques and Applications,» *ACM Computing Surveys*, vol. 51, n° 5, p. Article 92, 2018.
- [33] S. Gupta, «Locality Sensitive Hashing - Towards Data Science,» Towards Data Science, 29 June 2018. [En línea]. Available: <https://towardsdatascience.com/understanding-locality-sensitive-hashing-49f6d1f6134>.
- [34] J. Leskovec, A. Rajaraman y J. D. Ullman, *Mining of Massive Datasets*, Cambridge University Press, 2010.

- [35] W. Van der Aalst, *Process Mining*, Berlin, Germany: Springer-Verlag Berlin Heidelberg, 2016.
- [36] F. Biesinger, D. Meike, B. Kraß y M. Weyrich, «A Case Study for a Digital Twin of Body-in-White Production Systems,» de *IEEE 23rd International Conference on Emerging Technologies and Factory Automation (ETFA)*, Torino, Italy, 2018.
- [37] N. Mohammadi y J. E. Taylor, «Smart City Digital Twins,» de *IEEE Symposium Series on Computational Intelligence (SSCI)*, Atlanta, GA, 2017.
- [38] D. D. Clark, C. Partridge, J. C. Ramming y J. T. Wroclawski, «A Knowledge Plane For The Internet,» de *Conference on Applications, Technologies, Architectures and Protocols for Computer Communications (SIGCOMM)*, New York, NY, 2003.
- [39] K. R. Sollins, «An Architecture for Network Management,» de *Workshop on Re-Architecting the Internet (ReArch)*, New York, NY, 2009.
- [40] A. Mestres, A. Rodriguez-Natal, J. Carner, P. Barlet-Ros, E. Alarcón, M. Solé, V. Muntés-Mulero, D. Meyer, S. Barkai, M. J. Hibbett, G. Estrada, K. Ma'ruf, F. Coras, V. Ermagan, H. Latapie y C. Cassar, «Knowledge-Defined Networking,» *SIGCOMM Computer Communications*, vol. 47, nº 3, pp. 2-10, 2017.
- [41] B. Levy y B. Graham, «TM Forum Future Architecture Strategy,» 2017. [En línea]. Available: [https://www.tmforum.org/wp-content/uploads/2017/09/TM-FORUM-FUTURE-ARCHITECTURE-STRATEGY-v5\\_final.pdf](https://www.tmforum.org/wp-content/uploads/2017/09/TM-FORUM-FUTURE-ARCHITECTURE-STRATEGY-v5_final.pdf).
- [42] Z. Zhu, «Knowledge-Defined Network Orchestration in a Hybrid Optical/Electrical Datacenter Network,» de *Conference on Optical Network Design and Modeling*, Dublin, Ireland, 2018.
- [43] J. Hyun y J. Won-Ki Hong, «Knowledge-Defined Networking Using In-Band Network Telemetry,» de *Asia-Pacific Network Operations and Management Symposium (APNOMS)*, Seoul, Korea, 2017.
- [44] TM Forum, «Frameworks Technical Report, Data Governance Functions and Implementations (TR261 Release 16.0.1),» 2016.



# Customer First: CX-Driven Augmented Operations

A Technical Paper prepared for SCTE•ISBE by

**Roger Brooks, Ph.D.**  
Chief Scientist  
Guavus, a Thales company  
2860 Junction Avenue  
San Jose, CA 95134 USA  
roger.brooks@guavus.com

Pankaj Kumar, Sr. Mgr. Analytics, Guavus

Mudit Jain, Mgr. Analytics, Guavus

Megha Vij, Data Scientist, Guavus

Nandit Jain, Data Scientist, Guavus

Andrew Colby, Field CTO, Guavus

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
CPE versus Network Repair .....	3
1. Use Case .....	3
2. Data .....	5
3. Machine Learning Approach .....	5
4. Results .....	5
5. Challenges Addressed .....	6
6. Related Proof of Concept .....	6
Network Equipment Failure Prediction .....	7
1. Use Case .....	7
2. Data .....	8
3. Machine Learning Approach .....	9
4. Results .....	9
5. Challenges Addressed .....	10
Conclusion .....	10
Abbreviations .....	12
Bibliograph & References .....	12

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Recommending that a subscriber's issue is best addressed at the customer's premises or in the network. ....	4
Figure 2 – The percentage of actual Line Tech Truck Rolls predicted by the ML model run at various levels of confidence.....	6
Figure 3 – The top KPI categories which are predictive of future cable modems care events.....	7
Figure 4 - Predicting equipment failures from syslog data.....	8
Figure 5 – The percentage of actual site interventions predicted by the ML model run at various levels of confidence. ....	9
Figure 6 - Driving profits via machine learning on state data from network devices. ....	11

# Introduction

While service plan price will attract customers, it is their quality of experience (QoE) which will determine whether they churn (Ovum, 2017). Subscribers hold operators responsible for everything, from the content provider to their home, that affects their experience. When asked “What characteristics are important for a high-quality broadband service?” the top two responses were

- “100% reliable broadband service”
- “Good customer service”

When we look at the first of these, survey results (Incognitio, 2016) indicate that subscribers base their decision to recommend their broadband service provider on three factors:

- Service speed: 45%
- WiFi reliability: 31%
- Pricing and service bundling options: 21%

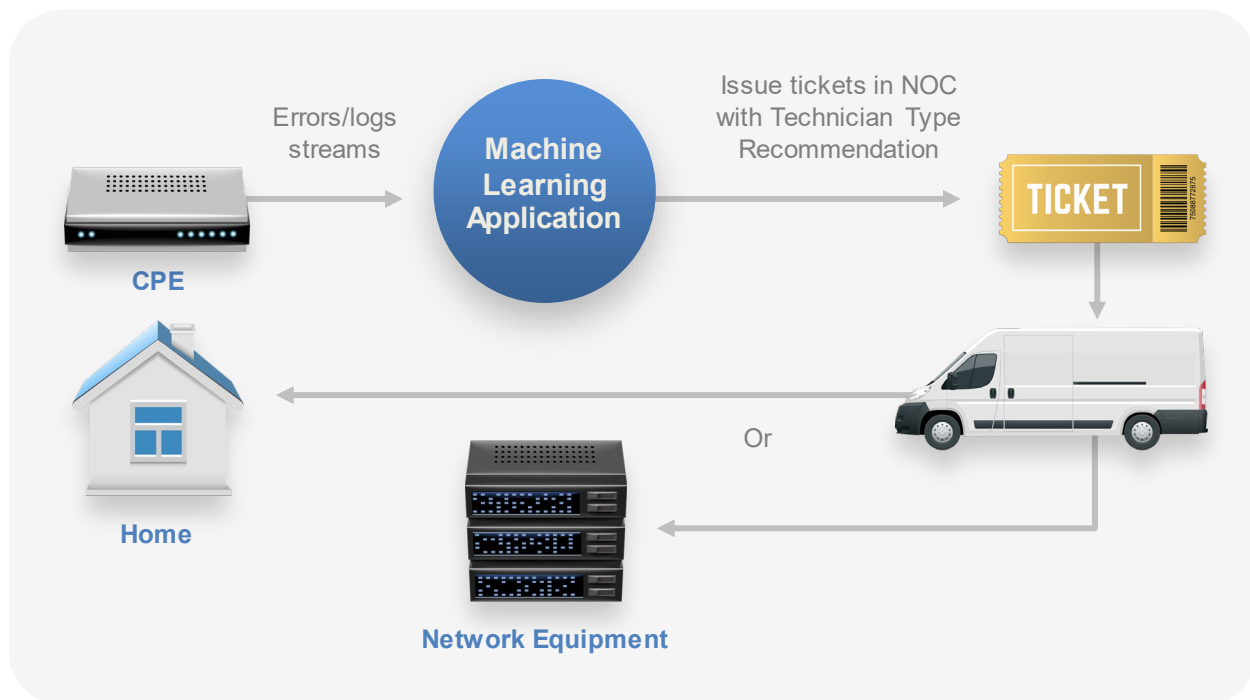
We can imagine numerous use cases to affect these key quality indicators (KQIs) and to which machine learning might be needed. However, appropriate data is not always available to support the use cases as significant portions of those data must come from the CPE and network devices themselves.

In this paper, we report on two examples of how machine learning can leverage data generally available from CPE and network equipment to address technical customer experience use cases.

## CPE versus Network Repair

### 1. Use Case

A typical cable operator schedules hundreds of thousands of truck rolls annually only to find out that the service impairments cannot be resolved by the technician (Field Tech) at the home. This discovery is made after the visit to the home and requires scheduling a second truck roll for a network maintenance technician (Line Tech). Apart from the wasted Opex of the truck roll to the home, this also negatively impacts NPS due to the frustration and delayed satisfaction of the customer. Conversely, network impairments have a potential to impact multiple customers, and delays in identifying these impacts further increase the likelihood of unnecessary truck rolls to customers impacted by the same network impairment and negative customer experiences. At \$60+ per truck roll, solving this one problem equates to savings in the order of \$10sM per year in addition to the impact on NPS and churn. Our goal is to make the prediction at least 24 hours before the Line Tech’s truck roll would have been scheduled (e.g., when the Field Tech is at the home).



**Figure 1 - Recommending that a subscriber's issue is best addressed at the customer's premises or in the network.**

### **Present Mode of Operations (PMO)**

1. Customer calls in seeking resolution of an Issue.
2. Troubleshooting results in a scheduled Field Tech Truck Roll to the home.
3. Field Tech arrives and determines that a Line Tech is needed to completely resolve the issue. (It is possible that BOTH a Field and Line Tech may be needed.)
4. Field Tech submits a request for a Line Tech to be scheduled.
5. A Line Tech is dispatched to diagnose and address the Issue in the HFC/Access network.
6. Attempts are made to verify that all services have been restored.
7. Other customers who are experiencing or will succumb to the same Issue either seek support resulting in additional Field and/or Line Tech Truck Rolls or remain silent with downgraded QoE.

### **Future Mode of Operations**

1. --- Same as 1 of PMO ---
2. The troubleshooting process is informed by an AI prediction that the issue will require a Line Tech Truck Roll to the HFC/Access network.
3. --- Same as 5 of PMO ---
4. --- Same as 6 of PMO ---
5. Other customers who are experiencing or will succumb to the same Issue also find their services are back to normal without the need for additional truck rolls.

**Operations Benefit:** Unnecessary Field Tech Truck Rolls are preempted. Needed Line Tech Truck Rolls can be automatically scheduled. A decrease in support calls, from customers sharing the same network issues, is also expected.

**Business Value:** \$10sM/year are saved in unnecessary Field Tech Truck rolls. Customers no longer experience unnecessary home visits and some issues are discovered and addressed without customers experiencing service impairments. This is expected to reduce churn and help stabilize NPS.

## 2. Data

In addition to data drawn from customer databases and network telemetry common to most MSOs, we leveraged telemetry from the CPE. Thematically, the data encompassed:

- Customer care event data
  - Calls – timestamp, account ID, service type(s), call type, etc.
  - Tickets – timestamp, account ID, problem type, prior resolution types, etc.
- CPE Telemetry
  - timestamp, CPE identifier, issue indicators, etc.
  - timestamp, CPE identifier, network connectivity events, etc.
- CPE Descriptor data
  - Service type, Device Type, Manufacturer, Model, MAC address, etc.


For the results reported, we utilized data from a MSO.

## 3. Machine Learning Approach

We created an analytics solution based on Supervised Machine Learning whose key modelling elements were an advanced clustering and classification on those clusters. To prove the concept, we constructed a training pipeline in which each subscriber's CPE data records were transformed into features reflecting the nature of the issues, the history (seven days prior to the prediction) and context of the device and customer. These features were broken into rolling windows of various time periods. The care event data was reflective of whether a Field Tech alone resolved the issue, or whether a Line Tech was needed, was used as the labels for the model.

## 4. Results

Separate from the data used to train the model, predictions were checked against recorded care events. It was found that roughly 90% of the subscriber accounts predicted to need a Line Tech, did in fact have the predicted truck roll after the interval (24 hours) on which the prediction was made. Despite unavailable (beyond those reflected in the data) factors leading to the scheduling of truck rolls, the 90% confidence predictions from the model captured 41% of Line Tech Truck Rolls for all subscribers who had at least one issue, within 24 - 168 hours, prior to the prediction; see Figure 2.




**Figure 2 – The percentage of actual Line Tech Truck Rolls predicted by the ML model run at various levels of confidence.**

## 5. Challenges Addressed

Data quality and completeness is always a challenge with these real-world projects and particularly so in operations analytics where one is essentially trying to interpret issues which would otherwise appear as statistical noise. The processes by which tickets are generated and updated further complicated the data preparation and analytics design as it led to ambiguous associations between the issues and those tickets. The machine learning had to be robust to these challenges; a task difficult even for human SMEs.

## 6. Related Proof of Concept

In a related proof-of-concept, we developed a similar ML model to predict which subscriber-cable modems were experiencing issues that would, if left unaddressed, result in a care event (call or truck roll). For the particular MSO customer, we used KPI telemetry data from the cable modems and found such predictions if acted upon would result in a seven-digit annual OPEX savings. The top categories of KPI attributes which drove those predictions are depicted in Figure 3.

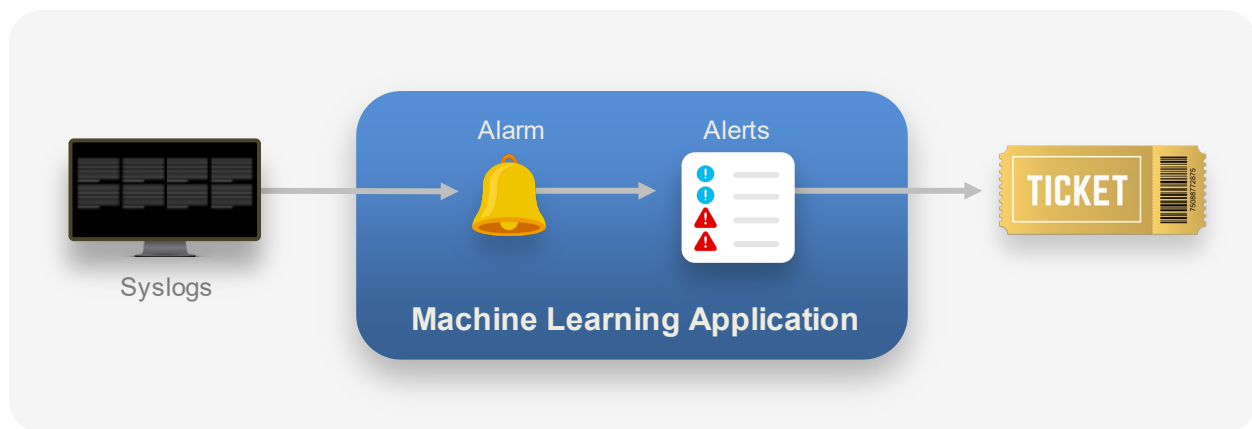


**Figure 3 – The top KPI categories which are predictive of future cable modems care events.**

## Network Equipment Failure Prediction

### 1. Use Case

We seek to predict if equipment critical to the operations of services may fail and why. We know how to apply machine learning to predict incidents from the alarms that are being issued by faulty devices; but what about non-alarm related failures? Logs are ubiquitous so in this use case, we predict service impacting equipment failures from syslog data. In particular, for tens of thousands of network node equipment, we sought failures associated with incidents which do not require on-site actions under the premise that those can be remotely and quickly resolved to take advantage of the advanced notice afforded by the predictions.



**Figure 4 - Predicting equipment failures from syslog data.**

### **Present Mode of Operations**

1. Syslogs are generated by node equipment.
2. Syslogs are automatically collected, supplemented and made available for manual review by the operations team.
3. The Operations team associates ongoing incident tickets with the logs. Due to variations in time of review, interpretations of logs, frequent updates, etc. there is inconsistency in when and how these logs are associated with the corresponding incident tickets.
4. Technicians attempt intervention/non-intervention to resolve the issues. Tickets are closed at a varying numbers of days later.

### **Future Mode of Operations**

1. --- Same as 1 of PMO ---
2. At configurable intervals, the Operations team receives a list of 'At Risk' equipment ids with risk score and risk drivers for the associated incident.
3. Staff takes remote action on 'At Risk' equipment ids to proactively solve issues before the incidents materialize.

**Operations Benefit:** Maintenance action on 'At Risk' equipment can be proactively taken in advance of those equipment leading to an incident.

**Business Value:** Incidents which impair services and hence customers' experiences are prevented. Such adverse contributions to churn and NPS are reduced.

## **2. Data**

The data consisted of system generated equipment syslogs and the incidents created (manually) by the operations team. Thematically, the data encompassed:

- Syslogs – timestamp, equipment ID, log type, log entry (machine generated text)
- Incidents – timestamp, element ID, incident type, on-site flag (includes text fields)



It should be noted that not all incidents are reflected in the log data. Hence, we did not expect to be able to predict all incidents; indeed, we will see this in the results reported below.


For the results reported, we utilized data from a telecom operator.

### 3. Machine Learning Approach

Much as in the prior use case, we created an analytics solution based on Supervised Machine Learning whose key modelling element was an advanced classifier. To prove the concept, we constructed a training pipeline in which the syslog's fields were transformed into features reflecting the lexicals extracted from the text in the log entries along with their incident history. To aid in the system noise reduction, only logs which persisted for greater than a configured time period were used. Furthermore, spectral clustering was used for further reduction. Two binary labels were assigned to each training record; one based on whether an incident occurred in the subsequent 24-hour period and the other depicting whether an intervention was taken. Two different classifiers were trained, one for each label. The first classifier reflected classes of equipment organized by the probability of their members to fail. Predictions were made, at each hour, for each equipment logged during the period 15-75 minutes prior. The output was a daily prediction list for 'At Risk' equipment and whether the nature of the failure likely required an on-site visit or not.

### 4. Results

A daily list of equipment which are likely to result in an incident within the subsequent 3 days was predicted. On average it was found that approximately 4% of the equipment account for about 75% of the total incidents. In particular, we found that approximately 39% of the non-intervention incidents could be predicted at least 12 hours before the incident would have occurred. Furthermore, with just this limited information, the model was also able to predict, at 77% confidence, 31% of the incidents for which no on-site intervention was needed; see Figure 5.



**Figure 5 – The percentage of actual site interventions predicted by the ML model run at various levels of confidence.**

## 5. Challenges Addressed

Familiar operations analytics challenges presented themselves in this project and were part of the reason we adopted the ML approach described above. Noteworthy amongst them were:

- Inconsistency in the manual incident ticketing process.
- Significant and varying time lags between the occurrences of the logs and the supposed incidents.

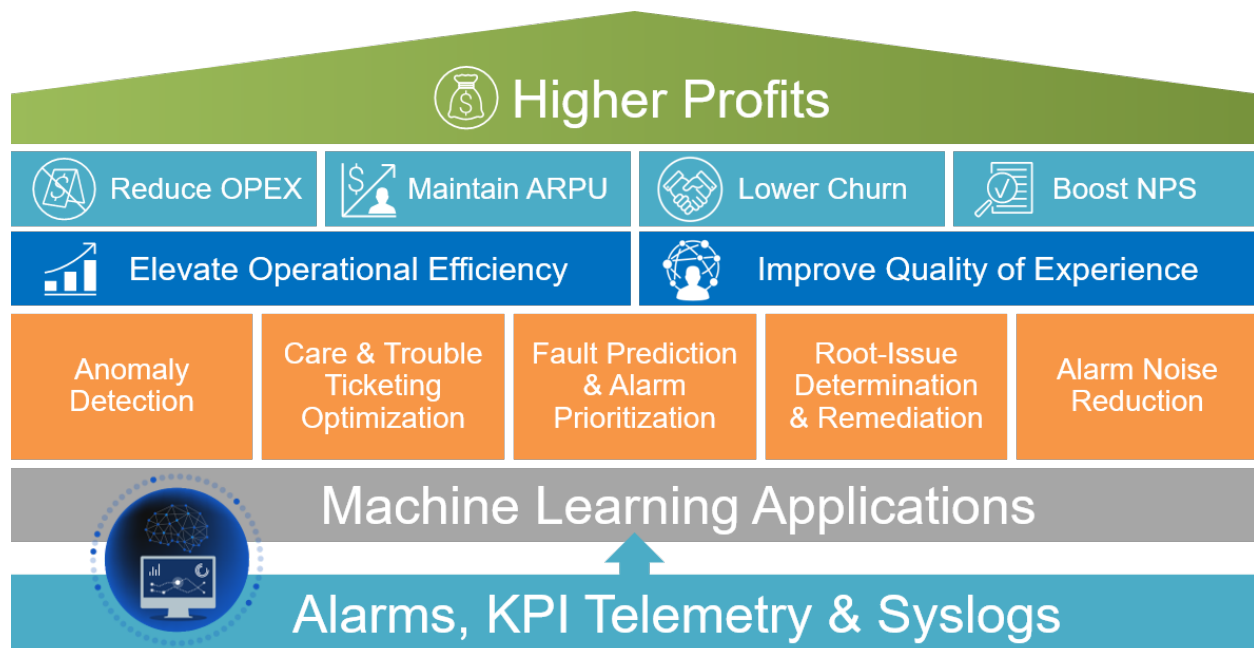
## Conclusion

Studies have shown that the reduction of churn and increase in NPS are driven, in part, by improvements in QoE. We have described two equipment failure use cases which, in our work with customers, we have found to (a) impact customer's experiences and (b) were feasible given their data and operator processes. Given that no two operator's ecosystems are the same, mileage will vary.

Beyond the use cases discussed here, we have also found a number of analytics use cases with QoE impact and for which data is often available, including

- Prediction of incidents from alarms and thus prioritizing to which to attend.
- Facilitating faster service restoration through diagnostic steps such as discovering the root issue driving incidents.
- Discovering issues, such as alarms, which do not negatively influence QoE so that attention can be focused on the QoE impacting ones.

All of these use cases can be delivered via the application of machine learning driven by real-time alarms, KPIs and Syslog data streams from devices.



**Figure 6 - Driving profits via machine learning on state data from network devices.**

## Abbreviations

AI	Artificial Intelligence
CPE	Customer Premises Equipment
HFC	Hybrid Fiber-Coax
KPI	Key Performance Indicator
KQI	Key Quality Indicator
MAC	Media Access Control
ML	Machine Learning
MSO	Multiple Systems Operator
NPS	Net Promoter Score
PMO	Present Mode of Operations
QoE	Quality of Experience

## Bibliograph & References

Incognitio. (2016). *2016 BROADBAND CONSUMER QoE SURVEY REPORT*. Incognito Software Systems Inc. Retrieved from Incognitio.com.

Ovum. (2017, February). *Ensuring high quality broadband applications: a necessity for consumers and business critical for service providers*. Retrieved from <https://ovum.informa.com/~media/informa-shop-window/tmt/files/whitepapers/ensuring-high-quality-broadband-applications-feb-17-extended-version.pdf>

# **Segment Routing Proof of Concept for Business Services**

## **Does it work for us**

A Technical Paper prepared for SCTE•ISBE by

**Elaine Yeo**  
Principal Engineer  
Charter Communications  
14810 Grasslands Drive | Englewood, CO 80112-7138  
720.536.1357  
Elaine.Yeo@charter.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
Problem Statement.....	4
1. Segment Routing .....	5
1.1. Concept.....	5
1.2. Operations.....	5
1.3. Segment Types .....	6
2. Proof of Concept Lab .....	6
2.1. Virtual Environment .....	7
2.2. Physical Environment .....	7
2.3. SR POC Topology .....	8
2.4. Test Equipment, Traffic Analyzer, and Impairment Tools.....	10
2.5. Network, Services and Traffic Path .....	11
3. Vendor Platform Operations.....	13
3.1. Maximum SID Depth (MSD).....	14
3.2. Vendor Operations and Comparison .....	14
3.3. Controller Use Case and Interoperability.....	14
3.4. Wireshark Captures.....	17
3.5. Segment Routing Global Block (SRGB) .....	18
3.6. Segment Routing Mapping Server .....	21
3.6.1. SR/LDP Domain Topology .....	23
3.6.2. Test Case 1 – SR-LDP-SR Interworking with vs without border routers.....	24
3.6.3. Test Cases 3 – Redundant Mapping Server .....	32
3.6.4. Test Case 4 – Services over SR/LDP Domain .....	33
3.6.5. Assessment of SRMS and SR/LDP Interworking with multiple vendors.....	34
4. Design Considerations .....	35
4.1. MPLS MTU.....	35
Conclusion .....	37
Abbreviations.....	38
Bibliography & References .....	39

## List of Figures

Title	Page Number
Figure 1 – CML Client View of SR POC .....	8
Figure 2 – Logical Topology of SR POC.....	9
Figure 3 – Test Equipment Placement .....	10
Figure 4 – Sample Wireshark Capture from CML .....	10
Figure 5 – Impairment Parameters.....	11
Figure 6 – Logical SR POC Topology with T and D Hub Specification .....	12
Figure 7 – Primary Path Traffic Pattern .....	13
Figure 8 – Back-up Path Traffic Pattern .....	13
Figure 9 - Devices to Controller Interoperability for Node MSD Discover and Signaling .....	15

Figure 10 – Example capture of MSD TLV .....	17
Figure 11 – SRGB Vendor Comparison .....	18
Figure 12 –SR-LDP-SR Topology .....	23
Figure 13 – SR/LDP with Node SIDs.....	24
Figure 14 – SR Topology with and without SR/LDP border routers.....	25
Figure 15 – Traffic Direction SR/LDP Topology without SR/LDP router .....	26
Figure 16 – Traffic direction with SR and LDP Labels.....	28
Figure 17 – SR to LDP Operations.....	29
Figure 18 – Redundant Mapping Server Topology .....	32
Figure 19 – L2VPN service over SR/LDP domain .....	33
Figure 20 – Dynamic and Explicit SR Path.....	35

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – SR vs MPLS Label Operations .....	5
Table 2 – Equipment Roles.....	7
Table 3 – Node Type and Role .....	9
Table 4 – MSD Range per Vendor .....	14
Table 5 – Controller and Node MSD Signaling Protocols.....	15
Table 6 – Node and Controller Metric Change Capability .....	16
Table 7 – Prefix-to-SID Index Comparison .....	30
Table 8 – Conflict Resolution Preference Rules Comparison.....	35

# Introduction

Segment Routing is a source based routing methodology that uses a list of unique segment IDs stacked in an order of arrival along a traffic path. It leverages the existing MPLS data plane by encoding a segment ID in each MPLS label, thus creating a stack of labels in the packet header which instructs each node along the path to execute the information within the labels upon receipt and forwards it along to the next node. This technology can be used to simplify and optimize the network, meeting key performance objectives such as latency and at the same time enhance operational efficiencies around capacity reporting, change control management via automation with the use of a controller. The interest lies in whether can Segment Routing be beneficial to business services, how does this simplify and affect a large service provider consisting of multiple network types, various vendor platforms and software. This technology has been chosen as a proof of concept that focuses on interoperability of existing multiple vendors within a service provider network, working together with a controller to maximize the benefits of segment routing, particularly the traffic engineering aspect. During the proof of concept, some observations were made and several key points were brought to attention where cascading effects from a system event were seen to have a potential impact that affects CAPEX, OPEX, and architectural processes. These findings from the proof of concept also provides information that can influence business decisions to move forward with Segment Routing or retain the existing mechanisms used today in the network.

## Problem Statement

Charter was initially presented the opportunity to provide a solution to meet latency requirements per contractually agreed Service Level Agreement (SLA) for Cell Tower Backhaul (CTBH) services under a new architecture design. Existing services that had originally met the necessary SLAs must continue to meet the requirements if an architecture were to change to a Hub and Spoke design. The new design would cause traffic to double back to the direction it had traversed, thus increasing the latency to its final destination. For example, with a Hub and Spoke architecture, a cell tower's traffic will have to traverse east to a hub location before doubling back past the source over to the west where the end destination of the Mobile Switching Center (MSC) resides.

RSVP-TE is the common choice, given the history of successful deployment and the ability to provide fast-reroute. However, with the modern IP networks we have today, the need to keep up with the growing demands of network capacity and service quality makes it difficult to scale without compromising network resources to support traffic engineering and other pertinent applications. Segment Routing (SR), a fairly new mechanism simplifies the need of separate protocols such as IGP, LDP and RSVP-TE interacting in a single network and alleviates network resources to hold network state within the core. It serves to remove state from the core network and keeping it in the packet and the ingress node. These features offered by Segment Routing peaks the interest of a potential alternative to RSVP-TE. Segment Routing has the prospect to keep things simple and making room to scale for future enhancements.

In addition for the need to counter latency with a simple approach, other benefits such as using Topology Independent Loop Free Alternate TI-LFA with SR provides for 100% coverage of the network, making it possible to compute, instantiate traffic engineering paths and restore traffic optimally. This is best accomplished when coupled with a controller to provide path computation for optimal routes during a network failure with 100% visibility of the network. From a trending perspective, the lack of network visibility makes it difficult to determine usage patterns and flow characteristics. Without these trending information, it is challenging to learn and evolve our network while planing for growth. Operationally, route stability has been a manual process, potentially introducing human error, i.e configuration errors and delay reaction to act quickly to re-route traffic to an optimal path. Historically, service impacting



outages are triggered by fiber cuts resulting in secondary root causing of sub-optimal back-up failover to latent or congested path in the absence of traffic protection and global path reoptimization.

To provide options of either RSVP-TE or SR-TE, a proof of concept was first pursued for Segment Routing alone to see if it would fit Charter's business model and simplify operational process. The proof of concept lab was built based on an example Charter market. To ensure that Segment Routing would fit into Charter's networks, the proof of concept was vetted against 3 existing major vendor platforms deployed across all legacy companies. This document outlines the differences and interoperability of the multiple platforms in addition to the functionality of Segment Routing and its effect on Charter.

## 1. Segment Routing

There are two types of SR technologies, SR over MPLS (SR-MPLS) and SR over IPv6 (SRv6). This document will focus on SR-MPLS. SR can be implemented over the existing MPLS architecture.

### 1.1. Concept

Segment Routing is a source based paradigm that uses an ordered list of segments appended to the packet header. These lists of segments serve as an abstraction instruction sets for the source node to process and execute on the path to take. A segment is encoded as an MPLS label and a list of segments are essentially a stack of labels. The segments are processed from top to bottom. With the list of segments holding the instructions on traffic path, the state is no longer held in the network but rather within the packet. Only the source node is required to compute and encode the instruction list, while the transit nodes simply reads the top most label before passing it along.

### 1.2. Operations

SR's top most segment is known as the Active Segment. The active segment is the segment that is processed by the receiving node. Segment Routing's segment list operations uses the same existing MPLS forwarding method. The table below details the correlation of the operation in a row within SR to the same row within MPLS label operations.

**Table 1 – SR vs MPLS Label Operations**

<b>SR Segment List Operations</b>	<b>MPLS Label Stack Operations</b>
<b>PUSH</b> Inserts an active segment over the list of segments pushes the label stack forward.	<b>PUSH</b> Injects a label over the label stack
<b>CONTINUE</b> The active segment is not completed, remains active and continues to next destination	<b>SWAP</b> Replace the top label with a new label
<b>NEXT</b> The active segment is completed. The next segment on the list is the active segment	<b>POP</b> Removes top label from label stack

### 1.3. Segment Types

There are many segment types in SR for specific functions. This document will use terms listed below that are relevant to only what was used in the SR Proof of Concept (POC) .

#### Segment Routing Global Block (SRGB)

Globally unique range of labels recognized within a node. Configured within IGP. The SRGB can be configured to the desired range. When the SRGB is set, a range of labels are set aside to be used, unique to all nodes within the SR domain. The beginning number of the range is known as the SRGB base.

*Example:*

*Manually configured SRGB = 16000 – 19000. Therefore 16000 is the SRGB base.*

#### SID

Segment Identifier (SID). The SID is encoded as an MPLS label that identifies the segment and sets it apart from other nodes or links.

#### Node SID

Allocated from the pool of SRGB. Globally significant and unique in within the SR domain. Similar to a router ID, it typically is attached to the loopback of the node. See Prefix-SID for how a Node SID is derived.

#### Adj-SID

Also known as IGP Adjacency Segment. A segment local only to the node. Dynamically allocated and advertised only to its direct neighbor via the adjacency link. The Adj-SID, when dynamically allocated uses the next label after the SRGB. Each Adj-SID is unique with the node only and mapped to its link to the next neighbor.

*Example:*

If SRGB = 16000 – 19000, then the first dynamically allocated Adj-SID is 19001.

#### Prefix-SID

Also known as prefix segment or Node SID. Global segment attached to a prefix.

Prefix-SID = SRGB base + Index table = Absolute value

*Example:*

*SRGB base = 16000*

*Index = 3*

*Prefix-SID = 16000 + 3 = 16003*

#### Binding-SID

An outer SID that nests a segment list, commonly used to stitch across various domains. This can be thought as a form of label compression to reduce the label stack depth

## 2. Proof of Concept Lab

The SR Proof of Concept (POC) lab consist of a virtual and physical environment which mocks up the Charter market. The purpose in using both environment is to eliminate the need for too many physical

devices to make up a simulated network. A Nexus 5548UP switch was used to bridge between the virtual and the physical environment.

## 2.1. Virtual Environment

Cisco Modeling Lab (CML) was used as the virtual environment platform. Each node had a mix of a CRR or a DTR.

Platform – Cisco Modeling Lab

Platform Version – 1.3

Nodes – 12

Node Image – IOS-XRv9K

Image Software – 6.3.2

Roles – Distribution Routers (DTR) and Core Routers (CRR)

## 2.2. Physical Environment

Various models of devices were used from three main vendors. These vendors were selected since they were widely used across the former merged companies. The models used in the POC were a representative of the roles they play in current production.

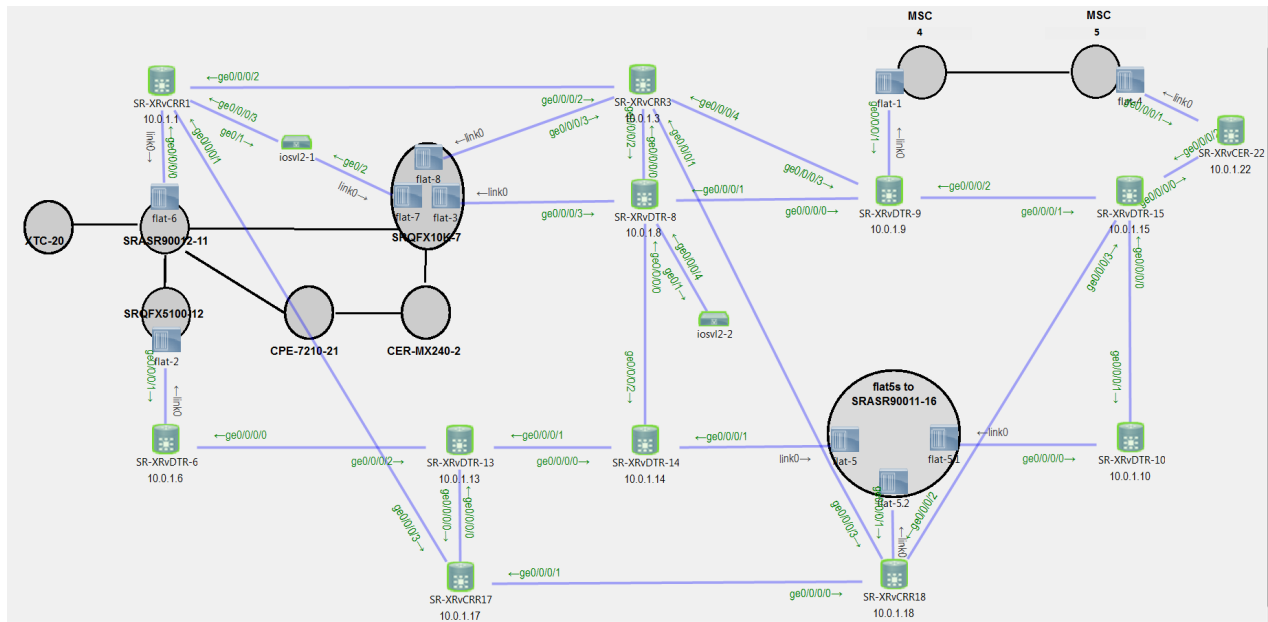
**Table 2 – Equipment Roles**

Equipment	Software	Role
MX240	18.2	CER
7750 SR-7	15.1	MSC
7750 SR-7	15.1	MSC
QFX10K	18.2	DTR
QFX5K	18.3	DTR
ASR9001	6.4.1	DTR
ASR9001	6.4.1	DTR
ASR9001	6.4.1	DTR
ASR9001	6.4.1	DTR
7210 SAS-M	8.0	CPE

Device roles justification:

1. The Juniper MX series router was placed as a Commercial Edge Router (CER) in the SR POC which have a role as a PE router and an aggregation router for services.
2. The Nokia 7750 SR routers are used as Mobile Switching Center (MSC) routers to simulate CTBH production.
3. The Juniper QFX series are used as DTRs where a CER will connect to it.

### 2.3. SR POC Topology



**Figure 1 – CML Client View of SR POC**

Diagram above is a layout of the SR topology in the CML client view.



Represents an XRv9K node



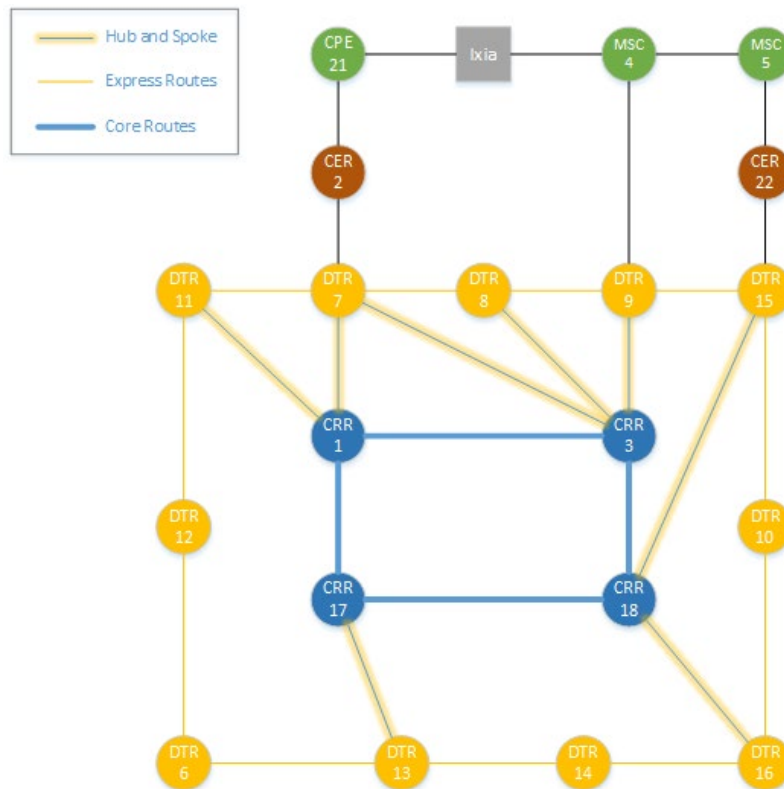
Represents a physical equipment

The SR POC is designed and built based on an example Charter market with an added scenario of an East/West connection. The topology is a hybrid of the hub and spoke topology and the East West express route, where selected DTRs are assigned back to select CRRs (represented by the Hub and Spoke connections) and a high metric cost on the East West express route.

Below is the logical representation of the topology within CML client view.

**Table 3 – Node Type and Role**

Node Type	Role
XRv9K	CRR 1
MX240	CER 2
XRv9K	CRR 3
7750 SR-7	MSC 4
7750 SR-7	MSC 5
XRv9K	DTR 6
QFX10K	DTR 7
XRv9K	DTR 8
XRv9K	DTR 9
XRv9K	DTR 10
ASR9001	DTR 11
QFX5K	DTR 12
XRv9K	DTR 13
XRv9K	DTR 14
XRv9K	DTR 15
ASR9001	DTR 16
XRv9K	CRR 17
XRv9K	CRR 18
7210 SAS-M	CPE 21
XRv9K	CER 22



**Figure 2 – Logical Topology of SR POC**

A simpler topology was built in comparison to the typical hub and spoke design where only a few DTR nodes were homed back to select CRRs. Since traffic was focused from the CPE to the MSC node, only the DTR 7 node is designed to home back to two separate CRRs to allow for simulation of primary path failure.

## 2.4. Test Equipment, Traffic Analyzer, and Impairment Tools

### Test Equipment

Ixia's IxNetwork was used to generate traffic to the SR POC. The type of traffic emulates the CTBH traffic from the subscriber in production today. Ixia was placed between the CPE and the MSC 4 node to allow for bi-directional traffic.




Figure 3 – Test Equipment Placement

### Traffic Analyzer

The packet capture feature within CML was used to analyze and understand SR traffic. Since the feature is available within the virtual environment, traffic between physical equipment was not captured.

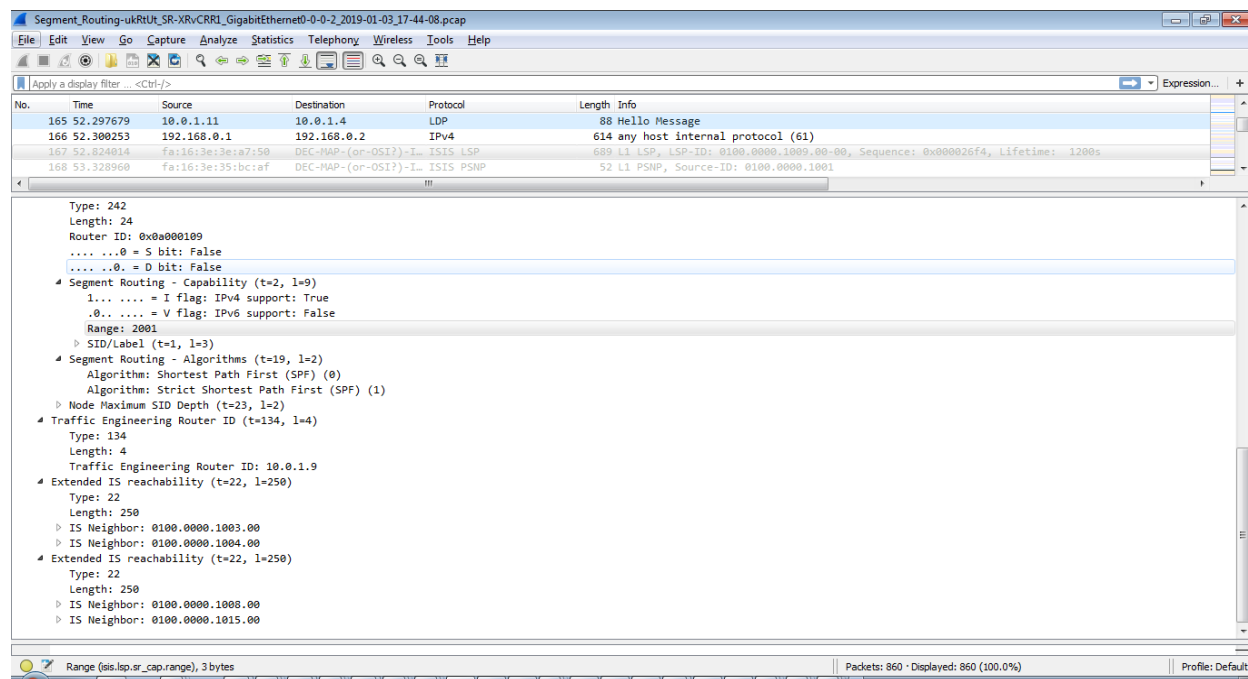



Figure 4 – Sample Wireshark Capture from CML

### Impairment Tools

CML provides the ability to apply latency, jitter, and packet loss at the link level parameters between the XRv9K nodes. Latency was used on the East West direction. The purpose of adding latency is to simulate a physical long path around the ring in the event of a failure on the shorter path from CPE to the MSC and understand the dynamics of SR-TE metric delay constraints.



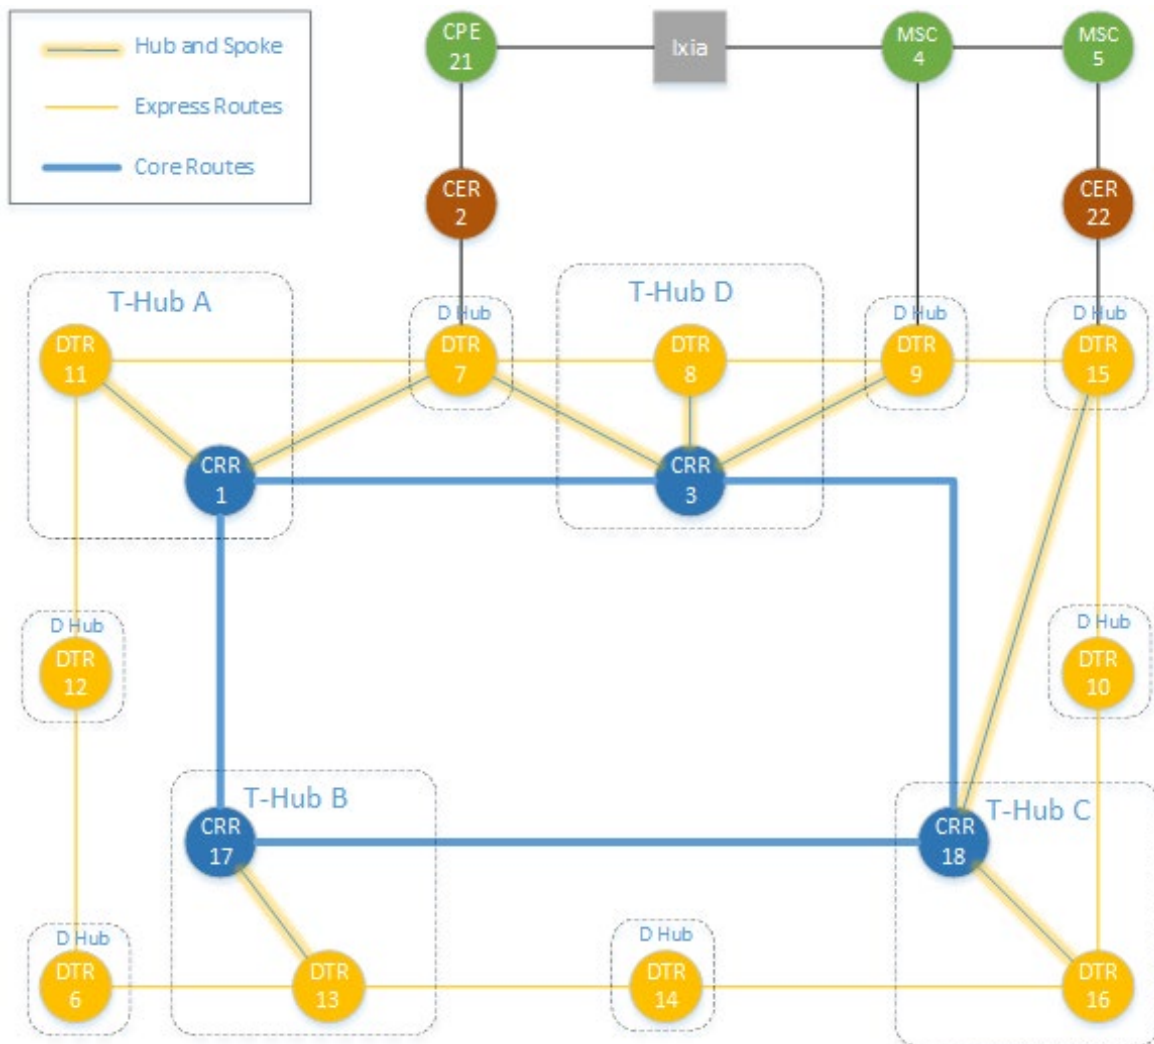
**Figure 5 – Impairment Parameters**

## **2.5. Network, Services and Traffic Path**

To simulate CTBH services across the TN market, Ethernet over MPLS were implemented from the MSC (7750s) across the IS-IS network to the CER (MX240) with a layer 2 hand-off to the CPE. Simulated CTBH Layer 2 traffic is generated with Ixia's IxNetwork.

Traffic between CPE 21 and MSC 4 will traverse via their local DTR hubs across the SR POC network.


Topology below is a hybrid of logical and physical environment. The T-Hub and D-Hub locations are specified to show the physical location and direction of traffic. The DTR homes back to its assigned CRR for forwarding via the Hub and Spoke connections.



**Figure 6 – Logical SR POC Topology with T and D Hub Specification**


Test run was performed using the primary and back-up path as shown below to determine the SR characteristics and functionality during a failover. The traffic path is from CPE 21 to MSC 4 via CRR3 and vice versa based on dynamic SR





**Figure 7 – Primary Path Traffic Pattern**

The traffic path from CPE 21 to MSC 4 via CRR1, CRR 17, CRR 18, and CRR 3 in that order during a failover as depicted in the red line and vice versa based on dynamic SR



**Figure 8 – Back-up Path Traffic Pattern**

### 3. Vendor Platform Operations

Charter utilizes multiple vendor platforms in various networks. To understand how the different platform work and interoperate, three main vendors were selected for the proof of concept testing. These three vendors are Cisco, Juniper, and Nokia.

### 3.1. Maximum SID Depth (MSD)

MSD is the maximum number of SIDs supported by a node or link. Since each SID is encoded in an MPLS label, the MSD can be referred to as the maximum number of labels supported. The MSD of the device determines the Base MPLS Imposition (BMI), where BMI is the total number of labels imposed inclusive of all service and transport labels.

### 3.2. Vendor Operations and Comparison

The MSD is defined by the dataplane capability of each vendor platform. There is a variance of MSD supported across multiple vendor platforms due to the type of network processors (NPU) used. Depending on the vendor platforms, MSD at each node can be provisioned if not already set at the maximum. Table 1 below shows a range of MSD from SR capable line cards in production with their the label depth and operational limits. MSD in Table 4 is the BMI.

**Table 4 – MSD Range per Vendor**

Vendor	Platforms Role	MSD	PUSH	POP
Vendor X	PE	3-16	3-16	2-16
Vendor Y	Core, PE	10	10	10
Vendor Z	Core, PE, CPE	6-12	6-12	6-12

It introduces some complexity when trying to establish an SR LSP or SR-TE LSP along a path that consists of multiple vendors with varying MSDs as shown in Table 1 without exceeding the lowest supported MSD. Careful planning will be required when designing the LSP paths and manual analysis of node MSDs where the LSP traverses can be labor intensive along with the record keeping of MSDs per type of device. Additionally, MSDs can change after software upgrades requiring engineers to keep track of the changes.

Factors that introduces complexity:

1. Tracking of multiple vendor platforms used in the network with different MSDs
2. Inconsistent software on same device models within a network can lead to different MSDs
3. Node MSD and link MSD are not homogenous leading to additional leg work in tracking MSD types
4. Different line cards within a node supports different MSD

All the complexity mentioned above can be mitigated with the use of a controller as defined in following section.

### 3.3. Controller Use Case and Interoperability

When using a controller to compute the SR paths, the controller can learn the MSDs of each node and ensure the segment list depth does not exceed the MSD of the nodes on the computed path. The controller can receive the MSD of nodes via advertisement methods below.

### Node MSD Advertisement Methods

There are 4 ways to advertise MSD capabilities.

1. IS-IS
  - Using the **Node MSD sub-TLV** within the **Router Capability TLV**
2. OSPF
  - Using the **Node MSD sub-TLV** within the LSA Type Opaque
3. Path Computation Element Protocol (PCEP)
  - Using the **SR-PCE-Capability sub-TLV** within the **Path-Setup-Type-Capability TLV**
4. BGP Link State
  - Using **Node Attribute TLV**

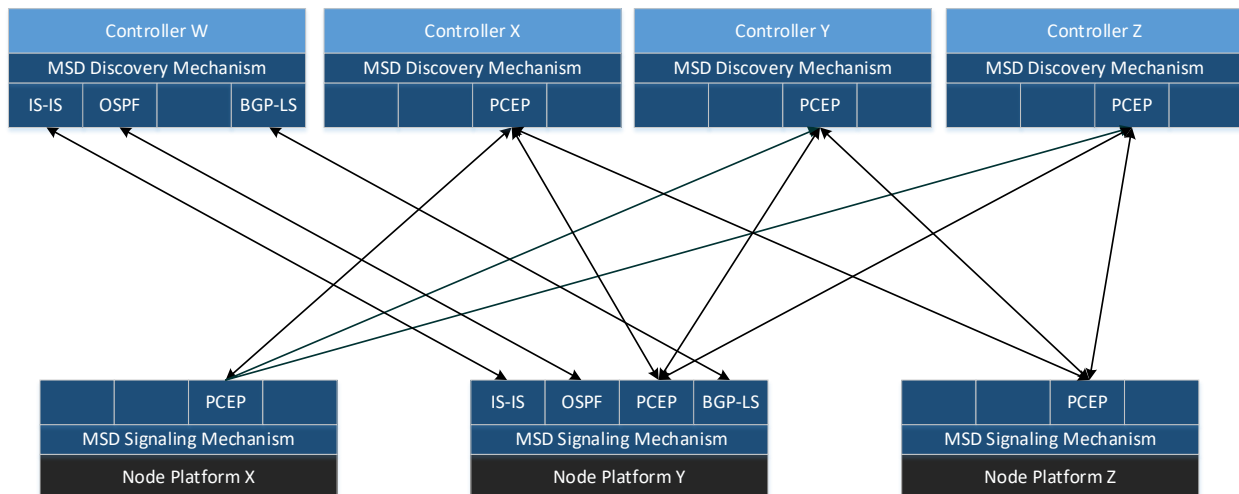
However, each controller's learning capability varies per vendor. See section “*Gap Analysis across vendor platforms*” for specific signaling protocols used for learning the MSD.

### Gap Analysis across vendor platforms

Four vendor controllers and three vendor devices are selected as the subject for interoperability of MSD signaling between controllers and devices in this document. Currently, the learning and advertisement of MSD varies across multiple vendor platforms for both controller and nodes.

**Table 5 – Controller and Node MSD Signaling Protocols**

Signaling Protocols	Controller W	Controller X	Controller Y	Controller Z
IS-IS	Yes			Future
OSPF	Yes			Future
PCEP		Yes	Yes	Yes
BGP-LS	Yes			Future



**Figure 9 - Devices to Controller Interoperability for Node MSD Discover and Signaling**

As shown above, not all vendor platforms support the same signaling protocol for MSD. While all controllers support learning of the MSD via PCEP except Controller W, which only supports IGP and BGP-LS, Controller Y additionally supports protocols such as IGP and BGP-LS.

Based on the diagram above, Controller W will have interoperability issues in discovering the MSD with Node X and Node Z device platform.

#### Transport and Service Labels with MSD

Since the MSD is the BMI that includes all transport and service labels, the controllers currently do not make a distinction between the types of labels. Establishing the MSD without that distinction between transport and service labels can result in the Path Computation Element (PCE) computing a sub-optimal path and/or returning a path that exceeds the MSD of a node without taking into account the service labels. Hence, a constraint can be set using the metric object in an exchange between the PCE and Path Computation Client (PCC) to reduce the label depth of a computed path.

A node can consist of the default MSD and the configured MSD. To signal a reduced label depth, the devices must allow configuration of the MSD.

Below is an example of Vendor Y's MSD signaling operations.

#### Option #1: Learning the default MSD

For the PCE to learn the MSD, the PCC will have a PCEP open session with the PCE.  
During the open session, the PCC will signal its default MSD X.

#### Option #2: Learning the configured MSD

PCC signals MSD Y during a path computation request via PCEP within the metric object.  
MSD Y overrides MSD X if the latter is already learned by the PCE.

Note: Vendor Y's MSD signaling operations complies with IETF Draft - <https://tools.ietf.org/html/draft-ietf-pce-segment-routing-16>

The table below summarizes the interaction of each vendor type controller with a different vendor device on what can be signaled. The exception would be Controller W which is not interoperable with node platform X and Z, but will learn the default MSD only from node platform Y.

**Table 6 – Node and Controller Metric Change Capability**

	<b>Controller W</b>	<b>Controller X</b>	<b>Controller Y</b>	<b>Controller Z</b>
Node X	Not Interoperable	Default/Configurable	Default/Configurable	Default/Configurable
Node Y	Default	Default/Configurable	Default/Configurable	Default/Configurable
Node Z	Not Interoperable	Default/Configurable	Default/Configurable	Default/Configurable

### 3.4. Wireshark Captures

Below is a capture of the Node MSD type 23 capability advertised by a Cisco ASR9K as specified in RFC8491.

```

> Frame 16592: 417 bytes on wire (3336 bits), 417 bytes captured (3336 bits)
> Ethernet II, Src: Cisco_58:30:64 (b0:26:80:58:30:64), Dst: DEC-MAP-(or-OSI?)-Intermediate-System-Hello? (09:00:2b:00:00:05)
> 802.1Q Virtual LAN, PRI: 6, DEI: 0, ID: 16
> Logical-Link Control
> ISO 10589 ISIS InTRA Domain Routeing Information Exchange Protocol
  ISO 10589 ISIS Link State Protocol Data Unit
    PDU length: 396
    Remaining lifetime: 1199
    LSP-ID: 0100.0000.1011.00-00
    Sequence number: 0x00002b3f
    Checksum: 0xacaf [correct]
    [Checksum Status: Good]
    > Type block(0x01): Partition Repair:0, Attached bits:0, Overload bit:0, IS type:1
    > Area address(es) (t=1, l=4)
    > Protocols supported (t=129, l=1)
    > IP Interface address(es) (t=132, l=4)
    > Traffic Engineering Router ID (t=134, l=4)
    > Extended IP Reachability (t=135, l=53)
    > Hostname (t=137, l=13)
  1 Router Capability (t=242, l=24)
    Type: 242
    Length: 24
    Router ID: 0x0a00010b
    .... ..0 = S bit: False
    .... ..0 = D bit: False
    > Segment Routing - Capability (t=2, l=9)
    > Segment Routing - Algorithms (t=19, l=2)
  2 Node Maximum SID Depth (t=23, l=2)
    MSD Type: Base_MPLS Imposition MSD (1) 3
    MSD Value: 10 4
    > Extended IS reachability (t=22, l=250)
```

**Figure 10 – Example capture of MSD TLV**

- 1 All SR capabilities and information are stored in the IS-IS Router Capability TLV Type 242.
- 2 Sub-TLV for node MSD consisting of MSD Type and value.
- 3 Supported MSD Type of BMI MSD specifies that the MSD is based on the total amount the device imposition of labels.
- 4 MSD value of 10 indicates a supported Maximum Depth of 10 SIDs.

*Note: Wireshark captures of Nokia and Juniper do not show MSD capabilities within the IS-IS sub-TLVs as they were not supported at the time of POC testing.*

### 3.5. Segment Routing Global Block (SRGB)

Segment Routing Global Block (SRGB) is a local range of blocks recognized within a node. While SRGB is globally unique, the three vendors were found to have their own default SRGB or SRGB configurable block. The diagram below shows the difference between three vendors and the common SRGB space.

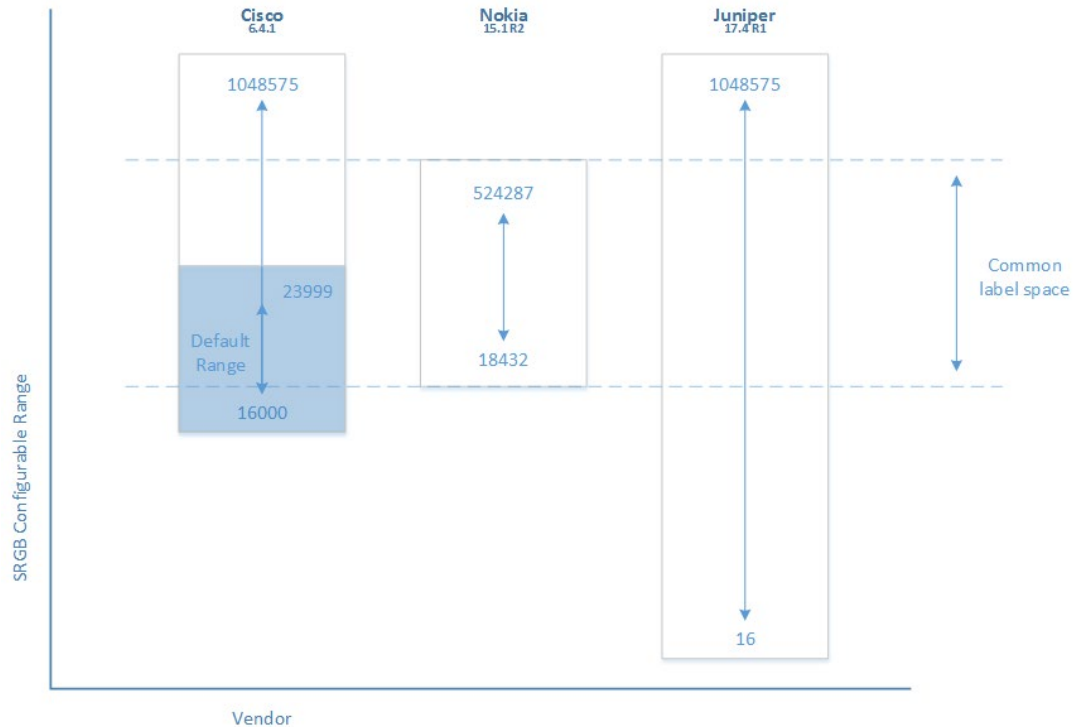


Figure 11 – SRGB Vendor Comparison

#### Cisco

Cisco, by default have their SRGB set between 16000 and 23999. However, a different SRGB other than the default is configurable between 16000 and 1048575. Cisco's configuration of the SRGB resides within the IGP instance.

```
router isis default
 is-type level-1
 net 49.1850.0100.0000.1001.00
 segment-routing global-block 16000 278142
```

SRGB of nodes within the IGP database can be seen using the show command below. Within the CRR 1 node is an excerpt of the SRGB of the router capability of the QFX10K.

Node - SR-XRvCRR1

Show Command: *show isis database verbose*

```
SRQFX10002-7.00-00    0x00004be5    0xb0c6    1128    0/0/0

Area Address:    49.1850

TLV 14:          Length: 2

NLPID:           0xcc

NLPID:           0x8e

Router ID:       10.0.1.7

IP Address:      10.0.1.7

Hostname:        SRQFX10002-7

Router Cap:      10.0.1.7, D:0, S:0

    Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 4000

    SR Algorithm:

        Algorithm: 0

Metric: 50        IS-Extended SR-XRvCRR3.00

Interface IP Address: 172.16.2.23

Neighbor IP Address: 172.16.2.22
```

Note the SRGB base of 16000 and Range of 4000 configured above.

## Nokia

Nokia does not have a default SRGB. A configurable value is allowed between 18432 and 524287 within the device's dynamic label range. Configuring outside of the range results in the error shown below.

```
*B:MSC-4>config>router>mpls-labels># sr-labels start 18000 end 20000

^

Error: Invalid parameter. Label value not in allowed range
```

Configuring the SRGB differs from Juniper and Cisco. Rather than configuring the SRGB in the IGP instance, it is configured in a new category called “mpls-labels” under the router field as shown. Note: Activating the SR is still required in the IGP instance.

```
#-----
echo "ISIS (Inst: 1) Configuration"
#-----

isis 1

    router-id 10.0.1.4

    level-capability level-1

    area-id 49.1850

    advertise-passive-only

    advertise-router-capability as

    level 1

        wide-metrics-only

    exit

    segment-routing

        prefix-sid-range global

        no shutdown

    exit
```

### Juniper

Juniper does not have a default SRGB. A configurable range between 16 and 1048575 is allowed. Juniper devices' SRGB configuration resides within the IGP instance.

```
protocols {
    isis {
        source-packet-routing {
            srgb start-label 16000 index-range 4000;
            node-segment ipv4-index 2;
        }
    }
}
```

Node – SRQFX10002-7



Show Command: *show isis database database extensive*

```
root@SRQFX10002-7> show isis database extensive
IS-IS level 1 link-state database:

SR-XRVCRR1.00-00 Sequence: 0x625, Checksum: 0x3599, Lifetime: 1072 secs
IPv4 Index: 1
Node Segment Blocks Advertised:
  Start Index : 0, Size : 262143, Label-Range: [ 16000, 278142 ]
  IS neighbor: SR-XRVCRR3.00 Metric: 25
    Two-way fragment: SR-XRVCRR3.00-00, Two-way first fragment: SR-XRVCRR3.00-00
    P2P IPv4 Adj-SID: 278146, weight: 0, Flags: --VL--
  IS neighbor: SRQFX10002-7.00 Metric: 50
    Two-way fragment: SRQFX10002-7.00-00, Two-way first fragment: SRQFX10002-7.00-00
    P2P IPv4 Adj-SID: 278148, weight: 0, Flags: --VL--
  IS neighbor: SR-XRVCRR17.00 Metric: 25
    Two-way fragment: SR-XRVCRR17.00-00, Two-way first fragment: SR-XRVCRR17.00-00
    P2P IPv4 Adj-SID: 278144, weight: 0, Flags: --VL--
  IP prefix: 10.0.0.16/30 Metric: 25 Internal up
  IP prefix: 10.0.0.56/30 Metric: 25 Internal up
  IP prefix: 10.0.1.1/32 Metric: 0 Internal up
  IP prefix: 10.200.0.2/31 Metric: 50 Internal up

Header: LSP ID: SR-XRVCRR1.00-00, Length: 246 bytes
  Allocated length: 284 bytes, Router ID: 10.0.1.1
  Remaining lifetime: 1072 secs, Level: 1, Interface: 558
  Estimated free bytes: 257, Actual free bytes: 38
  Aging timer expires in: 1072 secs
  Protocols: IP

Packet: LSP ID: SR-XRVCRR1.00-00, Length: 246 bytes, Lifetime : 1198 secs
  Checksum: 0x3599, Sequence: 0x625, Attributes: 0x1 <L1>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 18, Packet version: 1, Max area: 0

TLVs:
  Area address: 49.1850 (3)
  IS extended neighbor: SR-XRVCRR3.00, Metric: default 25
    IP address: 10.0.0.57
    Neighbor's IP address: 10.0.0.58
    Unknown sub-TLV type 15, length 2
    P2P IPv4 Adj-SID - Flags: 0x30(F:0,B:0,V:1,L:1,S:0,P:0), weight:0, Label: 278146
    P2P IPv4 Adj-SID: 278146, weight: 0, Flags: --VL--
  IS extended neighbor: SRQFX10002-7.00, Metric: default 50
    IP address: 10.200.0.2
    Neighbor's IP address: 10.200.0.3
    Unknown sub-TLV type 15, length 2
    P2P IPv4 Adj-SID - Flags: 0x30(F:0,B:0,V:1,L:1,S:0,P:0), weight:0, Label: 278148
    P2P IPv4 Adj-SID: 278148, weight: 0, Flags: --VL--
  IS extended neighbor: SR-XRVCRR17.00, Metric: default 25
    IP address: 10.0.0.17
    Neighbor's IP address: 10.0.0.18
    Unknown sub-TLV type 15, length 2
    P2P IPv4 Adj-SID - Flags: 0x30(F:0,B:0,V:1,L:1,S:0,P:0), weight:0, Label: 278144
    P2P IPv4 Adj-SID: 278144, weight: 0, Flags: --VL--
  Speaks: IP
    IP address: 10.0.1.1
    IP extended prefix: 10.0.0.16/30 metric 25 up
    3 bytes of subtlvs
    IP extended prefix: 10.0.0.56/30 metric 25 up
    3 bytes of subtlvs
    IP extended prefix: 10.200.0.2/31 metric 50 up
    3 bytes of subtlvs
    IP extended prefix: 10.0.1.1/32 metric 0 up
    11 bytes of subtlvs
    Node SID, Flags: 0x40(R:0,N:1,P:0,E:0,V:0,L:0), Algo: SPF(0), value: 1
  Hostname: SR-XRVCRR1
  Router Capability: Router ID 10.0.1.1, Flags: 0x00
    SPRING Capability - Flags: 0x80(I:1,V:0), Range: 262143, SID-Label: 16000
    SPRING Algorithm - Algo: 0
    SPRING Algorithm - Algo: 1
  No queued transmissions
```

Note: All non-SR nodes will start receiving SR router capabilities in their LSDB when there are SR nodes in the network.

### 3.6. Segment Routing Mapping Server

Purpose: Interoperability between SR and LDP

Functions:

1. Creates a database of prefixes which are not SR capable for both mapping servers and clients.

2. Advertises prefix to SID mappings of non SR routers to SR routers.
3. A control plane mechanism.
4. Part of IGP extensions encoded in SID/Label Binding TLV and Extended Prefix range TLV for ISIS and OSPF, respectively.

Restrictions:

1. Mapping Server must have IGP adjacency to the network.
2. For a network that relies on mapping servers to interop between protocols, a redundant mapping server is recommended.
3. SID-mapping entries learned from one IGP process or instance, cannot be used to learn or calculate prefix-SIDs from another IGP process or instance. Each mapping server is required to be configured per IGP instance.
4. Does not support VRFs.
5. For traffic path from SR domain to LDP domain, a border router between both domains must be enabled with SR and LDP. This document refers to the border router as an SR/LDP border router.

Deployment Methods:

1. Dedicated physical device not inline
2. Inline device
3. Virtualized

Best practice:

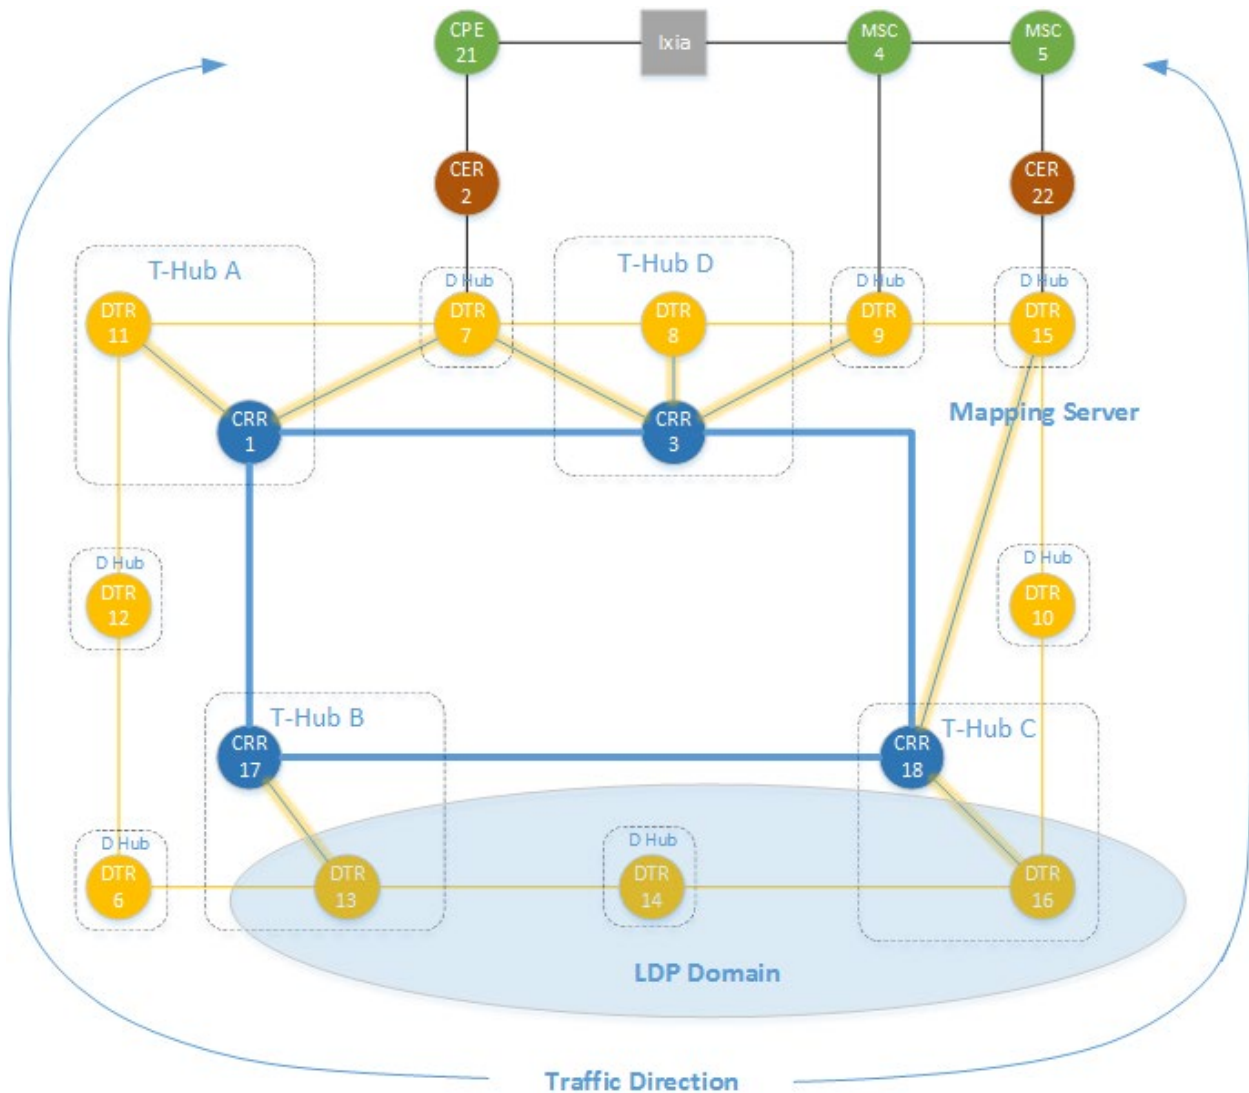
1. No more than two mapping servers. Too many is counter-productive and results in having to track the Segment Routing Mapping Server (SRMS) prefix configurations to ensure they are the same across the board.
2. Placement of SRMS where the only two are in a single hub poses a single point of failure.

When is a Mapping Server (SRMS) required:

1. LDP to SR
  - Does not require SRMS
  - Why is SRMS not required?
    - When Independent Label Distribution Control Mode ([RFC5036](#)) is active on the router that is on the border of the LDP and SR domain.
    - Any node on the LDP to Segment Routing border automatically installs LDP-to-SR forwarding entries
2. SR to LDP
  - Requires SRMS
  - Why is SRMS required?

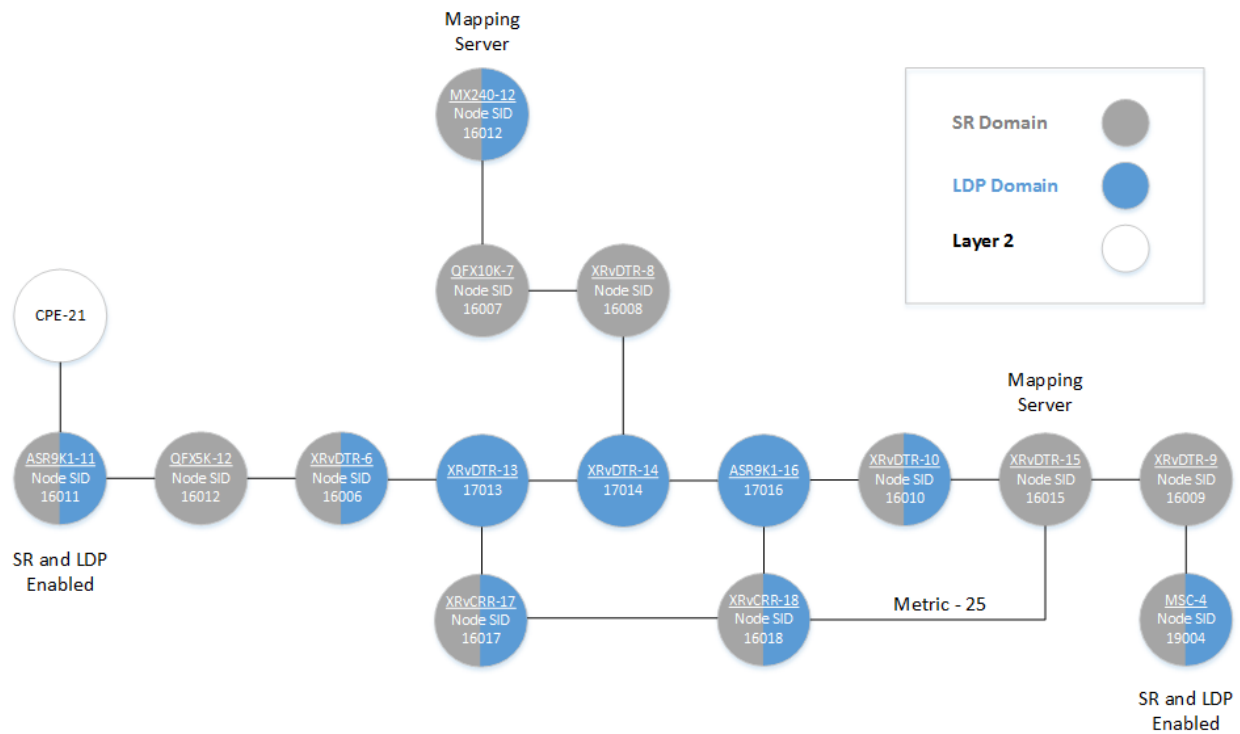
- Since LDP nodes are not capable of advertising a prefix SID, the SRMS acts as a translator for all other SR nodes by mapping prefixes of LDP nodes to SIDs.
- It advertises the prefix-to-SID mappings to all other SR nodes, which are mapping clients.

### 3.6.1. SR/LDP Domain Topology



**Figure 12 –SR-LDP-SR Topology**

The LDP domain consist of Node 13, 14, and 16, while all other nodes are SR nodes. The exception pertains to DTR 9 and CER 2, where both nodes are the termination point of the L2VPN service. Hence SR and LDP are enabled for targeted LDP sessions. Traffic is steered around the outer ring of the topology as depicted in figure 12.



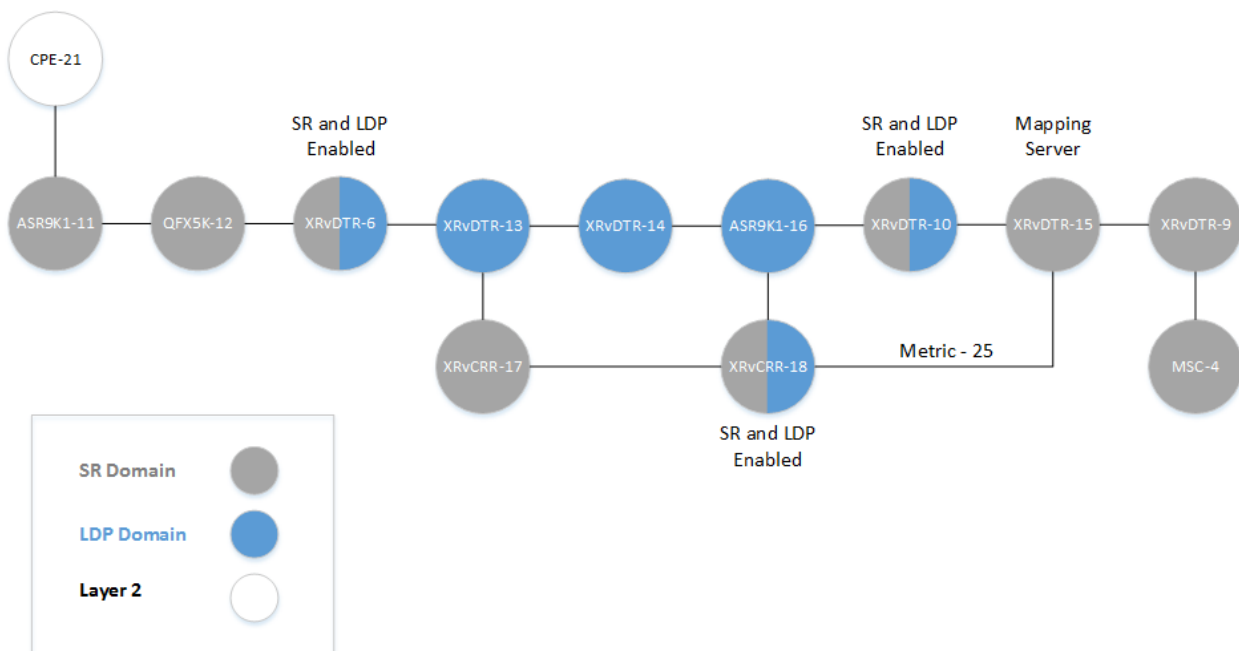
**Figure 13 – SR/LDP with Node SIDs**

Figure 13 represents SR/LDP topology with node SIDs. Index of 1013, 1014 and 1016 was used for prefix-to-SID mapping of LDP devices. A second mapping server is added for SRMS redundancy testing for a later test case.

### **3.6.2. Test Case 1 – SR-LDP-SR Interworking with vs without border routers**

The purpose of this test case is to test the behavior of mapping server advertisements and mapping client interaction. Note that XRVDR-17 node is not set as an SR/LDP border router.

*(While all vendors documentation stipulate that the border between SR and LDP regions run both LDP and SR, acting as SRMS client and perform stitching; this test case was created to further understand SR/LDP interworking and scan for any unknowns.)*



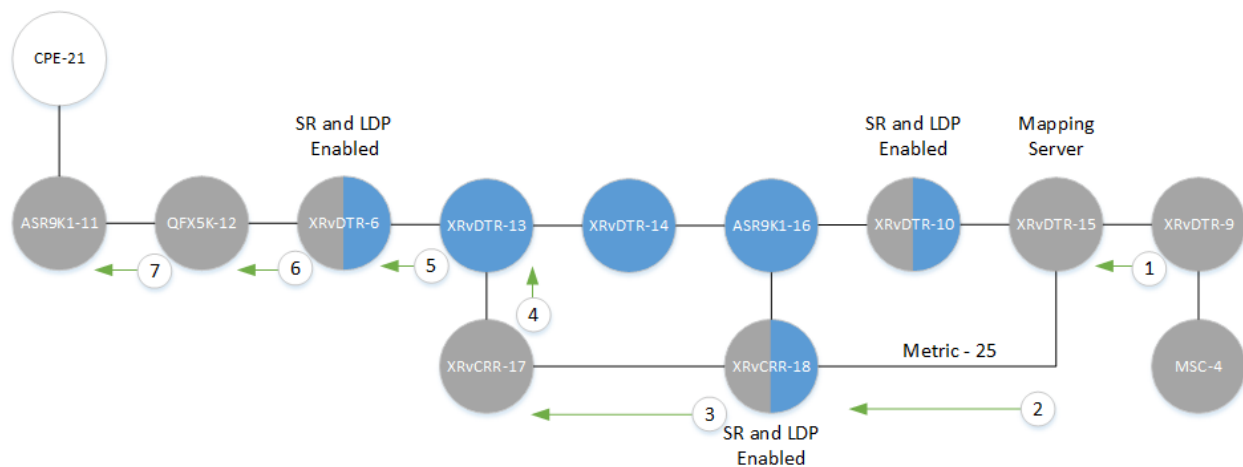
**Figure 14 – SR Topology with and without SR/LDP border routers**

1. Only XrvDTR-6, XrvDTR-10, and XrvDTR-18 have both SR and LDP enabled.
2. All other devices within SR domain are SR enabled only.
3. All other devices within LDP domain are LDP enabled only.
4. XrvDTR-15 is the mapping server, the rest of the nodes are mapping clients.
5. XrvCRR-17 is not an SR/LDP border router.

**Outcome when traffic passes through a non SR/LDP border router:**

With XrvCRR-17 set as an SR node instead of an SR/LDP border router, traffic failed when steered through the lower cost path as shown in step 4 below. Should Node 17 be enabled as SR/LDP border router, traffic will continue to be forwarded as shown in the following steps 5 through 7.

Traffic Path: Destination is ASR9K1-11 with a destination SR label of 16011.



**Figure 15 – Traffic Direction SR/LDP Topology without SR/LDP router**

1. XRvDTR-9 pushes label 16011 which is ASR9K1-11's Node SID.
2. XRvDTR-15 continues path to ASR9K1-11 using label 16011.
3. XRvCRR-18 continues path to ASR9K1-11 using label 16011.
4. XRvCRR-17, an SR only node did not receive a label binding FEC from XRvDTR-13, hence do not have label. **This is because XRvCRR-17 is not enabled as a border router for SR/LDP.**  
Even when node 17 has received a prefix-to-SID mapping label of 17013 to node 13 (10.0.1.13), no labels were imposed.

```
RP/0/RP0/CPU0:SR-XRvCRR17#show cef 10.0.1.13
Fri Nov  9 18:27:44.894 UTC
10.0.1.13/32, version 5105, labeled SR, internal 0x1000001 0x81 (ptr 0xd486aa0) [1], 0x0 (0xd61ff
Updated Nov  8 21:51:51.518
remote adjacency to GigabitEthernet0/0/0/0
Prefix Len 32, traffic index 0, precedence n/a, priority 1
via 10.0.0.26/32, GigabitEthernet0/0/0/0, 8 dependencies, weight 0, class 0 [flags 0x0]
path-idx 0 NHID 0x0 [0xdc131b0 0x0]
next hop 10.0.0.26/32
remote adjacency
  local label 17013      labels imposed {None}
```

Looking further, SR/LDP merge is requested but has no active flag. There were no operations to replace the SR label with an LDP label.

```

RP/0/RP0/CPU0:SR-XRvCRR17#show cef 10.0.1.13 flags
Fri Nov 9 20:10:53.809 UTC
10.0.1.13/32, version 5105, labeled SR, internal 0x1000001 0x81 (ptr 0xd486aa0) [1], 0x0 (0xd61
leaf flags: owner locked, inserted
leaf flags2: LDP/SR merge requested,sr-pfx,
leaf ext flags: EXTERNAL_REACH LC.sr-mpls.sr-pfx-sid,
Updated Nov 8 21:51:51.518
remote adjacency to GigabitEthernet0/0/0/0
Prefix Len 32, traffic index 0, precedence n/a, priority 1
via 10.0.0.26/32, GigabitEthernet0/0/0/0, 8 dependencies, weight 0, class 0 [flags 0x0]
path-idx 0 NHID 0x0 [0xdc131b0 0x0]
next hop 10.0.0.26/32
remote adjacency
local label 17013 labels imposed {None}

```

To demonstrate further that the mapping client (CRR-17 node) is receiving the mapping policy from the SRMS, the following command shows the local label's source is from the RIB but not the LSD which helps provide the operations to replace the SR label with an LDP label.

```

RP/0/RP0/CPU0:SR-XRvCRR17#show cef 10.0.1.13 detail | i "source"
Fri Nov 9 20:17:17.395 UTC
gateway array (0xd4b4260) reference count 10, flags 0x8068, source rib (?), 0 backups
RP/0/RP0/CPU0:SR-XRvCRR17#

```

- Assuming CRR-17 is an SR/LDP border router, XRvDTR-13 pushes LDP label 24039 towards XRvDTR-6.

```

RP/0/RP0/CPU0:SR-XRvDTR-9#traceroute 10.0.1.11
Mon Nov 12 21:17:33.418 UTC

Type escape sequence to abort.
Tracing the route to 10.0.1.11

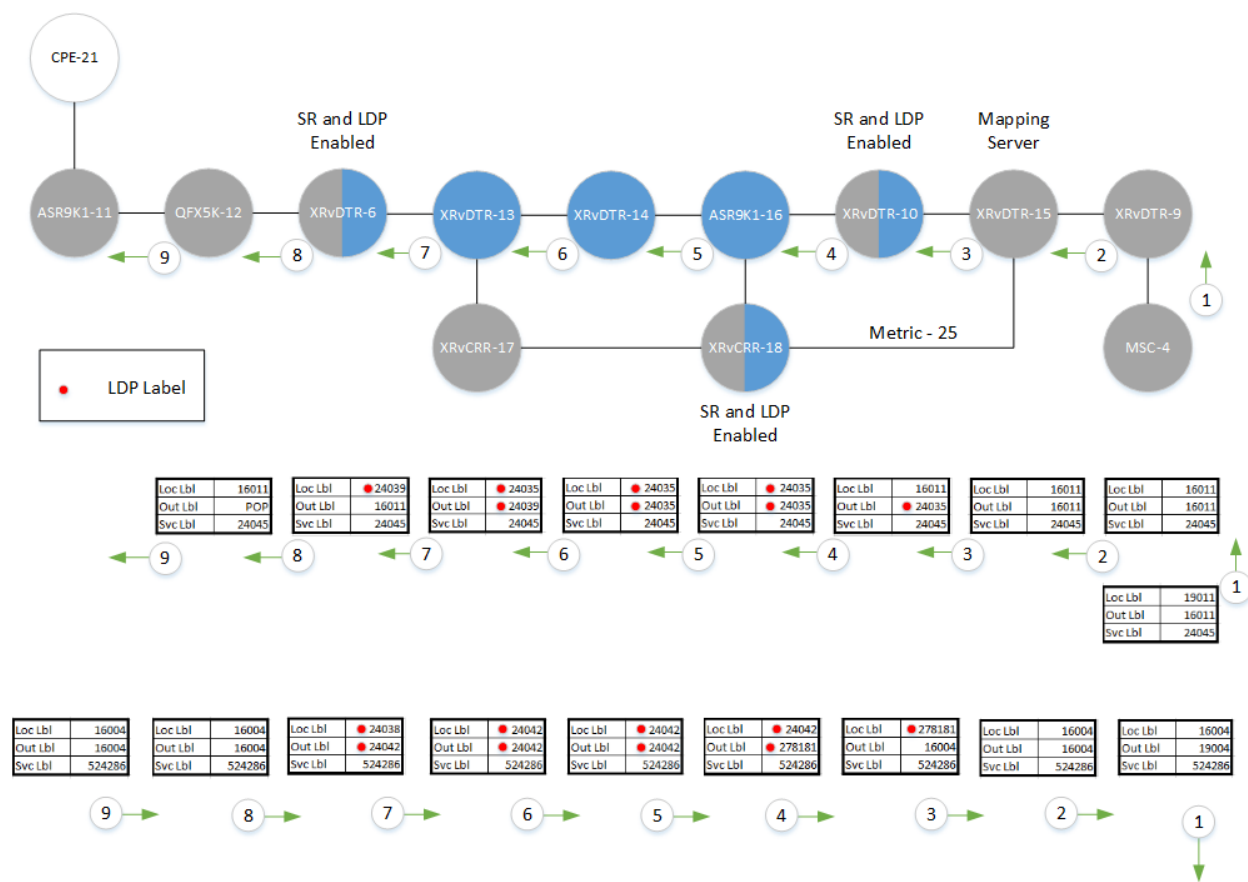
 1  10.0.0.33 [MPLS: Label 16011 Exp 0] 19 msec  3 msec  3 msec
 2  10.0.0.53 [MPLS: Label 16011 Exp 0]  2 msec  3 msec  6 msec
 3  10.0.0.37 [MPLS: Label 16011 Exp 0]  3 msec 26 msec  9 msec
 4  10.0.0.26 [MPLS: Label 24035 Exp 0]  2 msec  3 msec  3 msec
 5  10.0.0.10 [MPLS: Label 24039 Exp 0]  3 msec  3 msec  2 msec
 6  172.16.2.5 [MPLS: Label 16011 Exp 0]  7 msec  7 msec  7 msec
 7  10.200.0.97 3 msec * 3 msec
RP/0/RP0/CPU0:SR-XRvDTR-9#

```

- XRvDTR-6, being the SR/LDP border router, swaps LDP label 24039 with node SID label 16011 towards SR node QFX5K-12.
- QFX5K-12 pops label 16011 upon receipt and sends traffic towards ASR9K1-11.

## Outcome when traffic passes through an SR/LDP border router:

Note that each hop is assigned a label. Traffic path goes through SR/LDP border routers, XRvDTR-10 and XRvDTR-6, which allows for SR-to-LDP forwarding and vice versa.



**Figure 16 – Traffic direction with SR and LDP Labels**

Looking at how XRvDTR-10 works when receiving it receives an SR labeled traffic, any incoming SR label is stitched to an LDP label.



\*Diagram below is a format taken out of the Segment Routing Part 1 book by Clarence Filsfils, Kris Michielsens, Ketan Talaulikar to illustrate operations of an incoming SR label transition to an LDP label within the SR POC Lab. Labels relevant to the SR POC Lab were replaced.

Function of an SR/LDP border router:

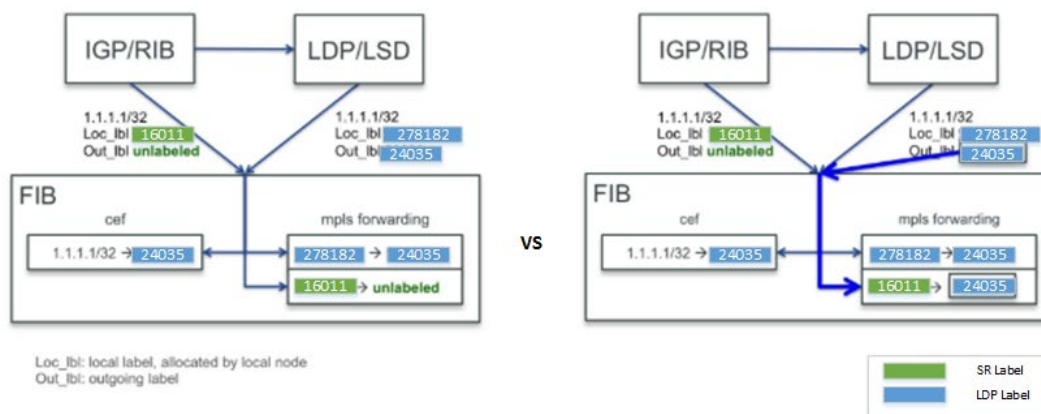


Figure 17 – SR to LDP Operations

If ASR9K1-16 is not SR enabled, and XRvDTR-10 is not SR/LDP enabled, the latter will not get an outgoing Prefix-SID label for ASR9K1-11 (10.0.1.11). Without knowing a label to reach 10.0.1.11, XRvDTR-10 provides “unlabeled” outgoing label in mpls forwarding entry for the 10.0.1.11 prefix.



If XRvDTR-10 is SR/LDP enabled, it receives a valid LDP label advertised by ASR9K1-16 on how to reach 10.0.1.11. The LSD provides the received label 24035 to the FIB and replaces the “unlabeled” entry.

Additional illustration to show the ingress and egress labels related to figure 18.

```
RP/0/RP0/CPU0:SR-XRvDTR-10#show mpls ldp forwarding | i "Prefix|In|10.0.1.11"
Fri Nov  9 23:17:29.644 UTC
Prefix          Label  Label(s)  Outgoing  Next Hop  Flags
                In    Out       Interface
10.0.1.11/32    278182 24035     Gi0/0/0/0.1016 172.16.2.13  G S R
```

```

RP/0/RP0/CPU0:SR-XRVDTR-10#show cef 10.0.1.11 flags
Fri Nov  9 23:06:18.666 UTC
10.0.1.11/32, version 106, labeled SR, internal 0x1000001 0x81 (ptr 0xd48e388) [1], 0x0 (0xd629948), 0xa28 (0xd822338)
leaf flags: owner locked, inserted

leaf flags2: LDP/SR merge requested, sr-pfx,
leaf ext flags: Prichange, EXTERNAL_REACH_LC, sr-mpIs, sr-pfx-sid,
Updated Nov  9 21:24:55.882
remote adjacency to GigabitEthernet0/0/0/0.1016
Prefix Len 32, traffic index 0, precedence n/a, priority 3
Extensions: context-label:16011
  via 172.16.2.13/32, GigabitEthernet0/0/0/0.1016, 23 dependencies, weight 0, class 0 [flags 0x0]
  path-idx 0 NHID 0x0 [0xdbf52c0 0x0]
  next hop 172.16.2.13/32
  remote adjacency
  local label 278182      labels imposed {24035}

```

There is no SR/LDP active because LDP provided a valid outgoing label to FIB.

Overall results show that interworking of SR/LDP is successful despite multiple vendor type used in a single topology.

### Outcome when SRGB is out of range of other nodes:

During testing, several mapping clients failed in installing the entries advertised by the SRMS. Configuring a prefix-to-SID mapping index on the SRMS that exceeds the SRGB of the mapping clients, will prevent the installation of the mapping server policies in the clients forwarding table. Table 7 below shows the configured SRGB of each node. When a prefix-to-SID mapping index is configured on the SRMS, it adds to the based SRGB of the mapping clients.

*Example:*

*SRGB base = 16000*

*Index = 17013*

*Prefix-to-SID mapping index = SRGB Base + Index = 16000 + 17013 = 33013*

When the Prefix-to-SID mapping index falls outside of the mapping client's SRGB range, the entry will not be installed in the forwarding table as depicted in **red** in the table below.

**Table 7 – Prefix-to-SID Index Comparison**

Nodes	SRGB	Index – 17013 (Prefix as seen on Clients)	Index – 1013 (Prefix as seen on Clients)	Index – 3013 (Prefix as seen on Clients)
XRvCER-22	16000 - 278142	33013	17013	19013
SR-XRvCRR1	16000 - 278142	33013	17013	19013
XRvCRR17	16000 - 278142	33013	17013	19013
XRvCRR18	16000 - 278142	33013	17013	19013
XRvCRR3	16000 - 278142	33013	17013	19013
XRvDTR-10	16000 - 278142	33013	17013	19013
XRvDTR-15	16000 - 18000	33013	17013	19013
XRvDTR-6	16000 - 18000	33013	17013	19013
XRvDTR-8	16000 - 18000	33013	17013	19013
MSC-4	19000 - 21000	52013	17013	22013
MSC-5	19000 - 21000	52013	17013	22013
ASR90012-11	16000 - 19000	33013	17013	19013

MX240-2	16000 - 19999	33013	17013	19013
SRQFX5100-12	16000 - 19999	33013	17013	19013
SRQFX10002-7	16000 - 20999	33013	17013	19013
XRvDTR-9	16000 - 18000	33013	17013	19013

- delete this page since all vendors have BCP on the interworking between SR and LDP with border node between SR and LDP regions run both LDP and SR and act as SRMR client

Example Symptom:

Node 6 receives the mapping server advertisements of the prefix-to-SID mappings but does not install them in the forwarding table. See below in (A) and (B) for example.

RP/0/RP0/CPU0:SR-XRVOTR-6#show isis segment-routing prefix-sid-map active-policy Mon Nov 5 19:33:32.686 UTC					RP/0/RP0/CPU0:SR-XRVCER-22#show isis segment-routing prefix-sid-map active-policy Mon Nov 5 19:34:58.620 UTC									
IS-IS default active policy					IS-IS default active policy									
Prefix	SID Index	Range	Flags		Prefix	SID Index	Range	Flags						
10.0.1.13/32	17013	2			10.0.1.13/32	17013								
10.0.1.16/32	17016	1			10.0.1.16/32	17016	1							
Number of mapping entries: 2					Number of mapping entries: 2									
RP/0/RP0/CPU0:SR-XRVOTR-6# RP/0/RP0/CPU0:SR-XRVOTR-6# RP/0/RP0/CPU0:SR-XRVOTR-6#show mpls forwarding Mon Nov 5 19:33:42.238 UTC					RP/0/RP0/CPU0:SR-XRVCER-22# RP/0/RP0/CPU0:SR-XRVCER-22# RP/0/RP0/CPU0:SR-XRVCER-22#show mpls forwarding Mon Nov 5 19:35:02.854 UTC									
Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched	Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched			
16001	unlabelled	SR Pfx (idx 1)	Gi0/0/0/0	10.0.0.9	1908169	16001	16001	SR Pfx (idx 1)	Gi0/0/0/0	10.0.0.62	7192176			
16002	unlabelled	SR Pfx (idx 2)	Gi0/0/0/0	10.0.0.9	0	16002	16002	SR Pfx (idx 2)	Gi0/0/0/0	10.0.0.62	416149			
16003	unlabelled	SR Pfx (idx 3)	Gi0/0/0/0	10.0.0.9	1902169	16003	16003	SR Pfx (idx 3)	Gi0/0/0/0	10.0.0.62	7192704			
16004	unlabelled	SR Pfx (idx 4)	Gi0/0/0/0	10.0.0.9	286640	16004	16004	SR Pfx (idx 4)	Gi0/0/0/0	10.0.0.62	0			
16005	unlabelled	SR Pfx (idx 5)	Gi0/0/0/0	10.0.0.9	0	16005	19005	SR Pfx (idx 5)	Gi0/0/0/1.522	172.16.2.15	0			
16007	unlabelled	SR Pfx (idx 7)	Gi0/0/0/0	10.0.0.9	0	16006	16006	SR Pfx (idx 6)	Gi0/0/0/0	10.0.0.62	7192704			
16008	unlabelled	SR Pfx (idx 8)	Gi0/0/0/0	10.0.0.9	1902429	16007	16007	SR Pfx (idx 7)	Gi0/0/0/0	10.0.0.62	0			
16009	unlabelled	SR Pfx (idx 9)	Gi0/0/0/0	10.0.0.9	1907562	16008	16008	SR Pfx (idx 8)	Gi0/0/0/0	10.0.0.62	7192320			
16010	unlabelled	SR Pfx (idx 10)	Gi0/0/0/0	10.0.0.9	2757662	16009	16009	SR Pfx (idx 9)	Gi0/0/0/0	10.0.0.62	7193280			
16011	16011	SR Pfx (idx 11)	Gi0/0/0/1.1206	172.16.2.5	1228742	16010	16010	SR Pfx (idx 10)	Gi0/0/0/0	10.0.0.62	7192896			
16012	Pop	SR Pfx (idx 12)	Gi0/0/0/1.1206	172.16.2.5	0	16011	16011	SR Pfx (idx 11)	Gi0/0/0/0	10.0.0.62	596251			
16015	unlabelled	SR Pfx (idx 15)	Gi0/0/0/0	10.0.0.9	1907503	16012	16012	SR Pfx (idx 12)	Gi0/0/0/0	10.0.0.62	0			
16017	unlabelled	SR Pfx (idx 17)	Gi0/0/0/0	10.0.0.9	1907444	16015	Pop	SR Pfx (idx 15)	Gi0/0/0/0	10.0.0.62	7193472			
16018	unlabelled	SR Pfx (idx 18)	Gi0/0/0/0	10.0.0.9	1907486	16017	16017	SR Pfx (idx 17)	Gi0/0/0/0	10.0.0.62	7193088			
16022	unlabelled	SR Pfx (idx 22)	Gi0/0/0/0	10.0.0.9	0	16018	16018	SR Pfx (idx 18)	Gi0/0/0/0	10.0.0.62	7191744			
24000	Pop	SR Adj (idx 0)	Gi0/0/0/0	10.0.0.9	0	33013	unlabelled	SR Pfx (idx 17013)	Gi0/0/0/0	10.0.0.62	1968			
24001	Pop	SR Adj (idx 2)	Gi0/0/0/0	10.0.0.9	0	33014	unlabelled	SR Pfx (idx 17014)	Gi0/0/0/0	10.0.0.62	1968			
24002	Pop	SR Adj (idx 0)	Gi0/0/0/1.1206	172.16.2.5	0	33016	unlabelled	SR Pfx (idx 17016)	Gi0/0/0/0	10.0.0.62	0			
24003	Pop	SR Adj (idx 2)	Gi0/0/0/1.1206	172.16.2.5	0	278143	Pop	SR Adj (idx 0)	Gi0/0/0/1.522	172.16.2.15	0			
RP/0/RP0/CPU0:SR-XRVOTR-6#					278144					Pop	SR Adj (idx 2)	Gi0/0/0/1.522	172.16.2.15	0
					278145					Pop	SR Adj (idx 0)	Gi0/0/0/0	10.0.0.62	0
					278146					Pop	SR Adj (idx 2)	Gi0/0/0/0	10.0.0.62	0
					RP/0/RP0/CPU0:SR-XRVCER-22#									

(A) Non-working Mapping Client

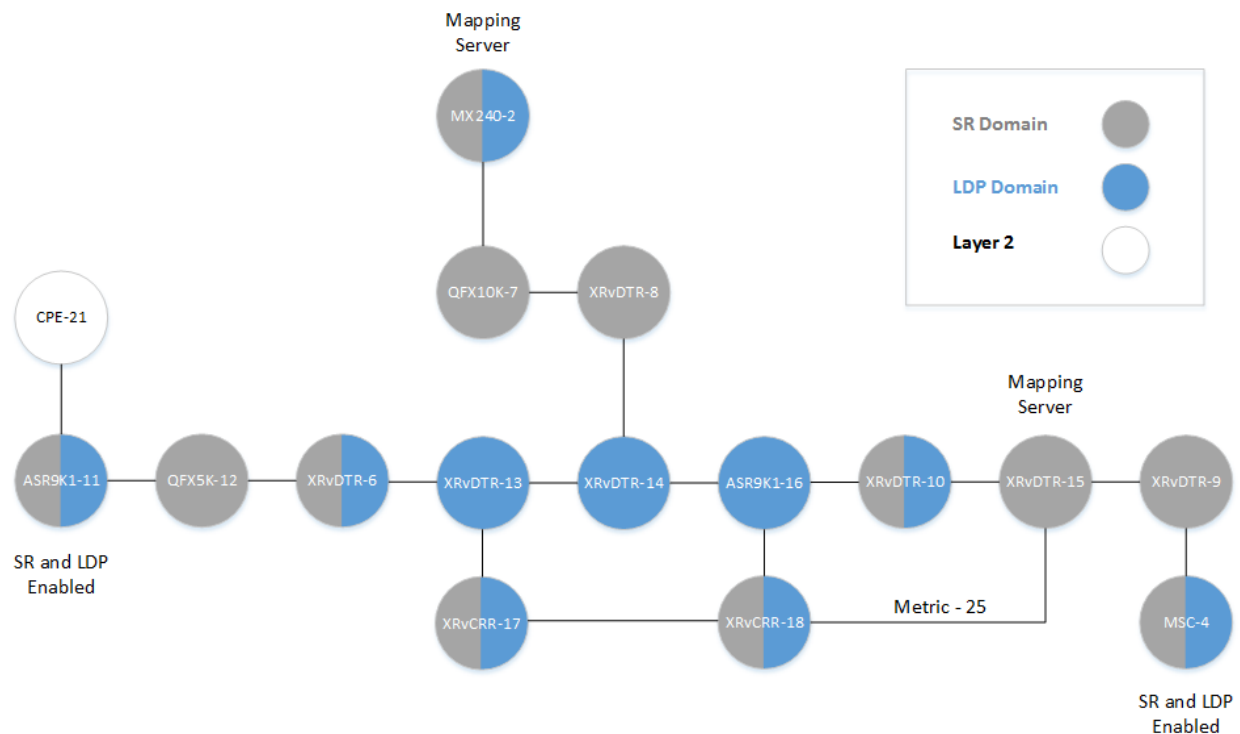
Mapping policy received, highlighted in blue. However, when looking at the forwarding table under “show mpls forwarding” command, the entries for SID index 17013 and 17016 from the mapping policy are missing.

(B) Working Mapping Client

Mapping policy received, highlighted in red. When looking at the forwarding table under the “show mpls forwarding” command, the entries for SID index are installed, highlighted in green.

Note: While the working mapping client installs the SRMS policy in the forwarding table, the outgoing label is unlabeled as highlighted in green due to the lack of SR/LDP border router. Test Case 2 performs this testing further.

### 3.6.3. Test Cases 3 – Redundant Mapping Server



**Figure 18 – Redundant Mapping Server Topology**

1. All devices within SR domain are SR enabled only. Exceptions are ASR9K1-11 and MSC-4 where L2VPN terminates and the SR/LDP border routers.
2. All devices within LDP domain are LDP enabled only.
3. XRvDTR-15 is the original mapping server, the rest of the nodes are mapping clients.
4. MX240-2 is added as redundant mapping server.

#### **Outcome with second mapping server:**

All devices receive the prefix-to-SID advertisement from the second SRMS, in this case MX240-2. The redundant mapping server is configured with one less entry. Note figure below shows backup-policy with

one less entry as received by the mapping client.

```
RP/0/RP0/CPU0:SR-XRVCRR1#
RP/0/RP0/CPU0:SR-XRVCRR1#show isis segment-routing prefix-sid-map active-policy
Thu Nov 1 18:36:13.716 UTC

IS-IS default active policy
Prefix          SID Index      Range      Flags
10.0.1.13/32    17013          2
10.0.1.16/32    17016          1
10.0.1.17/32    17017          1

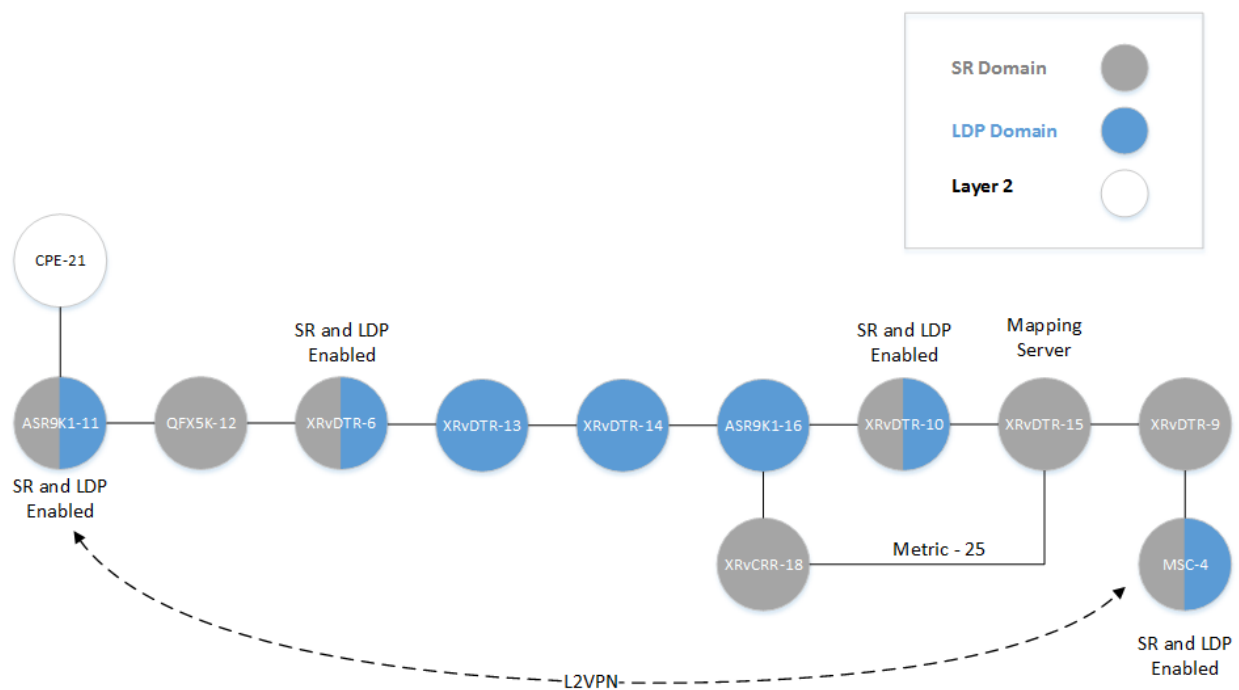
Number of mapping entries: 3

RP/0/RP0/CPU0:SR-XRVCRR1#show isis segment-routing prefix-sid-map backup-policy
Thu Nov 1 18:36:19.568 UTC

IS-IS default backup policy
Prefix          SID Index      Range      Flags
10.0.1.13/32    17013          2
10.0.1.16/32    17016          1

Number of mapping entries: 2
```

### 3.6.4. Test Case 4 – Services over SR/LDP Domain



**Figure 19 – L2VPN service over SR/LDP domain**

1. All devices within SR domain are SR enabled only. Exceptions are ASR9K1-11 and MSC-4 where L2VPN terminates for targeted LDP session and the SR/LDP border routers.
2. All devices within LDP domain are LDP enabled only.

3. XRvDTR-15 is the mapping server, the rest of the nodes are mapping clients.
4. MX240-2 is added as redundant mapping server now shown in figure 19.

**Outcome:**

L2VPN service successfully traversed across the LDP/SR domain.

### **3.6.5. Assessment of SRMS and SR/LDP Interworking with multiple vendors**

#### Cautionary Practices and Deployment

1. Prefix-to-SID-mapping range
  - Rather than configuring many prefix-sid-mapping entries for each node, it is easier to use the prefix, range, and start index command to represent a wide range of prefixes in a single entry. This works best if production loopback addresses are in contiguous fashion.
2. Any LDP only nodes must have all their direct SR neighbors to be an SR/LDP border router. Leaving out one node as SR/LDP could cause traffic loss should that node be the best IGP path after a primary failure.

#### Mapping Server and Client Risks

1. Risks of Conflicts and Overlapping prefix-to-SID-mappings
  - Forwarding loops can occur
  - Traffic blackholes
2. Traffic loss and service impact could occur if different vendor platforms interpret and perform conflict resolution differently as this could lead to inconsistent forwarding state across the network. (See Table 8 for vendor differences in accordance with IETF draft)
3. Vendor Z does appear to perform conflict resolutions for SID conflicts while it supports Prefix conflicts.
4. Troubleshooting commands are limited.
5. Using different SRGB values and ranges would require tracking of every node's SRGB since a mapping server could advertise an index that is outside of the receiving mapping client's SRGB range. Therefore it is better to use the common label space across three platforms.

**Table 8 – Conflict Resolution Preference Rules Comparison**

Node Platform X	Node Platform Y	Node Platform Z
Largest router-id (OSPF) or system-id (ISIS) is preferred	Largest router-id (OSPF) or system-id (ISIS) is preferred	
	Smallest area-id (OSPF) or level (ISIS) is preferred	
	IPv4 range is preferred over IPv6 range	IPv4 range is preferred
Smallest prefix length is preferred	Smallest prefix length is preferred	
Smallest IP address is preferred	Smallest IP address is preferred	Smallest IP address is preferred (IPv4 Only)
Smallest SID index is preferred	Smallest SID index is preferred	Smallest SID index is preferred
Smallest range is preferred	Smallest range is preferred	Smallest range is preferred
	First received range is preferred	

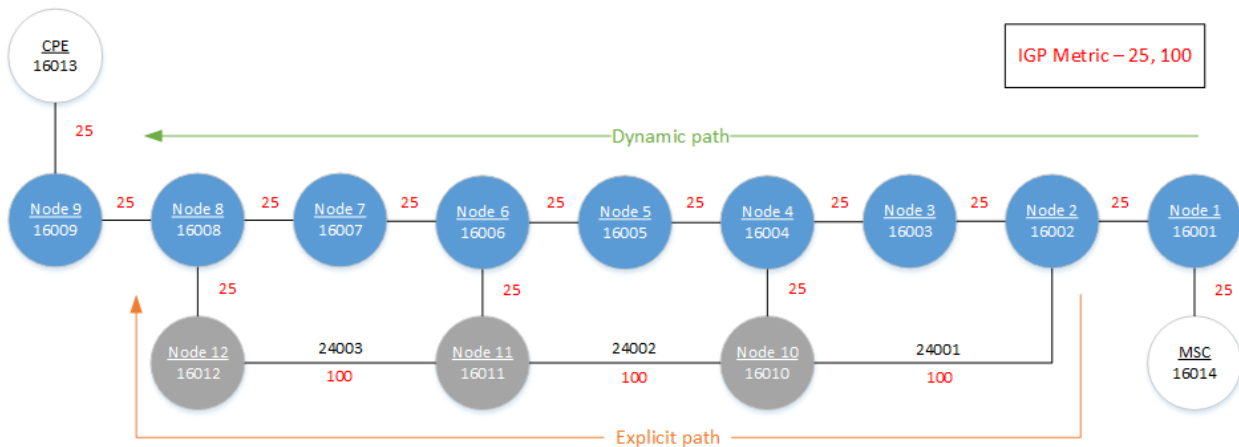
Mapping ranges conflict - <https://tools.ietf.org/html/draft-ietf-spring-conflict-resolution-05>

## 4. Design Considerations

### 4.1. MPLS MTU

While SR with IGP path forwarding with just the transport label do not impose additional overhead, it is the other mechanisms like Traffic Engineering or TI-LFA that should be taken into consideration when designing Maximum Transmission Unit (MTU) across the network. When using Traffic Engineering or TI-LFA, the MTU size could grow as the segment label stack increases.


With the MTU conditions of a baseline IP and layer 2 header with a single label, it will be of no if the network is configured with jumbo frames of 9000 or higher. For certain scenarios that require packets with higher MTU together with Traffic Engineering or TI-LFA, it will be necessary to evaluate the traffic path and the MTU across the network.



**Figure 20 – Dynamic and Explicit SR Path**

Based on the topology in figure 1, the following shows the multiple scenarios of how SR labels could impact MTU sizing. The scenarios below are built on the following assumptions where the payload is the typical customer service frame with C-VLAN tag.

#### Low MTU, dynamically forwarded




As every SR node is aware of other SR node's unique SID, only the destination node SID 16013 is imposed on the label stack. This scenario has no impact to Charter's MTU restrictions

#### Low MTU, explicitly forwarded




Steering this frame along the higher metric path calls for an increased segment list of six labels. Adjacency SIDs 24001, 24002, and 24003 were added to ensure traffic flows through nodes 11 through 12. Despite the growing label stack, the low MTU size of the service frame keeps the entire Ethernet frame well under 2000. This scenario has little to no impact to Charter's MTU restrictions.

#### High MTU, dynamically forwarded



Similarly to the low MTU size, dynamically forwarded scenario, only a single label is imposed for transport. While in this scenario the service frame is significantly larger in MTU, adding a single label still has little impact as it simply replaces a previously used transport label such as LDP or RSVP.

#### High MTU, explicitly forwarded



In a scenario where a high MTU service frame is steered, the label stack can potentially grow. This adds to the MTU size of the entire Ethernet frame. Adjacency SIDs 24001, 24002 and 24003 were added to ensure traffic flows from node 10 through 12, thus inflating the label stack.



**Recommendations:**

For best practice of MTU size considerations in SR-TE, it is recommended to account for the needed headroom against the maximum label stack and MTU restrictions across the network when setting up a service end to end. This maximum label stack or MSD is defined by each vendor platform. Please see MSD section for more information.

## Conclusion

The intent was to review the option for SR-TE as an alternate traffic engineering solution for business services. In the first phase of planning for an SR POC Lab, an audit of Charter's network was performed to better account for all types of major hardware and linecards used in production so that the POC is built according to deployment. During this audit, it was discovered that 50% of two vendor platforms' hardware/linecard would require replacement to support a full adoption of SR. However, with the successful testing of utilizing mapping servers in a mix of SR and LDP domains within a network, the impact of hardware replacement can be lessened. Caveat to the previous statement would be if the PE falls into the 50% replacement of hardware, deeming it difficult to depoly SR since the PE would be used as the SR headend node.

During the testing of SR, there were multiple instances where some features were not supported prompting several upgrades. Careful planning of what is required for a successful SR deployment is imperative. Even the smallest detail or a feature that didn't seem relevant at that point but is crucial to enabling the bigger feature can be easily overlooked. To date, there are still some vendors that may not have certain features available to align with another vendor that may be capable. SR may have been introduced for a few years, but some vendors are still catching up to the latest spec making it difficult to harmonize all vendors in a single deployment for the same feature. This is particularly important to providers that use multiple vendors in their network.

It is observed that some features could use a controller to alleviate some of the operational work such as visibility of MSD across the network and assist in reducing user errors. The major assist in having a controller would be path computation for optimal routing and constraints to meet the customer SLA.

Overall view of SR appears to be feasible given the multiple vendor platforms for baseline SR deployment. Though CAPEX would be seemingly high at initial SR rollout due to hardware support, the trade-offs are operational expense with less complexity to maintain the network with traffic engineering, lesser penalty fees for missing MTTR, and scalability to use a controller for full visibility of a network which includes efficient capacity planning, trend reporting, telemetry, and change control modeling. This also introduces opportunities for automation. An incremental or partial deployment of SR, or even a greenfield market, would be a better and more cost effective approach by Charter so as not to dive into a full investment while continuing to use existing transport methodology until SR is fully baked-in by all vendors. Charter is still pursuing further in-depth testing, particularly with the use of controllers and a higher subset of different interdomain networks operating together.

# Abbreviations

BMI	Base MPLS Imposition
CER	commercial edge router
CTBH	cell tower backhaul
IGP	Interior Gateway Protocol
LIB	label information base
LSD	label switching database
MPLS	multi-protocol label switching
MSC	Mobile Switching Center
MSD	Maximum SID Depth
MTTR	mean time to repair
MTU	maximum transmission unit
PCC	path computation client
PCE	path computation element
PCEP	Path Computation Element Protocol
PE	provider edge
POC	proof of concept
RSVP-TE	resource reservation protocol – traffic engineering
SLA	service level agreement
SID	segment identifier
SR	segment routing
SRGB	segment routing global block
SRMS	Segment Routing Mapping Server
SR-TE	Segment Routing Traffic Engineering
TI-LFA	topology independent – loop free alternate

## Bibliography & References

*PCEP Extensions for Segment Routing* - <https://tools.ietf.org/html/draft-ietf-pce-segment-routing-16>

*Segment Routing MPLS Conflict Resolution* - <https://tools.ietf.org/html/draft-ietf-spring-conflict-resolution-05>

*LDP Specification* - <https://tools.ietf.org/html/rfc5036>

*Segment Routing, Part 1*; Clarence Filsfils, Kris Michielsen, Ketan Talaulikar

# **IoT Device Energy Harvesting Technologies and Implementations**

A Technical Paper prepared for SCTE/ISBE by

Joe Rodolico, Principal Engineer, Comcast, SCTE/ISBE Member  
1800 Arch Street  
Philadelphia, PA 19103  
Joseph\_Rodolico@cable.comcast.com  
215-300-2516

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents	2
1. Introduction	3
2. Sources of Power	3
3. Transferring Environmental Energy to Power IoT Devices	3
3.1. Hybrid Mode (Ability to charge a battery)	3
3.2. Assist Mode (Ability to store limited energy)	4
3.3. Standalone Mode (Operated by harvested energy only)	5
4. Device/Sensor Design Considerations	5
5. Device/Sensor Examples	6
6. Mechanical Energy Harvesting	7
7. Solar Cell Energy Harvesting	7
8. Temperature Transfer Energy Harvesting	8
9. RF Energy Harvesting	8
10. Advanced Energy Harvesting	9
10.1. MEMS Pyroelectric Capacitor	9
10.2. Nano-Antennas (Nantennas) for Solar Energy Harvesting	9
11. Summary	10

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Charging circuit and battery for storing energy from a solar panel	4
Figure 2 – Example of a supercapacitor	4
Figure 3 – Solar-powered calculator	5
Figure 4 – ZigBee wall switch	6
Figure 5 – Example of a shake flashlight	7
Figure 6 – Remote control with solar cell	8
Figure 7 – Flashlight with Peltier tiles	8
Figure 8 – Harvesting RF energy to generate DC power	8
Figure 9 – MEMS Pyroelectric Capacitor Concept	9
Figure 10 – Two Examples of Nano-Antennas	9

# 1. Introduction

The proliferation of wireless home security and automation sensors and devices, i.e., the so-called Internet of Things (IoT), has advanced the need for improvements to battery life. The cable industry has deployed IoT sensor solutions for many years (in home security devices such as motion sensors, for example), but it needs to reduce costly service calls for replacing batteries and to work towards addressing global environmental concerns regarding the disposal of spent batteries.

Currently, the capacities of inexpensive chemical batteries are reaching their physical limits, and the concurrent drive to create ever-smaller devices calls for further innovations to extend battery-powered sensor life. Chief among these is energy harvesting. This paper focuses on promising energy harvesting methods and solutions that may be applied to IoT devices. Business and consumer drivers for improved efficiency sensor-powering solutions are also discussed.

Sample test results examining both performance and economic impacts are evaluated. Of particular importance is an assessment of user and system operator experiences, the perspectives of time and cost savings. Finally, recommendations are proposed that identify implementation opportunities for Multiple System Operators (MSOs).

## 2. Sources of Power

A variety of energy harvesting sources will be discussed in this paper, such as wind, rain, heat, cold, vibration, water flow in pipes, ocean currents, induction (temperature/electrical), motion, sunlight, and artificial light. Some of these energy sources can provide a direct power input, while others, depending on the application, can be utilized for target implementations. For example, solar energy can be converted directly into power to be used for long-term functions, such as a solar-powered light, or to top off batteries such as those used in solar calculators. Separately, the water flow in home plumbing pipes can be used to power a sensor (to send a status message) to indicate that water flow is present.

## 3. Transferring Environmental Energy to Power IoT Devices

There are three modes of providing harvested energy to a device:

- Hybrid Mode
- Assist Mode
- Standalone Mode


These modes are described more fully in the subsections that follow.

### 3.1. Hybrid Mode (Ability to charge a battery)

Hybrid mode extends battery life and, in some cases, allows for a battery of a smaller form factor to be used to provide for the same design functionality and expected battery life as a larger battery. The device operates either on energy harvested in real time or on harvested energy that is stored, depending on the amount of energy available for harvesting.

Using harvested energy for charging a battery has some drawbacks. First, there is a loss of energy to implement a step-up or step-down circuit; the device requires the addition of a rechargeable battery and a

charging circuit (an example is shown in Figure 1, below); and power for the device is dependent upon favorable environmental conditions (e.g., solar cells typically yield the most energy in sunny climates).



**Figure 1 – Charging circuit and battery for storing energy from a solar panel**

### 3.2. Assist Mode (Ability to Store Limited Energy)

Assist mode is an energy-efficient method where a smart circuit utilizes environmental energy. Such energy may be very low power or have infrequent availability, such as a solar cell utilized indoors that can assist with device function or “store” energy for future use in a capacitor or supercapacitor (i.e., a capacitor with extremely high capacity and low leakage loss over time).

It is challenging to use coin cell batteries in an assist-mode configuration since they have high internal resistance which causes a voltage drop during high-current demand, e.g., during transmissions to and from IoT devices. IoT SoCs (Systems on a Chip) are rated for a “cutoff voltage,” i.e., the voltage at which the device will cease to operate. The addition of a supercapacitor significantly reduces the voltage drop at high-demand periods, consequently raising the cutoff voltage of the device and thereby extending the battery life.

Assist mode can also be implemented with the addition of a small solar cell utilized to “top off” a battery for extremely low cost, for example, in basic handheld calculators or wearable devices where motion is converted to supplemental power for a device.



**Figure 2 – Example of a supercapacitor**

### 3.3. Standalone Mode (Operated by harvested energy only)

In standalone mode, only the harvested external energy powers the device. Examples of devices using this method are IoT devices without internal energy storage, such as low-cost Zigbee light switches (mechanical-to-piezoelectric) and solar-only powered IoT devices. Figure 3, below, shows a solar-only powered calculator which works well in bright light, but is challenged to operate effectively in low-light conditions.




Figure 3 – Solar-powered calculator

## 4. Device/Sensor Design Considerations

Circuits, including SoCs with wireless capability in IoT designs, are utilizing less power over time for increased functionality. Sleep currents in SoCs and microcontroller units (MCUs) are now operating in the nanoamps range, and the operational voltage required is decreasing over time; it is currently as low as around 1.7 V for contemporary SoCs. If a voltage requirement of less than 1.5 V is accomplished in the future, SoC operational voltage requirements would line up with commonly available 1.5 V batteries. Utilizing 1.5 V batteries would be of benefit since devices would generally not require boost circuits or batteries in series to support a SoC's operational voltage requirements.

As the illustration below indicates, efficiency of energy harvesting technologies is increasing while the wattage required for the electronics is decreasing, so we are likely to see smaller and more energy-efficient devices continue to develop.





The third leg of the stool for the above illustration is the most challenging for designers to implement: crafting software that goes beyond basic functionality to achieve a highly energy-efficient device. For example, turning off circuits when they're not in use and decreasing the CPU cycles for targeted functions, if supported by the SoC, decreases energy use and hence extends battery life.

## 5. Device/Sensor Examples

Device/sensor energy harvesting can be one of three types: conductive, mechanical, or radiant/electromagnetic energy. Conductive energy harvesting can be achieved through a temperature differential between two surfaces. Mechanical energy harvesting is achieved through motion/vibration. Radiant energy can come from the Sun as solar energy and from the Earth's natural electromagnetic field as Schumann resonances. Unfortunately, the energy harvested via Schumann resonances is impractical for sensors since the amount of energy collected would be very small and would require a very large antenna to collect.

A push-button Zigbee wall switch (Figure 4) utilizes mechanical energy, through the motion of pushing a switch, to generate enough energy to send out a Zigbee message. This technology is excellent for sending a simple one-way message. However, without having the power to receive an acknowledgement or retry message without another mechanical action limits the use of this type of harvesting method.



**Figure 4 – Zigbee wall switch**

## 6. Mechanical Energy Harvesting

Mechanical energy can be harvested from motion or vibration and converted into electrical energy for devices/sensors. Utilizing Peltier technology, applied vibration provides electrical power, such as in a Peltier device mounted to a vibrating motor. For example, door sensors can utilize the mechanical motion of the door opening and closing to power themselves. In the “shake flashlight” (Figure 5), magnets and wire coils combine with mechanical motion to induce a current that generates enough power to light an LED.



**Figure 5 – Example of a shake flashlight**

## 7. Solar Cell Energy Harvesting

Solar cell devices are ideal for outdoor areas where sunlight is prevalent for approximately 12 hours a day. Solar cells can also be utilized indoors to provide supplemental power to a device. Obviously, however, there are limitations for indoor use (indoor power is generally 1/1000 that of outdoor power), and efficiency is highly dependent on the application and placement.

For IoT devices, two-way and one-way communications may be utilized for applications. For simple one-way communicating devices with short messages, energy harvesting is ideal in standalone mode where it is the sole source of power, such as with the Zigbee light switch. For two-way communications, especially for security systems where there is frequent messaging, energy harvesting is more applicable as a supplemental energy source.


Most IoT devices tend to be black, dark gray, white, or off-white in color. Adding an energy harvesting solution such as a solar cell may change the aesthetics of the device, as solar cells are usually copper or black in tone with an overlaying metal grid. IoT device manufacturers strive to integrate the energy harvesting technology into the device with the best aesthetics possible. The remote control with a solar cell (Figure 6) uses ambient light in the home, when available, to provide supplemental power and charge the battery. With this approach, either the life of the remote control’s primary batteries can be extended or the number of primary batteries in a remote control can be reduced. The addition of a non-replaceable rechargeable battery or a supercapacitor storage mechanism is typically utilized to accomplish the aforementioned applications. The customer experience would thereby be enhanced by the extension of the battery replacement period and/or the overall reduced cost of primary cell replacement.



**Figure 6 – Remote control with solar cell**

## 8. Temperature Transfer Energy Harvesting


The concept of harvesting heat energy is used in the flashlight shown in Figure 7. This device uses Peltier tiles to convert heat from the hand holding the flashlight to electricity to power the light. Peltier tiles generate electricity when the temperature differential of its top and bottom layers is approximately 5° C.



**Figure 7 – Flashlight with Peltier tiles**

## 9. RF Energy Harvesting

Harvesting RF-transmitted energy depends on the level of transmissions and the conversion of RF energy to actual utilized power, such as paralleling the antenna and energy harvesting from the RF power provided by the antenna on transmission. Directing such power to a capacitor to be utilized for future transmission or other device is one harvesting method.



**Figure 8 – Harvesting RF energy to generate DC power**

## 10. Advanced Energy Harvesting

Advanced energy harvesting technologies may be cost- and implementation-challenged, but they are noted here for their promising potential. The examples MEMS pyroelectric capacitors and nano-antennas are summarized in the subsections below.

### 10.1. MEMS Pyroelectric Capacitor

A MEMS (microelectromechanical systems) pyroelectric capacitor is a technology for harvesting residual heat to power devices/peripherals, such as a USB port or HVAC system vent position.




Figure 9 – MEMS Pyroelectric Capacitor Concept

### 10.2. Nano-Antennas (Nantennas) for Solar Energy Harvesting

Nano-antennas achieve close to 90% efficiency compared to 10-20% for silicone-based solar cell energy harvesting. Using silver in nantennas produces higher efficiency and allows fine-tuning the dipole dimensions. Below are two examples of nantennas (Figure 10).

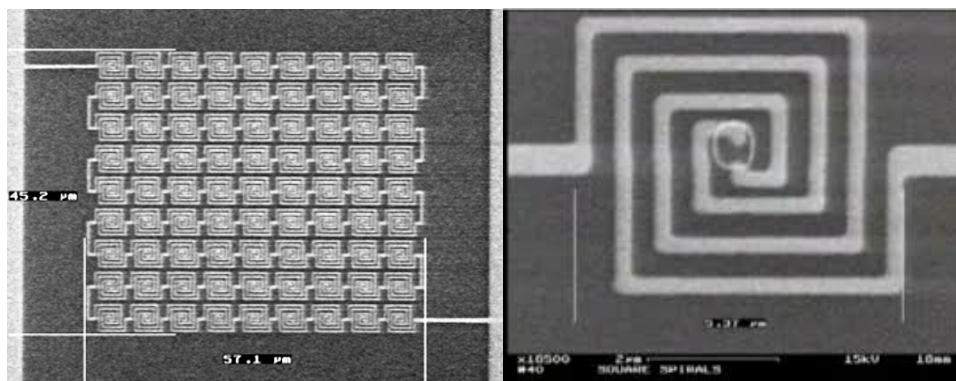


Figure 10 – Two Examples of Nano-Antennas

## 11. Summary

Battery technology capacities per form factor for IoT devices have stabilized somewhat for now. The IoT battery-life design challenge must be met with a coordinated multi-pronged approach that matches the design choice with the environmental energy available for harvesting. Electronics engineers need to choose low-energy electronics/designs and, even more importantly, craft smart software for energy conversation wherever possible.

Improving IoT battery life reduces overall battery change-out costs, resulting in fewer service calls (thus reducing operator cost) and increased customer satisfaction. Additionally, the extension of primary cell battery life or the reduction in the number of primary cells utilizing energy harvesting furthers the extensive opportunities for a “green initiative” for the cable industry.

# **Methods to Maximize IoT Battery Life**

A Technical Paper prepared for SCTE/ISBE by

Joe Rodolico, Principal Engineer, Comcast, SCTE/ISBE Member  
1800 Arch Street  
Philadelphia, PA 19103  
Joseph\_Rodolico@cable.comcast.com  
215-300-2516

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents	2
1. Introduction	3
2. Battery Procurement	3
2.1. Origin of Manufacture	3
2.2. Pricing	4
2.3. Storage	4
3. Battery Internal Resistance	4
4. Advanced Battery Designs	6
5. Storage and Disposal	6
6. Safety	6
7. Summary	7
8. Abbreviations	7

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Battery Schematic Illustrating Internal Resistance	4
Figure 2 – Variation of Voltage and IR over Time for Alkaline (top) and Lithium (bottom) Batteries	5
Figure 3 – Analysis of Battery Capacity over Time for Different Manufacturers/Products	6

# 1. Introduction

Addressing Internet of Things (IoT) energy consumption has a number of potential benefits, from the obvious environmental value of reducing the number of disposed batteries, to the customer experience benefits and resource savings that can be associated with fewer service calls. The concurrent migration to smaller devices requires action to extend the sensor battery life; chief among these is to improve battery selection techniques.

On the surface, choosing a battery for an IoT sensor may seem like a simple task, but in reality it involves a complex process of discovery, selection, and testing as outlined in this paper. Typically, replaceable lithium coin cell, cylindrical alkaline, and lithium batteries are utilized, as are other flat battery designs. While some IoT sensors use special high-cost chemical or rechargeable batteries, this paper is focused on maximizing battery life for low-cost wireless IoT sensors with replaceable batteries, deployed by the cable industry.

Simulated implementation for both performance and economic aspects are evaluated. Among important factors are assessing subscriber and system operator experiences from time, cost-savings, safety, and disposal perspectives. Recommendations are proposed to identify product design and deployment opportunities for Multiple System Operators (MSOs).

## 2. Battery Procurement

The objective for choosing a battery for an IoT sensor design is to find either the lowest-cost battery or, in some instances, to find a name-brand battery that meets your requirements. Battery manufacturers consistently advertise that their battery is superior and longer lasting, but the reality may be that there is only a marginal improvement of one over another. Also, often what is advertised differs drastically from what is delivered.

There are four false assumptions about procuring batteries: (1) major manufacturers have the highest-quality batteries; (2) lithium batteries are more powerful than alkaline batteries; (3) brand-name batteries cost more than no-name batteries; and (4) buying batteries from major vendor websites or distributors is good enough.

### 2.1. Origin of Manufacture

Major battery manufacturers may procure their batteries from their own factories or from third-party factories. From either source, the case and look of the battery can be identical. There are two methods to determine the quality of the battery. The first is to contact the battery vendor directly, which is usually performed under a non-disclosure agreement (NDA), and ask in what factory (or factories) the battery in question is manufactured and which factory carries out quality control; then confirm the country of manufacture on the battery label. If the battery manufacturer has multiple factories in China, for example, modify the procurement agreement to state that only batteries of this type are acceptable from this factory, unless notified and agreed to otherwise.



## 2.2. Pricing


Typically, information on battery pricing is found on bulk battery distributors' websites, but this pricing does not reflect the discounts that a large company would be provided for volume purchasing for a new IoT device and other products for which the company is currently procuring. Public website pricing is typically higher than what can be negotiated by a bulk customer, so it is suggested that pricing should become an exercise of discussions between major battery manufacturers and, in some cases, lower-cost manufacturers for cost comparison. Major manufacturers may sometimes charge the same price or a price within a couple of cents per unit as a lower-cost manufacturer, so that cost difference becomes a less important factor than other considerations when deciding on which battery to use.

## 2.3. Storage

Vendors must be able to provide documentation that their batteries are stored at a temperature close to 70 °F in a humidity-controlled environment, and that the batteries provided are fresh and not from dated warehouse stock. For manufacturer devices shipped, it is strongly suggested to require batteries be dated no more than 3 months from the date of manufacture as battery capacity generally degrades about 10% per year (see Figure 3). Recently, battery manufacturers have been marking batteries with an expiration date rather than a date of manufacture. For example, in 2016, a battery with a 2022 expiration date and a 10-year shelf life is likely to be 4 years old, and thus have an estimated 40% reduction in capacity. Battery manufacturers that have fully automated factories usually have less contamination; as a result, non-manual production and a better level of quality therefore exists. Battery contamination during manufacture, such as a dielectric marred by micro-pinholes, will result in a shorter battery lifespan than rated.

## 3. Battery Internal Resistance

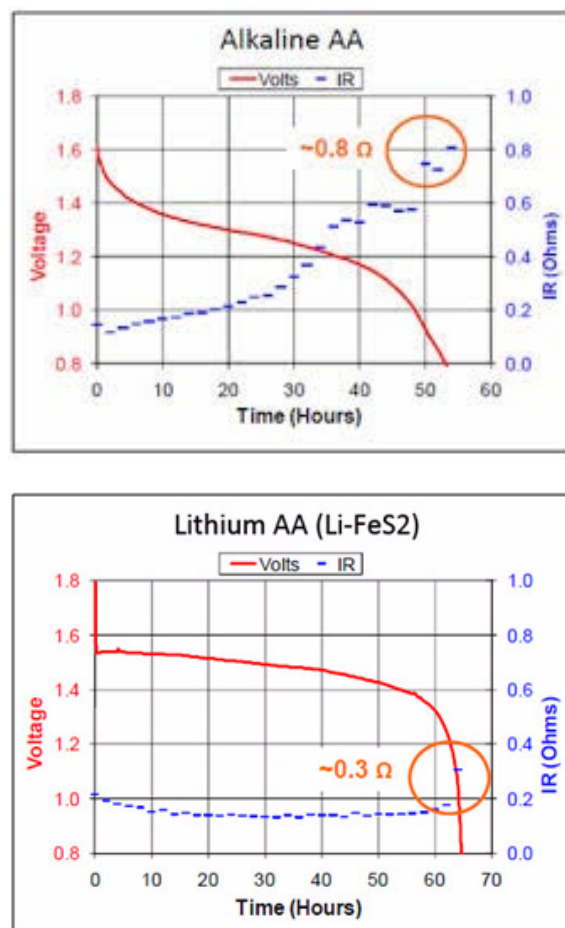
All batteries will have some amount of internal resistance (IR; see battery schematic in the figure below). For example, a CR2450 cell will have an IR of 20  $\Omega$  or less. IR can be measured by calculation or using a test meter.



**Figure 1 – Battery Schematic Illustrating Internal Resistance**

Internal resistance generally varies as the inverse of the battery voltage drop over the life of the battery, as shown in the below diagram. Battery charts such as those shown below typically show “light loading” over time, while IoT devices on transmission, for example, exhibit a high level of instantaneous current draw which results in a lower voltage provided by the battery, depending on the battery maximum current draw specification and the internal resistance of the battery. Hence, there is a need to select device circuits with the lowest current draw for a max-loading operation such as wireless transmissions. For example, for a 50 mA instantaneous draw, a 3 V lithium coin cell may drop 400 mV, while two 1.5 V AA alkaline batteries in series (i.e., 3 V) may only drop 100 mV, mostly due to the much lower internal resistance of the alkaline battery.

In some cases, depending on current drawn, alkaline batteries can outperform lithium coin cell batteries. Performing a load test with both lithium and alkaline batteries is an improved method to determine the best battery for a particular application.



**Figure 2 – Variation of Voltage and IR over Time for Alkaline (top) and Lithium (bottom) Batteries**

## 4. Advanced Battery Designs

New battery chemistries are more resilient against degradation over time, particularly in a high-heat environment. However, these batteries have the trade-offs of smaller capacity and slightly higher cost. Modern smoke alarms are using such batteries for a sealed 10-year life span. (Note: While the battery life of these devices is longer than that of standard batteries, the 10-year life span assumes a life spent in standby. Battery life is shortened when alarms occur.) Taking into account customer satisfaction, truck-roll costs, and disposal related to battery replacement, utilizing the new chemistry, longer-life batteries is well worth the small increase in cost.




Figure 3 – Analysis of Battery Capacity over Time for Different Manufacturers/Products

## 5. Storage and Disposal

It is suggested that only a limited number of batteries and devices with batteries be carried in trucks and that stock be regularly turned over, as trucks exposed to the extreme summer heat will reduce battery life. (Extreme heat degrades batteries; cold preserves batteries as long as the temperature is not at or below the manufacturer's extreme cold battery rating.)

Batteries removed from IoT sensors should be wrapped/stored as recommended by the battery manufacturer and returned for recycling or proper disposal as advised by the battery manufacturer. One benefit of battery recycling is that some new batteries are composed of a small amount of recycled battery material, and are marked as such. Thus, battery recycling is encouraged to promote a “greener” environment for battery usage.

## 6. Safety

Accidental shorting of batteries must always be a consideration. Some battery types, but not all manufacturers within that type, incorporate a built-in positive temperature coefficient (PTC) circuit for additional protection. Other recommended protective measures include: the addition of battery protection circuits in the device design, mechanical reverse battery protection, and wrapping the batteries individually in plastic by the manufacturer. It is strongly suggested to contact the battery manufacturer for their complete list of battery safety recommendations.

## 7. Summary

To maximize IoT battery life the following steps are required:

1. Select device designs with the lowest battery cut-off voltage, lowest sleep current, lowest max-load draw, and adequate low equivalent series resistance (ESR) capacitance across the battery circuit.
2. Review the battery candidate's data sheets and select advanced battery chemistries with the lowest internal resistance, widest operating temperature range, largest capacity, and lowest capacity loss over storage time.
3. Pre-screen a large sample of batteries to find the product with the most consistent and lowest internal resistance.
4. Test and compare candidate batteries over a temperature range, loading down to the cut-off voltage of the device.
5. Follow the procurement procedures in this paper for the final battery vendor selection criteria.

When the steps above are taken for the IoT design and battery selection, the result will be optimal battery life (reduced overall battery change-out cost), fewer service calls (reduced operator cost), and increased customer satisfaction (customer retention benefit).

To truly maximize IoT device battery life beyond this paper's recommendations, device designers must employ smart software to conserve battery life: sleep the device for the maximum time the application will allow and ensure device functions take as little time as possible.

## 8. Abbreviations

ESR	equivalent series resistance
IoT	Internet of Things
IR	internal resistance
PTC	positive temperature coefficient

# On The Performance Of CBRS Fixed Wireless Access: Coverage And Capacity Field Study

A Technical Paper prepared for SCTE•ISBE by

**Mohamed Daoud**

Principal Engineer Wireless R&D  
Charter Communications  
6360 S Fiddlers Green Circle, Greenwood Village, CO 80111  
+13123639864  
Mohamed.Daoud@Charter.com

**Matthew Hubbard**

Principal Engineer Wireless R&D  
Charter Communications  
6360 S Fiddlers Green Circle, Greenwood Village, CO 80111  
+12142640100  
Matthew.G.Hubbard@Charter.com

**Rajeev Aggarwal**

Principal Engineer Wireless R&D  
Charter Communications  
6360 S Fiddlers Green Circle, Greenwood Village, CO 80111  
+17205369397  
Rajeev.Aggarwal@Charter.com

**Hossam Hmimy**

Sr. Director Wireless R&D  
Charter Communications  
6360 S Fiddlers Green Circle, Greenwood Village, CO 80111  
+17205369396  
Hossam.Hmimy@Charter.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Abstract.....	5
Introduction .....	5
1.1.    Citizen Broadband Radio Services CBRS .....	5
1.2.    Fixed Wireless Access (FWA).....	6
1.2.1.    CBRS FWA Opportunities .....	6
1.2.2.    CBRS FWA Challenges .....	7
1.3.    Charter CBRS FWA Trials and Inverstigations .....	7
1.4.    Paper Structure .....	7
Network Architecture .....	8
Technology Description .....	11
1.1.    LTE Equipment Specification .....	11
Test Results .....	12
1.1.    Denver Trial – Hilly Terrain .....	12
1.1.1.    Coverage in Hilly Terrain.....	13
1.1.2.    Capacity in Hilly Terrain .....	19
1.2.    Tampa Trial – High Foliage.....	27
1.2.1.    Coverage in High Foliage – LTE CBRS .....	27
1.2.2.    Coverage in High Foliage – Proprietary CBRS and 5GHz.....	31
1.3.    Coldwater Trial – Snow and Rain.....	33
1.3.1.    Snow and Rain Effect on CBRS .....	33
1.3.2.    Beamforming Gain – Coverage Test .....	36
1.3.3.    CBRS+5GHz Radio Prototype.....	37
1.4.    Lexington Trial – User Experience .....	39
Conclusion .....	39
Abbreviations.....	40
Bibliography & References .....	41

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – CBRS Band 48 Sharing Framework .....	6
Figure 2 - LTE Network Architecture .....	8
Figure 3 - Charter Testing Network Architecture.....	9
Figure 4 - Charter Custom Built Test Van.....	10
Figure 5 - Denver Coverage Trial Fixed Locations.....	13
Figure 6 - Denver Coverage Trial DL Throughput.....	14
Figure 7 - Denver Coverage Trial UL Throughput.....	14
Figure 8 - Denver Test Point 3 Elevation Profile .....	15
Figure 9 - Denver Test Point 5 Elevation Profile .....	16
Figure 10 - Proprietary Radio DL Throughput at 25 ft CPE .....	17

Figure 11 - Proprietary Radio UL Throughput at 25 ft CPE .....	17
Figure 12 - Proprietary Radio SNR Values at 25 ft CPE .....	18
Figure 13 - Proprietary Radio RSSI Values at 25 ft CPE .....	18
Figure 14 - Denver Field Trial First Sector.....	19
Figure 15 - Participant With Major Fresnel Zone Blockage .....	20
Figure 16 - Participant With Minor Fresnel Zone Blockage .....	20
Figure 17 - Participant First Sector Throughputs .....	21
Figure 18 - Denver Field Trial Second Sector.....	22
Figure 19 - Participants Second Sector Throughput .....	23
Figure 20 - Sector Throughput with 10 CPE's at 256QAM .....	24
Figure 21 - Sector Throughput with 10 CPE's at 64/16QAM .....	25
Figure 22 - Adding 10 CPE's at 256QAM to Existing Sector - Sector Throughput = 227 Mbps .....	26
Figure 23 - Adding 10 CPE's at 64QAM to Existing Sector - Sector Throughput = 152 Mbps .....	27
Figure 24 - RSRP at 12ft and 25ft CPE Height.....	28
Figure 25- SNR at 12ft and 25ft CPE Height .....	29
Figure 26 - DL Throughput at 12ft and 25ft CPE Height .....	29
Figure 27 - UL Throughput at 12ft and 25ft CPE Height .....	30
Figure 28 - Modulation at 12ft CPE Height .....	30
Figure 29 - Modulation at 25ft CPE Height .....	31
Figure 30 - Proprietary CBRS and 5GHz Throughput Results for 12ft CPE .....	32
Figure 31 - Proprietary CBRS and 5GHz Throughput Results for 25ft CPE .....	32
Figure 32 - Effect of Rain on RSRP at Different CPE Heights – Cell Near .....	33
Figure 33 - Effect of Rain on SNR at Different CPE Heights – Cell Near.....	34
Figure 34 - Effect of Rain on Throughput at Different CPE Heights – Cell Near .....	34
Figure 35 - Effect of Rain on RSRP at Different CPE Heights – Mid-cell.....	35
Figure 36 - Effect of Rain on SNR at Different CPE Heights – Mid-cell .....	35
Figure 37 - Effect of Rain on Throughput at Different CPE Heights – Mid-cell.....	36
Figure 38 - Effect of Beamforming on Throughput at Mid-cell.....	36
Figure 39 - Effect of Beamforming on Throughput at Cell Edge .....	37
Figure 40 - Prototype Radio Throughput at 12ft Receiver Height.....	38
Figure 41 - Prototype Radio Throughput at 25ft Receiver Height.....	38
Figure 42 - Various Internet Devices Running Simultaneously.....	39

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – LTE eNB Equipment Specifications .....	11
Table 2 - LTE CPE Equipment Specifications .....	11
Table 3 - Proprietary Equipment Specifications.....	12
Table 4 - Denver Test Points Distance from Base Station .....	13
Table 5 - Denver Trial First Sector Participants RF conditions and Throughput.....	20
Table 6 - Denver Trial Second Sector Participants RF conditions and Throughput .....	22

Table 7 - Adding 10 CPE's at 256QAM to Existing Sector .....	25
Table 8 - Adding 10 CPE's at 64QAM to Existing Sector .....	26
Table 9 - Tampa Test Points Distance from Base Station.....	28
Table 10 - Tampa Test Points for Proprietary Equipment .....	31



# Abstract

The rural broadband gap in the U.S. is real and can't be ignored. It's estimated that millions of Americans in rural communities lack broadband internet access, or at best they are underserved with limited connectivity below the Federal Communications Commission (FCC) definition of broadband. This broadband gap is affecting various aspects of rural community's life style that ranges from lack of internet in schools to less attractive investment opportunities, and lower standard of living.

Charter, being a leader in broadband connectivity, found in the Citizens Broadband Radio Service (CBRS) band a good opportunity to provide rural broadband internet access cost effectively up to 40 miles away from its Fiber network service area.

In the past two years, Charter conducted extensive field testing and studies on CBRS and 5GHz for Fixed Wireless Access (FWA) in different markets to cover varieties of terrain, and environmental impact on the signal propagation. The testing conducted in an end-to-end setup from Customer Premises Equipment (CPE) to the core network.

In this paper FWA field testing results, best practices, and some techniques to help expand the coverage for rural FWA are presented.

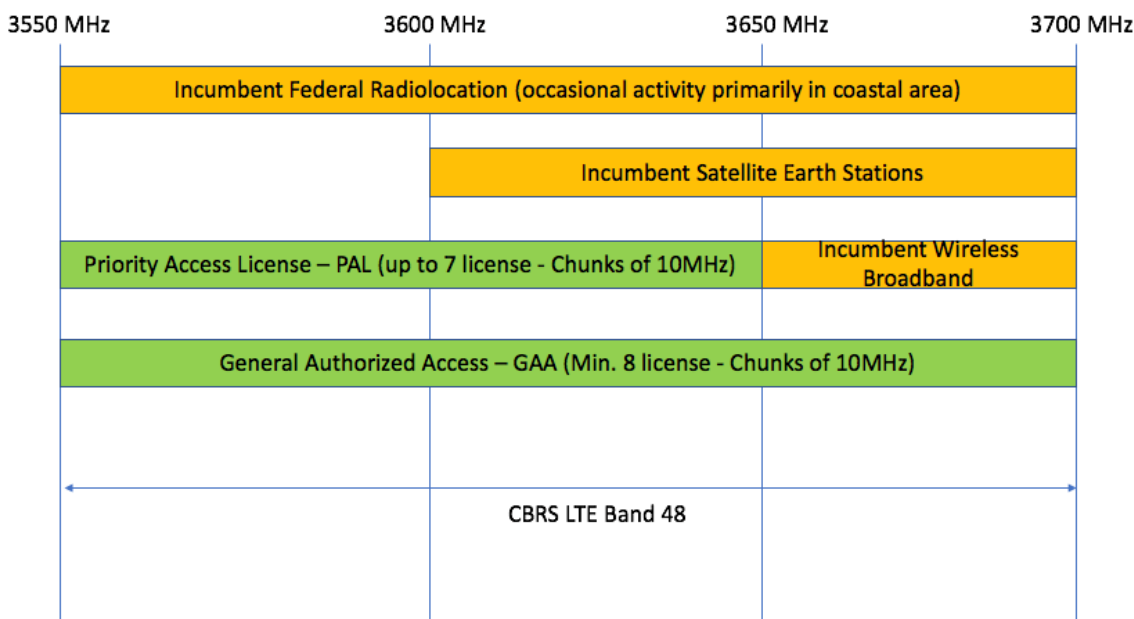
## Introduction

### 1.1. Citizen Broadband Radio Services CBRS

CBRS isn't a new band, rather it's a new framework of using the 3.55 to 3.7 GHz frequencies. Currently the use of 3.55 to 3.7GHz also known as Band 48 is limited due to the existence of legacy users like military, satellite earth stations, and some Wireless Internet Service Providers (WISPs). The FCC made a historical decision in 2015 to expand the use of the 3.5GHz band to other users including but not limited to broadband operators, venues, public and private entities, enterprises, startups, and service providers. Historically, spectrum has been a scarce resource and is auctioned off to service providers for billions of dollars. The FCC goal is to democratize the band and decrease the barriers of entry for any entity wanting to use the 3.5GHz band, this is expected to spur innovation and keep the United States on the leading edge of telecommunications technology.

The CBRS spectrum sharing framework allows users to access the full 150MHz on dynamic basis based on priority tiers. Legacy users retain the right to use the spectrum whenever they need it. Priority Access License (PAL) holders collectively retain access to as much as 70 MHz of spectrum in a license area, with up to 40 MHz of spectrum per PAL holder. They must protect legacy users from harmful interference, but they receive protection from interference by General Authorized Access (GAA) users. GAA users collectively have access to spectrum not being used by legacy users and PAL holders in a given area, which is as much as 150 MHz of bandwidth. GAA users do not receive interference protection from legacy users or PAL holders. Spectrum Access Systems (SASs) have been created to dynamically monitor and authorize use of specific spectrum resources for PAL and GAA users based on this priority order, using geolocation databases and policy management servers[1].

The CBRS sharing framework enables multiple users to share the CBRS spectrum, and while doing so, each has use of its assigned channel based on the priority of the tier the user is in. The SASs authorize users to use the spectrum and ensure that sharing among users is fair and try to decrease interference as much as possible.



**Figure 1 – CBRS Band 48 Sharing Framework**

## 1.2. Fixed Wireless Access (FWA)

CBRS spectrum sharing has many use cases, in this paper we are focusing on FWA use case due to its relevance and importance to Charter as a leading broadband provider in the United states.

The first logical target market for FWA is rural households especially that millions of them lack broadband internet speeds. The FCC made it one of its priorities to close the digital gap in the country using various technologies. CBRS band is positioned to play a major role in closing this digital gap. The idea of FWA is simply to deliver high speed internet using wireless technologies, in the case of CBRS it's LTE or 5G or a proprietary technology. A radio is installed on a tower that delivers high speed wireless internet to CPE attached on the outside of the customers' house. The CPE has to be oriented towards the radio and acts like a cellular device in the sense that it needs a Subscriber Identification Module (SIM) card to operate.

### 1.2.1. CBRS FWA Opportunities

To understand the business case of FWA CBRS we have to know what's meant by rural community and broadband service. While there are many definitions for rural the U.S Census bureau defines rural as area with population density less than 100 people per square mile[2]. According to the FCC, broadband speeds are 25Mbps downlink and 3Mbps uplink[3], it's expected that the definition of broadband will increase over time to account for all the new applications and use cases requiring higher data rates.

Charter is well positioned to capture this FWA opportunity because of the vast fiber rings Charter owns cross country, the big number of Charter owned towers, and the Charter – Spectrum brand equity. The fiber

will be used for backhaul while the towers used to install radios on them. The existence of fiber and towers in charter's portfolio will decrease the cost and time to roll out FWA service.

### **1.2.2. CBRS FWA Challenges**

As with any new opportunity some challenges arise. The challenges for the CBRS FWA are creating the appropriate framework for spectrum sharing which the CBRS Alliance (OnGo) and WinnForum are addressing. Another challenge is the spectrum management by SAS, with incumbent protection and coordination for GAA and for the PAL. Ecosystem is a challenge in terms of equipment and vendors availability and variability but we are seeing big improvements with many companies interested in bringing forward CBRS equipment and competitive roadmaps. The CBRS Alliance OnGo certification created a framework for vendors certification that helps in accelerating the equipment ecosystem. Another major challenge is the limited reach of the 3.5GHz signal.

It's worth noting that CBRS FWA is not the only answer to close the digital gap and connect all underserved in this country, rather it should be one of the tools available for service providers to extend their network reach. Depending on the house hold density different options are more economic viable. For example, in dense area with high household population density fiber is more appropriate means to deliver broadband internet and in areas of very low population density and very few households solutions based on sub GHz like TV White Space (TVWS) would be more viable. Operators should have a toolbox of solutions to serve their customers and CBRS FWA should be in this toolbox.

### **1.3. Charter CBRS FWA Trials and Inverstigations**

For the past two years Charter has been running trials across the country to investigate the opportunities and challenges of the CBRS band. The idea is to provide FWA in rural areas meeting the FCC definition of broadband. FWA isn't here to replace fiber rather to complement it.

Multiple technology including standards based and proprietary in both CBRS band and 5GHz using different morphologies have been evaluated.

The FWA trials were performed in the hilly suburbs of Denver, high foliage areas north of Tampa, mixed terrain of Bakersfield, snowy conditions in Northern Michigan, and rural farms of Kentucky. Proving that 25Mbps on the downlink and 3Mbps on the uplink can be achieved in a cell radius equal to and greater than 5 miles using Long Term Evolution (LTE) and wireless proprietary technologies.

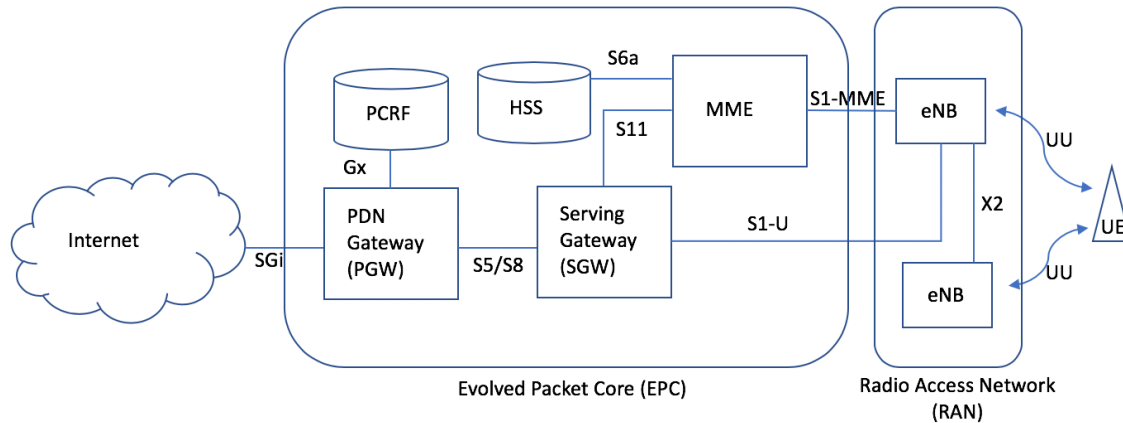
Multiple features have been tested including carrier aggregation, multi user Multiple Input Multiple Output (MIMO), and beam forming for cell performance and capacity.

### **1.4. Paper Structure**

The rest of this paper is organized as follows; section 1 the FWA network architecture is presented, section 2 details the technology and equipment used in the trials. In Section 3, coverage and capacity results are presented for testing in hilly terrain, high foliage, and snow and rain conditions with detailed analysis and discussion on findings. Section 4 concludes with the lessons learned from Charter FWA testing, followed with some recommendations.

# Network Architecture

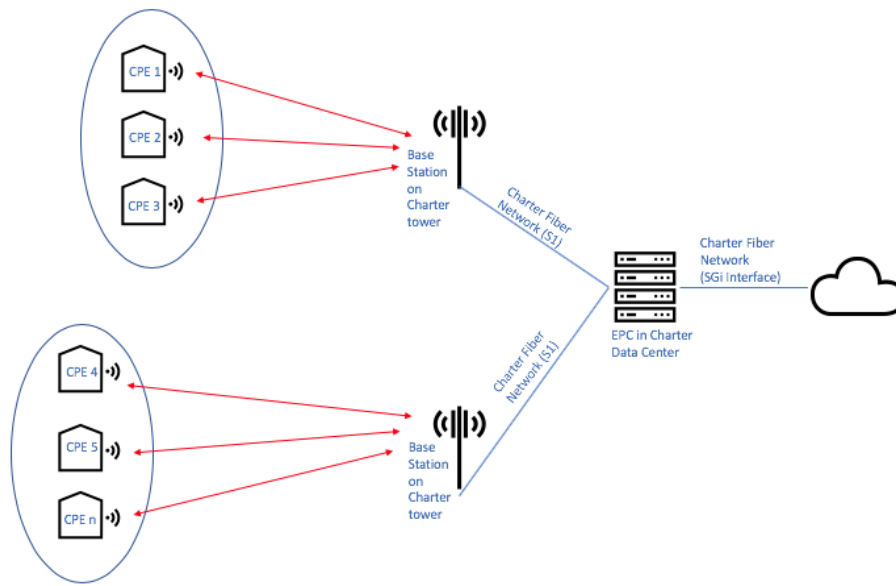
In this section we describe the FWA network architecture which is very much like the regular LTE network architecture as shown below. There is an Evolved Packet Core (EPC) and Radio Access Network (RAN) portions and instead of the User Equipment (UE) a CPE will be installed on the outside of the customer's home pointing towards the radio.



**Figure 2 - LTE Network Architecture**

Unlike mobile networks, some settings won't be necessary like thresholds for handovers and load balancing on X2 since the CPE will never change location.

For our testing we installed the EPC in one of Charter's data center and used Charter owned fiber as front haul to the radio as well as backhaul to the internet. In all our testing the radios were installed on Charter towers or on top of a Charter buildings. A high-level network diagram of our testing is shown below.



**Figure 3 - Charter Testing Network Architecture**

During field testing we used two custom-built test vans with hydraulic mast that goes up to 45 ft high. Each van was equipped with a test station containing state of the art computer running CPE debugging software. This software reads information from the CPE's chipset and records it, the information included Radio Frequency (RF) signal strength, Signal-to-noise ratio, Quadrature Amplitude Modulation (QAM), throughput at various layers, signaling messages, neighboring cells, frequencies, etc.

The vans were used to simulate various house heights at different locations. For the Denver employee field trial, the vans were instrumental in checking the RF signal at participants' homes prior to CPE installation. Below is a picture of one of the test vans.



**Figure 4 - Charter Custom Built Test Van**

To run throughput tests we used two methods; the first method is installing iperf, which is an industry standard traffic pushing software, in the data center and on mini-PC's connected to the CPE's. This setup allowed us to push downlink and uplink as Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) whenever needed and for as long as we wanted. This method accounted for the latency in the core and RAN.

The second method is using an industry standard hardware plugged into the CPE, when the hardware detects that the CPE is not being utilized for a while it runs speed tests towards a server in the internet. This second method was used only in the Denver employee field trial because the employees were using the FWA service and we didn't want to interrupt their use. This hardware ended up being useful to run speed tests when employees weren't using the FWA service. This method accounts for latency in the core, RAN, and internet.

In few instances we wanted to test the capacity of a loaded network so we forced the iperf method while employees were using the service in the peak hour.

In this paper we present the results of our testing in Denver where the terrain was very rough with many hills, in Tampa where the high trees foliage didn't allow the signal to travel far, in Cold water where we wanted to study the effect of extreme weather on the RF signal, and finally in Lexington where we ran a multi-user demo showcasing a typical household environment.

# Technology Description

In this section we present the technology used in our field testing. In the CBRS band we used both LTE 3<sup>rd</sup> Generation Partnership Project (3GPP) compliant equipment and non-3GPP proprietary equipment from multiple vendors. In the unlicensed band testing we used non-3GPP proprietary equipment transmitting in the 5GHz Unlicensed National Information Infrastructure (U-NII) band.

## 1.1. LTE Equipment Specification

For the LTE equipment vendors, the objective was to test the technology given the different stage of CBRS equipment development. We started testing with Band 42 and 43 then finally moved to band 48 equipment.

**Table 1 – LTE eNB Equipment Specifications**

Specification	Value
Product	LTE Macro eNB
Band Support	B48 (3.55 – 3.7 GHz)
Carrier Aggregation	Up to 3 CA
MIMO	2x2
Frame Configuration	TDD Frame Configuration 2 Special Sub Frame 7
IBW/OBW	150 MHz / 60 MHz
Output Power	2 x 10 W
Antenna	Built-in 11dBi OR utilize external antenna
Modulation QAM DL/UL	256 / 64
BF Capability	No
CBRS Classification	CBSD CAT B

For our field testing we used both the built-in 11dBi antenna and an external 17dBi antenna at different instances. Each time we ran tests we made sure to adjust the radio power output in order to stay in compliance with the FCC power limitation rules of 47dBm/10MHz for CAT B Citizen Broadband Radio Service Device (CBSD)[4][5].

The CPE used for testing was an outdoor CAT B CBSD-CPE, in 2019 this CPE became OnGo CBRS and FCC certified, it's worth noting that when we started our trials in 2017 there was no certification process in place. The specifications of the CPE used is presented in the below table.

**Table 2 - LTE CPE Equipment Specifications**

Specification	Value
Product	CPE
Setup	Outdoor Mounted
Band support	B42, B43, B48
Chipset	GCT
Carrier Aggregation DL/UL	4CA / 2 CA
MIMO	Up to 4x4
LTE Category	CAT 15
Output Power	23 dBm
Antenna Gain	10 dBi

Specification	Value
Modulation QAM DL/UL	256/64
CBRS Classification	CBSD-CPE CAT B

The total Equivalent Isotropically Radiated Power (EIRP) of this CPE is 33dBm which is less than the 47dBm/10MHz FCC rules. During testing we faced some coverage limitations because we were uplink challenged. A higher power CPE would have allowed successful RACH channel at further distances from the radio thus increasing the coverage of the tested eNB. This topic is further discussed in the next section, Test Results.

The proprietary equipment used for testing came in two different version. One version works in B48 (3.55-3.7GHz) and the other in the U-NII band (5.1-5.9GHz). The proprietary equipment spec's is presented in the below table.

**Table 3 - Proprietary Equipment Specifications**

Specification	Value
Product	Proprietary Wireless Equipment
Band Support	B48 and U-NII 5GHz
MIMO	4x4
Duplex	TDD 50:50
EIRP	CBRS: 43 dBi 5GHz: 33 dBi
Modulation QAM DL/UL	Up to 512/512 QAM
BF Capability	Digital Beam Forming 16 active RF chains
CBRS Classification	CBSD CAT B
5GHz Classification	Point-to-Multi-Point

This equipment transmitted as point to multi-point without the need of an EPC or Core. The transmitter and receiver acted as a layer 2 switch passing the data. The EIRP in the CBRS was 43dBm/20MHz which is short of the maximum allowed by the FCC. In the U-NII band the EIRP was 33 dBm as allowed by the FCC[6]. Both transmitter and receiver had the same specifications. Modulation QAM 1024 was also tested but was hard to achieve for locations far from the transmitter due to drop in Signal to Noise Ratio (SNR).

## Test Results

In this section we present the field test results from different markets we been to. We also present the reasoning of choosing to test CBRS in these markets and the purpose behind the executed testing.

### 1.1. Denver Trial – Hilly Terrain

The first test market was Denver, CO. The terrain in Denver is hilly with high percentage of Fresnel zone blockage which made it very challenging for coverage.

We executed two sets of tests, the first is to study the coverage in the hilly terrain environment and the second is to investigate the capacity in the same environment.



### 1.1.1. Coverage in Hilly Terrain

CBRS radio was mounted on the rooftop of one of Charter's building at 45 ft height and 7 test points were selected to reflect different terrain profile as per the below table. Only point 1 had Line Of Sight (LOS) to the radio and the rest were non-Line Of Sight (nLOS).

**Table 4 - Denver Test Points Distance from Base Station**

Test Point	Distance From Radio (Miles)
Point 1	0.22
Point 2	0.87
Point 3	1.27
Point 4	2.04
Point 5	2.47
Point 6	2.87
Point 7	3.21

The below map shows the seven test points and the radio location. We used one channel of 10MHz bandwidth to transmit. The EIRP used was 47dBm since we were using one channel only. All points were chosen to fall within the 65 degrees 3dB horizontal bandwidth of the antenna.




**Figure 5 - Denver Coverage Trial Fixed Locations**


At each test point we tested the CPE at two different heights, 12 ft and 20 ft. These two heights were chosen to mimic a one- and two-story house.

At 12 ft CPE high the CPE couldn't attach at points 5 and 7. At 20 ft CPE high the CPE couldn't attach at point 5. We had the qRexLevMin set to 128 dBm which is the minimum Reference Signal Received Power (RSRP) values measured by the UE in a cell to be able to get unrestricted coverage-based service in that cell. This means at points 5 and 7 (12 ft CPE) the CPE detected RSRP lower than 128 dBm.

The figures below show the Downlink (DL) throughput, and Uplink (UL) throughput at all test points for 12ft and 20ft CPE height.



**Figure 6 - Denver Coverage Trial DL Throughput**

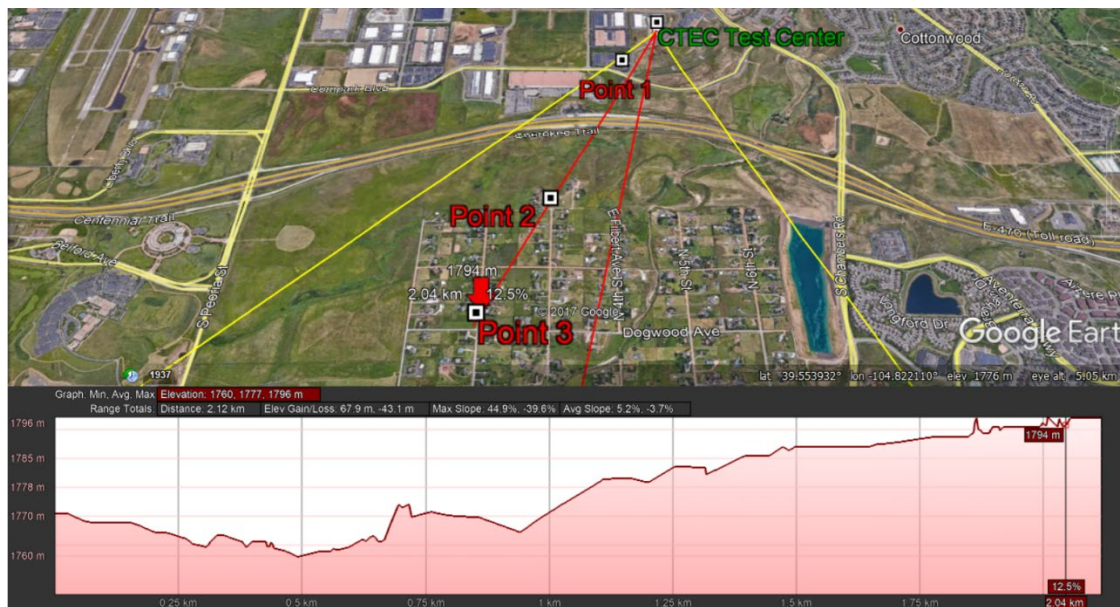


**Figure 7 - Denver Coverage Trial UL Throughput**

When we raised the CPE to 20 ft height, the CPE could attach at point 7 while still cannot attach at point 5. The RF conditions improved and the throughput increased due to raising the CPE.

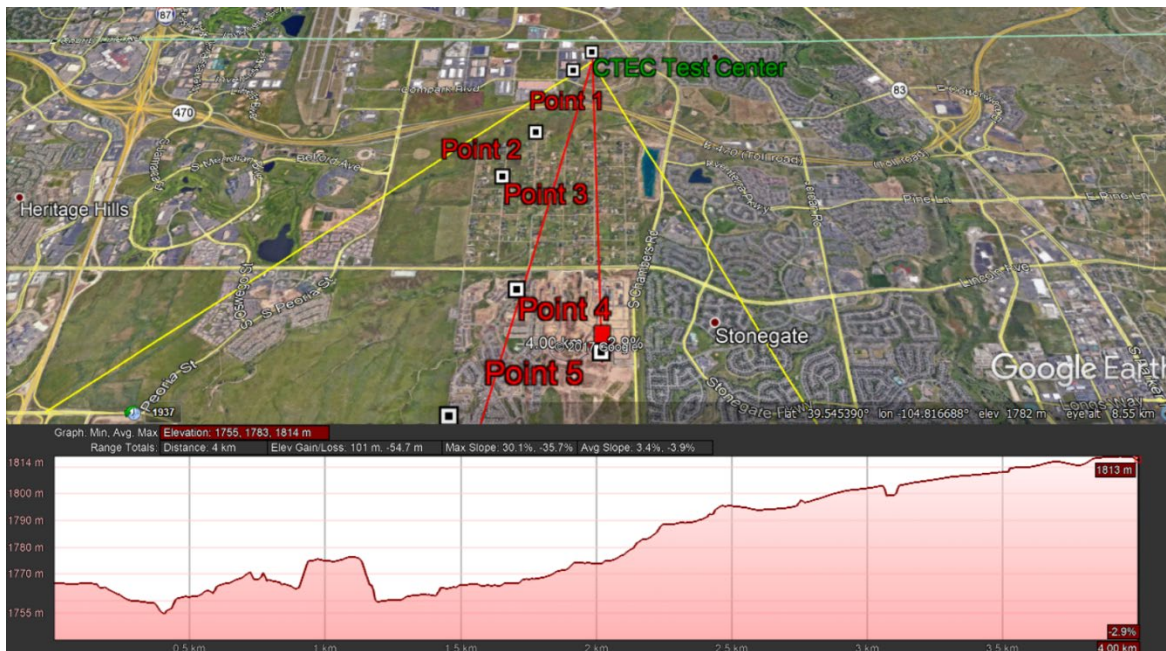
At 12 ft high CPE the QAM modulation varied between 64QAM and 16QAM with point 3 mostly 16QAM. At 20 ft high CPE we saw 64QAM most of the times at all points except point 7 where we saw mostly Quadrature Phase Shift Keying (QPSK) modulation which explains the big dip in throughput at point 7.

For both CPE heights we saw a dip in RF conditions at point 3 and no attach at point 5. Looking at the terrain profile for both points we can see the test points are at a much higher elevation than the radio with a lot of terrain blockage. Figure 9 shows the elevation profile for the fixed point 3.



**Figure 8 - Denver Test Point 3 Elevation Profile**

Point 3 was in an open area with few houses and high trees up to 20 ft high. Figure 10 shows the elevation profile for point 5.



**Figure 9 - Denver Test Point 5 Elevation Profile**

Point 5 was in a residential area with some mid-rise buildings.


The inability for the CPE to attach at point 5 made us think of the need to use different technologies to attach at hard to reach urban areas. At the time of testing there was no available LTE beam forming radios in band 48 so we turned to proprietary technology.

We installed a non-3GPP proprietary radio on the top of the Charter building instead of the macro eNB, and took the receiver to all 7 test points to test at 12 ft and 20 ft heights same as we did with the LTE CPE. The main differences between the proprietary equipment and the LTE is the proprietary supported digital beam forming with 16 active RF chains, transmitted at 43 dBm/20 MHz which is short of the FCC rules of maximum EIRP, and the receiver also transmitted at 43 dBi vs 33 dBi for the LTE CPE. The channel bandwidth was 20MHz and the frame configuration was set to 50:50.


With this setup we didn't expect UL challenges due to the high-power CPE also and wanted to understand the effect of beam forming in the CBRS band.

The below figures show the throughput, SNR, and Received Signal Strength Indicator (RSSI) at all test points for 12 ft and 20 ft high CPE.






**Figure 10 - Proprietary Radio DL Throughput at 25 ft CPE**




**Figure 11 - Proprietary Radio UL Throughput at 25 ft CPE**

Note the downlink and uplink are almost identical, that's because the equipment used TDD frame configuration 1:1 meaning half the resources are used for downlink and the other half used for uplink.



**Figure 12 - Proprietary Radio SNR Values at 25 ft CPE**



**Figure 13 - Proprietary Radio RSSI Values at 25 ft CPE**

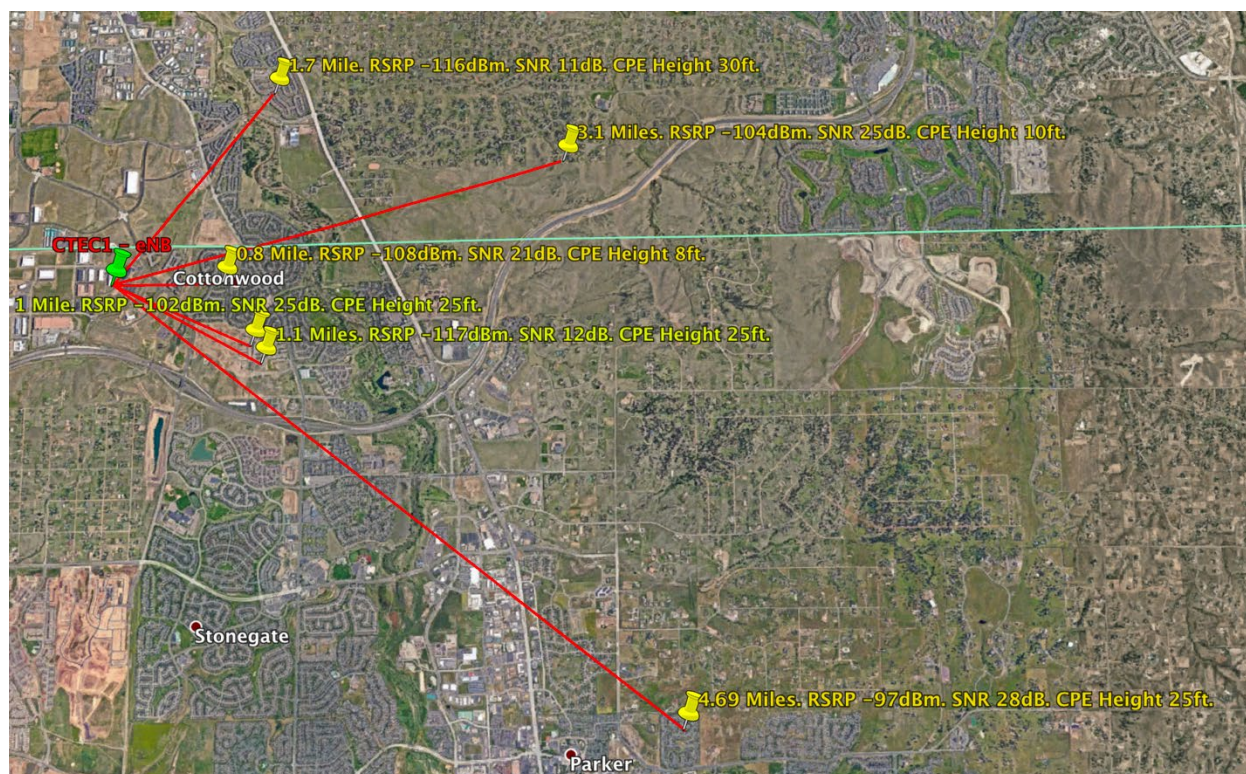
With the proprietary CBRS equipment we saw the same dip in RF conditions at point 3 similar to the LTE equipment tested earlier. The two interesting observations we saw was the CPE connected at point 5 and RF conditions at 12ft and 25 ft weren't very different except at point 5.

We concluded that due to beam forming effect we could connect at point 5 between the mid-rise buildings. This conclusion encouraged us to work with the 3GPP LTE vendors to get beam forming LTE CBRS equipment to test various scenarios in residential areas.

### 1.1.2. Capacity in Hilly Terrain

After we determined the CBRS coverage in Denver, it was time to study the capacity. We are conducting an employee field trial in Denver where we provided broadband speeds of minimum 25Mbps downlink and 3Mbps uplink. Many employees signed up to be part of the trial and agreed to have CPE's installed on their houses. In the Denver employee trial, we used the LTE CBRS equipment and had several sectors serving the employees, here we focus on the results of two sectors only.

The first sector was installed on Charter's building rooftop used earlier in the coverage testing. At the time of writing this paper this sector was used to serve 7 employees with more being added as shown in the below figure.



**Figure 14 - Denver Field Trial First Sector**


This sector served employees ranging from 0.8 mile to 4.7 miles away from the radio. The below table shows each house's distance from the eNB along with the RF conditions and throughputs.

**Table 5 - Denver Trial First Sector Participants RF conditions and Throughput**

Distance from eNB (miles)	RSRP(dBm)	SNR(dB)	CPE Height (ft)	Downlink (Mbps)	Uplink (Mbps)
0.8	-108	21	8	174	4
1	-102	25	25	58.4	2.8
1.11	-117	12	25	38.2	2.3
1.7	-116	11	30	91	3.8
3.1	-104	25	20	204	10.5
4.69	-97	28	25	243	10.3

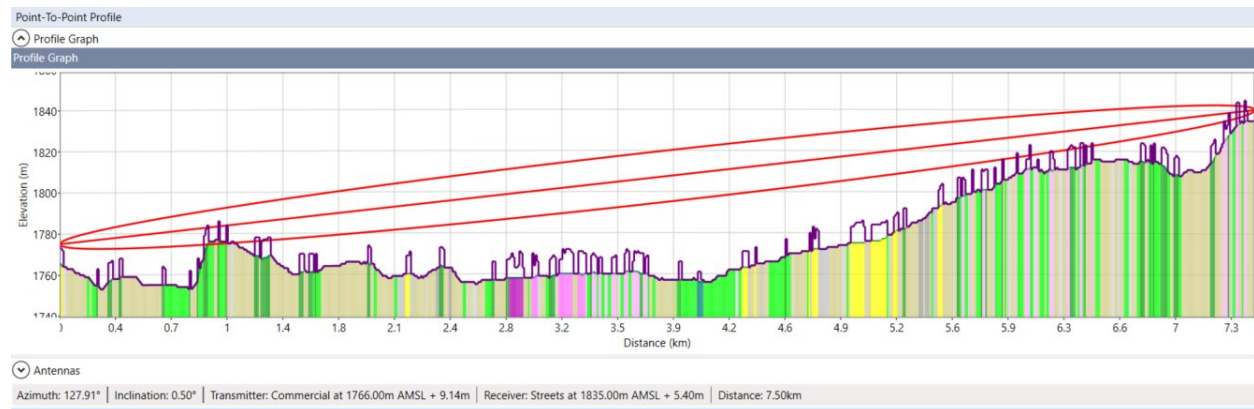
The employee field trial data showed us some very interesting insights for example there is a participant at 1.1 mile from the radio with CPE installed at 25ft getting worst RF conditions than a participant at 4.7 miles away with CPE also at 25ft high.

The below is a screenshot from a propagation modeling tool showing the CPE in this participant house 1.1 mile away from the radio. The model shows blockage of the Fresnel zone which in turns explains the CPE is getting RSRP -117dBm and SNR 12dB although fairly close to the radio.



**Figure 15 - Participant With Major Fresnel Zone Blockage**


The same tool showed the CPE in a participant house 4.7 miles away has little to no Fresnel zone blockage. That explains the good RF conditions RSRP -97dBm and SNR 28dB at 4.7 miles away from the radio.



**Figure 16 - Participant With Minor Fresnel Zone Blockage**




The throughputs the employees were getting are shown in the below graph. The graph shows the throughputs they were getting when the sector wasn't loaded and when the sector was fully loaded.



**Figure 17 - Participant First Sector Throughputs**

The blue line represents the throughput of each CPE performing maximum download for its location one CPE at a time while the orange represents the throughput when all CPE's are performing maximum download at the same time. There were different scheduler settings for the eNB to serve multiple CPE's simultaneously but we decided to keep it at proportional fair to guarantee fair resources distribution to all CPE's.

The second sector we focus on in this paper was installed on another Charter building at 125 ft and is currently serving 11 employees, more are being added. This sector served employees ranging from 0.6 mile to 5 miles away from the radio as shown in the below figure.



**Figure 18 - Denver Field Trial Second Sector**

Since this sector was at a higher elevation than the first sector we expected better performance in terms of CPE RF conditions.


RF conditions of most participants were good due to the radio installed on a high building. For example, there is a participant at 5 miles away getting RSRP -114dBm, SNR 24dB and another participant at 4.3 miles getting RSRP -99 dBm, SNR 29dB. The RF conditions along with throughputs of the CPE's are shown in the below table.

**Table 6 - Denver Trial Second Sector Participants RF conditions and Throughput**

Distance from eNB (miles)	RSRP(dBm)	SNR(dB)	Downlink (Mbps)	Uplink (Mbps)
0.6	-112	16	153	4.31
1.4	-110	19	210	5.13
1.4	-118	12	155	6.16
2.4	-116	13	121	1.29
2.6	-106	21	192	5.56
2.7	-95	29	248	10.2
3.4	-100	26	270	12

Distance from eNB (miles)	RSRP(dBm)	SNR(dB)	Downlink (Mbps)	Uplink (Mbps)
3.5	-113	15	146	3.49
3.8	-99	26	273	13.9
4.3	-96	28	257	11.9
5	-112	15	165	6.28

The throughputs the employees were getting are shown in the below graph. The graph shows the throughputs they were getting when the sector wasn't loaded and when the sector was fully loaded.




**Figure 19 - Participants Second Sector Throughput**

When loading all CPE's at once the LTE resource blocks were shared among CPE's thus some CPE's downlink fell below 25Mbps which was expected. During network planning there is a factor called over subscription which basically means that not all subscribers will be doing maximum throughput simultaneously even at peak hours. This over subscription factor differs from a network to another and is used to make sure the network is fully utilized while at the same time ensuring subscribers experience a high quality of service. Network planners have to carefully pick the over subscription factor in order not to affect users experience.


For the sake of sector capacity, we calculated that theoretically a 40MHz LTE CBRS channel can achieve around 300Mbps sector throughput at 256QAM. In the lab we got a maximum 285Mbps at perfect RF conditions. We turned to the field and ran a couple of test cases where we had 10 CPE's at 256QAM and measured the sector throughput, then moved the 10 CPE's away from the eNB till they dropped to 64/16QAM and again measured the sector throughput.

With all 10 CPE's at 256QAM and doing maximum downlink simultaneously the sector throughput was 253.9 Mbps as shown in figure 20.



**Figure 20 - Sector Throughput with 10 CPE's at 256QAM**

With all 10 CPE's at 64/16QAM and doing maximum downlink simultaneously the sector throughput dropped to 111.6 Mbps as shown in figure 21.



**Figure 21 - Sector Throughput with 10 CPE's at 64/16QAM**


The conclusion here is that the location and subsequently RF conditions of the CPE's will greatly affect the sector throughput because CPE's at bad RF conditions consumes more resource blocks. This is a point of consideration when designing FWA networks since CPE's will always be at the same location and their RF conditions can be pre-determined before installation.

To further understand the sector throughput in commercial deployments we turned to our Denver field trial to the sector serving 11 subscribers. On that sector we did the same previous couple of tests where we added 10 CPE's at 256QAM to the sector and measured the sector capacity then added 10 CPE's at 64QAM and measured the sector capacity. We felt these two tests are a better representation of a real-world scenario because we had several subscribers at various RF conditions.

Figure 22 shows the Denver trial sector when adding 10 more subscribers to it all at 256QAM, we ended up having 14 CPE's at 256QAM, 3 CPE's at 64QAM, and 4 CPE's at 16QAM. The calculated sector throughput was 227 Mbps. This can be considered an ideal scenario sector because most subscribers are at 256QAM.

**Table 7 - Adding 10 CPE's at 256QAM to Existing Sector**

QAM Modulation	Number of CPE's
16	4
64	3
256	14




**Figure 22 - Adding 10 CPE's at 256QAM to Existing Sector - Sector Throughput = 227 Mbps**

Figure 23 shows the Denver trial sector when adding 10 more subscribers to it all at 64QAM, we ended up having 4 CPE's at 256QAM, 13 CPE's at 64QAM, and 4 CPE's at 16QAM. The calculated sector throughput dropped significantly to 152 Mbps. This can be considered a realistic scenario sector because most subscribers are at 64QAM.

**Table 8 - Adding 10 CPE's at 64QAM to Existing Sector**

QAM Modulation	Number of CPE's
16	4
64	13
256	4



**Figure 23 - Adding 10 CPE's at 64QAM to Existing Sector - Sector Throughput = 152 Mbps**

These tests show the importance of FWA network planning for capacity purposes because the distribution of subscribers will affect the sector throughput and quality of service. Also, the theoretical maximum sector throughput will likely not be seen in the field. Ideally FWA network planners will want all subscribers at 256QAM which is not realistic because the cell radius will likely be small. A cell radius vs capacity tradeoff has to happen to ensure a good size cell radius with reasonable throughputs, most likely the majority of subscribers will have to be on 64QAM with very few on 256 and 16QAM.

We got a lot of learning from the Denver trial for example how to find the perfect place on the house to install the CPE, plan LTE CBRS network, and monitor various network components' performance. Also, to consider oversubscription when planning a FWA commercial network, and keeping in mind that sector throughput will vary depending on the subscriber's locations and RF conditions. We also learned the higher we install the radio the better RF coverage we can achieve.

## 1.2. Tampa Trial – High Foliage

After testing CBRS in Denver which is considered an urban hilly environment, we decided to study the effect of high foliage on CBRS coverage. We wanted to understand how can CBRS be used in rural areas to close the digital gap in the country. Our choice was New Port Richey which is north of Tampa, FL.

The test market had high dense trees going up more than 25 ft. One of the reasons we chose New Port Richey is because Charter owned a 140ft tower which made it a perfect environment to test CBRS FWA.

### 1.2.1. Coverage in High Foliage – LTE CBRS

To understand the CBRS coverage in foliage, we decided to use both the LTE and the proprietary equipment.


We started with the LTE CBRS radios and installed it at 130ft on the Charter owned tower. We used one channel 10MHz bandwidth and set the radio EIRP to 47dBm.

Test points were chosen as per the below table. All test points were non-line of sight. One at a time we took the test van to each test point and setup the CPE at 12ft and 25ft to test how far the radio can cover.

**Table 9 - Tampa Test Points Distance from Base Station**


Test Point	Distance From Radio (Miles)
Point 1	0.32
Point 2	0.58
Point 3	1.06
Point 4	1.24
Point 5	1.68
Point 6	2.67
Point 7	3.21

The RF and throughput results for both CPE heights are shown below graph.




**Figure 24 - RSRP at 12ft and 25ft CPE Height**






**Figure 25- SNR at 12ft and 25ft CPE Height**




**Figure 26 - DL Throughput at 12ft and 25ft CPE Height**




**Figure 27 - UL Throughput at 12ft and 25ft CPE Height**

An interesting observation is at points 3 and 4 as the height of the CPE didn't make much of a difference in terms of throughput. That's due to RF conditions being very similar at both CPE heights due to the high dense trees blocking the signal, going above tree line would have made the RF conditions better.

The below figures show the QAM modulation values at all test points for both CPE heights.



**Figure 28 - Modulation at 12ft CPE Height**



**Figure 29 - Modulation at 25ft CPE Height**

From the above modulation scheme graphs we see points 3 and 4 at both CPE heights are very similar which again explains the similar throughput at both heights.

Looking at point 5, It's clear that increasing the CPE height made the RF conditions better, thus changed from 16QAM at 12ft CPE height to 64QAM at 25ft CPE height.

The CPE couldn't attach at point 7 (3.2 miles) and beyond. We concluded the high foliage limited our sector coverage and to get good RF conditions the CPE must be installed above tree line.

We went beyond 3.5 miles and raised the CPE 45ft high and could connect, however, 45ft high CPE isn't practical in real live deployment scenarios. We wanted our testing to mimic a CPE installed at one- or two-story buildings.

### **1.2.2. Coverage in High Foliage – Proprietary CBRS and 5GHz**

We then moved to test the proprietary equipment and did two sets of testing. First, we tested the CBRS proprietary equipment then we tested the unlicensed 5GHz version. We mounted both CBRS and 5GHz radios at 130ft height on the tower. The CBRS maximum EIRP was 43dBm which is a bit short of the FCC maximum EIRP CBRS rules. The 5GHz radio used maximum EIRP 33 dBm which is the maximum power allowed for U-NII by the FCC. Both radios used 20MHz channels and TDD frame configuration 50:50.


Since these proprietary radios used beam forming technology and performed well in the Denver trial, we decided to challenge them. We chose a different yet more challenging test point set as below.

**Table 10 - Tampa Test Points for Proprietary Equipment**

Test Point	Distance From Radio (Miles)
Point 1	0.6
Point 2	1.7


Test Point	Distance From Radio (Miles)
Point 3	2.7
Point 4	3.2
Point 5	4
Point 6	5

At 12ft high CPE, the CBRS radios connected at points 1, 2, and 6. The 5GHz radio connected as points 1 and 6 only. It was expected to get better performance with the CBRS radio vs. the 5GHz radio because of the high power of the CBRS radio and lower frequency of CBRS compared to the 5GHz.



**Figure 30 - Proprietary CBRS and 5GHz Throughput Results for 12ft CPE**

Raising the CPE to 25ft high gave better performance for both the CBRS and 5GHz radios. The CBRS radio connected at all points except point 4, while the 5GHz radio connected at points 1, 2, and 6 only.



**Figure 31 - Proprietary CBRS and 5GHz Throughput Results for 25ft CPE**

The interesting points here were points 5 and 6 which are 4 and 5 miles away from the radio. Raising the CPE to the clutter height above the tree line resulted in the CPE connected with decent throughputs.

We concluded that in FWA scenarios when using a high tower to install the radio we can still connect in high foliage environments if the CPE is above the clutter height.

Also, the 5GHz CPE connecting at 5 miles away from the radio was interesting and encouraged us to think of building a prototype CBRS+5GHz combined radio.

### 1.3. Coldwater Trial – Snow and Rain

After testing CBRS in the high foliage environment of Florida, we went to Coldwater, MI to study the effect of snow and rain on CBRS RF conditions. We also had just gotten massive MIMO LTE CBRS equipment and wanted to understand the effect of LTE beam forming gain on coverage. Finally, we wanted to test the idea of combining CBRS with 5GHz radio.

Charter owns a tower in Coldwater, MI and we used it to install the radios at 130ft height. Similar to the other test markets we chose test points varying from 0.3 to 5 miles.

#### 1.3.1. Snow and Rain Effect on CBRS

We went to test in January some days were raining and snowing while other days were just gloomy without snow or rain, it was the perfect environment to understand if snow or rain had any effect on CBRS signal. We picked two test points at 0.8 and 1.4 miles from the radio and tested the LTE CBRS radio on two different days to capture the effect of rain. At both days we tested the CPE at 12ft and 25ft high.

At 0.8 mile from the tower the maximum fluctuation in RSRP was 3dBm, SNR 3dB, and DL throughput 2Mbps as seen in the below graphs.






Figure 32 - Effect of Rain on RSRP at Different CPE Heights – Cell Near




**Figure 33 - Effect of Rain on SNR at Different CPE Heights – Cell Near**




**Figure 34 - Effect of Rain on Throughput at Different CPE Heights – Cell Near**


At 1.4 mile from the tower the maximum fluctuation in RSRP was 4dBm, SNR 5dB, and DL throughput 6Mbps as seen in the below graphs.



**Figure 35 - Effect of Rain on RSRP at Different CPE Heights – Mid-cell**



**Figure 36 - Effect of Rain on SNR at Different CPE Heights – Mid-cell**



**Figure 37 - Effect of Rain on Throughput at Different CPE Heights – Mid-cell**


From the above graphs we concluded that unlike mmWave signal the rain had little to no effect on the LTE CBRS signal. The RF fluctuation was in the normal fluctuation range.

### 1.3.2. Beamforming Gain – Coverage Test

To understand the LTE CBRS beam forming gain effect on coverage we used a massive MIMO 64Tx64R at two test points and compared the results to the regular macro at the same test points.


We chose test points to represent cell middle and cell edge at 2.5 and 5 miles away from the tower. The RSRP's were -109dBm and -130dBm respectively. The macro was set to use open loop Transmission Mode (TM) TM 3 while the massive MIMO was set to use Single User MIMO (SU-MIMO) TM7.

As seen in the below graph, the macro performed better in terms of throughput at cell middle. However, at cell edge the beam forming gain gave 35% more throughput than the macro.



**Figure 38 - Effect of Beamforming on Throughput at Mid-cell**





**Figure 39 - Effect of Beamforming on Throughput at Cell Edge**

We concluded that beam forming gain using SU-MIMO TM7 can help users at cell edge to get better throughput but not at cell middle conditions. This encouraged us to run more tests in the future to further understand the effect of various transmission modes on coverage and capacity.


### **1.3.3. CBRS+5GHz Radio Prototype**

Using the proprietary equipment, we built a prototype of CBRS + 5GHz by aggregating the radios using a layer 2 switch. We did the same aggregation method to the receiver using a layer 2 switch. Since this was just a prototype this wasn't carrier aggregation between both bands rather just aggregating both radios at each end together.

Our goal was to make the link more reliable in case one of the technologies link goes down the other will stay up and also to increase the throughput delivered to the receiver.


We used two channels each is 20MHz, one for the CBRS and the other for the 5GHz. We noticed that the 5GHz didn't connect at one test point for the 12ft receiver but the CBRS did.

The below graphs show the aggregate maximum throughput we got from our prototype unit and also the breakdown of each radio on its own.



**Figure 40 - Prototype Radio Throughput at 12ft Receiver Height**

Notice at 3 miles away from the tower at 12ft CPE the 5GHz radio didn't connect but the CBRS did.



**Figure 41 - Prototype Radio Throughput at 25ft Receiver Height**

It's worth noting that at 3 miles the test point was heavily obstructed due to surrounding buildings and at 5 miles we were testing in an open area. That explains the dip in throughput at 3 miles from the tower then better performance at 5 miles from the tower.

The prototype combined radio delivered high throughput at all test points due to the aggregation of both bands also it showed reliability where one radio couldn't attach but the other did.

The prototype encouraged us to think of more innovative ways to aggregate different bands in order to provide a more reliable and faster link. Also, this directed us to explore 3GPP technologies like Licensed Assisted Access (LAA) where a licensed and unlicensed band are combined together.

## 1.4. Lexington Trial – User Experience

From all the previous trials we got a pretty good understanding of CBRS equipment coverage and capacity. What we didn't explore was how much throughput was enough for a typical household usage.

We were out in Lexington, KY for LTE CBRS vendor testing and decided to see if 50Mbps was enough for a typical household usage.

We had a test site at 1.6 miles away from the radio with no line of sight to the CPE. We throttled the speed to 50Mbps and ran all of the following at once: 4K TV running on demand streaming service, tablet running 4K online video service, gaming console playing online game, IP camera, and several laptops browsing the internet and watching online videos. The figure below is from the Lexington demo showing various devices running at the same time simulating a typical household.



**Figure 42 - Various Internet Devices Running Simultaneously**

We found that none of these services were affected or buffered while all running simultaneously on 50Mbps downlink speed. We monitored the CPE and found the throughput will spike to 50Mbps then drops and few seconds later peaks again, that's because all the mentioned video services have buffering capabilities and are not continuously requested maximum speeds all the time.

We concluded that 50Mbps is more than enough for the typical household needs.

## Conclusion

Over the past couple of years, Charter team conducted several CBRS trials to further understand the coverage and capacity of this new band for FWA. It is concluded that a typical cell radius can be 3.5 to 5 miles depending on the terrain and morphology. It is also concluded that foliage negatively affects the CBRS signal and in such areas the CPE's should be installed at or above clutter height. Snow and rain have almost no effect on CBRS signal, and that 50Mbps can run what a typical household family need. SU-MIMO would help give better throughput results for users at cell edge. However, more investigations are needed for massive MIMO technology capacity and how it can be used to further enhance FWA service. It's important to aggregate unlicensed band with CBRS to have a more reliable and faster link.

Finally, we conclude that CBRS can be used for FWA to bridge the digital divide in the nation and help connect millions of users in rural America. CBRS can't replace fiber but will definitely complement it and help reach the last mile users. Charter is in a very good position to use CBRS to extend its broadband offering and reach customers that couldn't be reached in the past due to the high cost of laying fiber. Charter can leverage not only its huge fiber network but also the towers it owns nationwide to reach new customers and provide them with FWA broadband internet speeds. Moreover, Charter could also provide IPTV services over the FWA to customers.

FWA on CBRS represent a very good opportunity for Charter to capture new customers and a new way to utilize Charter's current assets.

Charter will keep investing latest technologies in CBRS and other bands to help connect millions of Americans in rural areas and to provide better services to current customers.

## Abbreviations

3GPP	3rd Generation Partnership Project
CBRS	Citizens Broadband Radio Service
CBSD	Citizen Broadband Radio Service Device
CPE	Customer Premises Equipment
DL	Downlink
EIRP	Equivalent Isotropically Radiated Power
EPC	Evolved Packet Core
FCC	Federal Communications Commission
FWA	Fixed Wireless Access
GAA	General Authorized Access
LAA	Licensed Assisted Access
LOS	Line Of Sight
LTE	Long Term Evolution
MIMO	Multiple Input Multiple Output
nLOS	non-Line Of Sight
PAL	Priority Access License
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase Shift Keying
RAN	Radio Access Network
RF	Radio Frequency
RSRP	Reference Signal Received Power
RSSI	Received Signal Strength Indicator
SAS	Spectrum Access System
SIM	Subscriber Identification Module
SNR	Signal to Noise Ratio
SU-MIMO	Single User MIMO
TCP	Transmission Control Protocol
TM	Transmission Mode
TVWS	TV White Space
UDP	User Datagram Protocol
UE	User Equipment
UL	Uplink

U-NII	Unlicensed National Information Infrastructure
WISP	Wireless Internet Service Provider

## Bibliography & References

1. *Total Cost of Ownership for Fixed OnGo in the 3.5GHz CBRS Band*; CBRS Alliance. Link: [https://www.cbcrsalliance.org/wp-content/uploads/2018/06/The-TCO-of-fixed-OnGo-in-the-3.5-GHz-CBRS-band\\_Whitepaper.pdf](https://www.cbcrsalliance.org/wp-content/uploads/2018/06/The-TCO-of-fixed-OnGo-in-the-3.5-GHz-CBRS-band_Whitepaper.pdf)
2. *The Urban and Rural Classification*; United States Census Bureau. Link: <https://www2.census.gov/geo/pdfs/reference/GARM/Ch12GARM.pdf>
3. *2018 Broadband Deployment Report*; Federal Communications Commission. Link: <https://www.fcc.gov/reports-research/reports/broadband-progress-reports/2018-broadband-deployment-report>
4. Federal Communications Commission, Rulemaking 12-354 and 17-258. Link: <https://www.fcc.gov/wireless/bureau-divisions/mobility-division/35-ghz-band/35-ghz-band-citizens-broadband-radio-service>
5. Federal Communications Commission, Certification and Test Procedures for Citizens Broadband Radio Service Devices Authorized Under Part 96. Link: [https://apps.fcc.gov/kdb/GetAttachment.html?id=RV8R2eM861G%2FcoXTqUigyA%3D%3D&desc=940660%20D01%20Part%2096%20CBRS%20v02&tracking\\_number=229297](https://apps.fcc.gov/kdb/GetAttachment.html?id=RV8R2eM861G%2FcoXTqUigyA%3D%3D&desc=940660%20D01%20Part%2096%20CBRS%20v02&tracking_number=229297)
6. *Revision of Part 15 of the Commission's Rules to Permit Unlicensed National Information Infrastructure (U-NII) Devices in the 5GHz Band*; Federal Communications Commission, Link: <https://www.fcc.gov/document/5-ghz-unlicensed-spectrum-unii>

# **Building a Technology Platform for Smart Agriculture Deployments Using C-Band and Unlicensed Technologies**

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

**Joshua Sanders**

Network Engineer III  
Charter Communications  
6360 S. Fiddlers Green Circle, Greenwood Village, CO 80111  
(720) 536-9392  
Joshua.N.Sanders@charter.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
1. Background and Trial Purpose .....	3
2. Trial Network Architecture .....	3
3. IoT Sensors .....	4
4. C-Band Rural Broadband Connection.....	7
5. AR/VR User Interface.....	9
Conclusion .....	11
Abbreviations.....	11
Bibliography & References .....	12

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Smart Farm trial network architecture.....	4
Figure 2: left - LoRa gateway installed on the farm house, right: installing soil moisture sensor .....	5
Figure 3 - Main farm complex with grain bins (right), machine shed (center), and farm house (left). Hog house #1 is in the background to the left of the grain bins. ....	6
Figure 4 - Image from hog house #1 video camera. ....	7
Figure 5 - Wireless field test trailer with extendable radio mast for 5G NR and EPC equipment.....	8
Figure 6 - C-Band rural broadband connection.....	8
Figure 7 - Temperature and Moisture sensors in the grain bins reduce spoilage and increase yield.....	10
Figure 8 - Security cameras show live footage of the hogs. All pertinent information is available virtually. ....	10
Figure 9 - Need to visit a site in real-time? Remotely fly a drone to a pre-determined waypoint.....	11

# Introduction

This paper demonstrates how synergies between C-Band and unlicensed wireless technologies for massive connectivity in conjunction with IoT devices can enable a Smart Farm business vertical. The purpose of this project was to investigate areas where multiple access technologies along with IoT devices can demonstrate real savings and provide the basis for a model for managed IoT services in actual deployments along with investigating the range and performance of C-Band in rural broadband applications. Using available LoRa IoT sensors for monitoring various elements around the farm we designed and deployed a complete LoRa network to collect and house the device reports. Data visualization and presentation aspects were done through partnerships. The end user can view the data either through a web-based portal or with an app that runs on either a Virtual Reality HMD or Android device (i.e. smart phone or tablet). We have also seen a sizeable potential for a new business vertical for a cable MSO who would deploy such a system across a wide geographical footprint to serve US agricultural customers. These kinds of platforms could also be expanded to address other verticals like Industrial Automation, Health Care, and Smart Cities.

## Content

### 1. Background and Trial Purpose

Rural agricultural communities have historically been the last places to receive access to high-speed networks. Through low latency and massive connectivity, 5G is capable of improving farmers' lifestyles by making their tasks more efficient and cost effective.

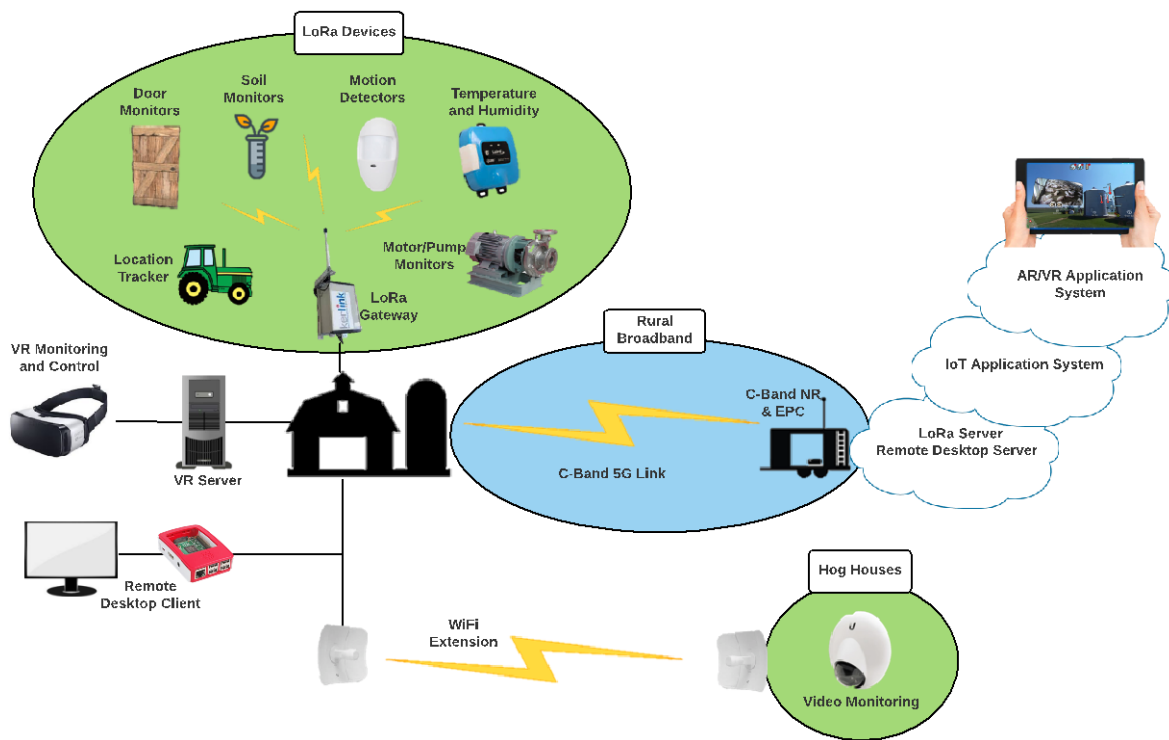
The Charter Wireless R&D team is conducting a trial on a working farm in Eastern Iowa which began in the fall of 2018. The farm is family owned and operated and is just outside the town of Keystone, which is about 30 miles west of Cedar Rapids. It consists of approximately 250 acres on which both corn and soybeans are grown. In addition to these crops there are 4 hog houses nearby with about 2500 hogs in each. Our goal for this project is to develop a widely deployable and scalable platform that provides demonstrable benefits with a positive business case for both grain farms and livestock farms but can also be extendable to other business segments and verticals. This will allow an MSO like Charter to develop new service offerings in areas that are beyond the traditional scope of a cable operator.

By making use of a cloud-based platform for collecting, analyzing, and displaying IoT sensor data in an immersive environment, farmers are able to remotely access and oversee their farms from anywhere. The developed platform provides building blocks for other use cases including Home Security, Industrial Automation, Smart Healthcare, and Smart Cities.

### 2. Trial Network Architecture

The overall Smart Farm trial solution platform consists of 4 major components: 1) the IoT sensors, 2) the cloud service that collects and manages the IoT data, 3) the fixed wireless link that provides rural broadband connectivity to the farm, and 4) the graphical user display application that includes overlaying the IoT data on a 360 degree live video.





**Figure 1 - Smart Farm trial network architecture**

For the trial system, the deployed LoRa IoT devices utilize the 900 MHz unlicensed spectrum band. The point-to-point network WiFi extensions utilize both the 5 GHz (802.11ac) and 60 GHz (802.11ad) unlicensed bands. The IoT RF data transmissions are received by a central LoRa gateway installed at the main farm complex which then sends the data packets through a secure IP tunnel to the cloud-based LoRa server and IoT application system. A separate cloud-based AR/VR application system powers the user interface which overlays the collected IoT data onto live or pre-recorded video.

### 3. IoT Sensors

The Smart Farm trial deployment consists of more than 70 IoT sensors and network devices. These devices include:

- Door monitors
- Soil monitors
- Motion detectors
- Motor activity monitoring
  - Grain bin blower
  - Hog house window shade
  - Water pump monitor
- Temperature and humidity monitors
- Location tracker
- Cameras: security and thermal

The trial system includes a total of 14 moisture sensors in the surrounding fields. Moisture probes at multiple soil depths are deployed to provide readings at 1ft, 2ft, and 3ft levels. All of the soil sensors operate on the LoRa system and provide soil moisture measurements every 1 hour. The purpose of using sensors at multiple depths is to provide an understanding of the moisture levels for the entire root depths of the crops.



**Figure 2: left - LoRa gateway installed on the farm house, right: installing soil moisture sensor**

There are 5 grain storage bins at the main farm complex. These bins are used to store the harvested grain and provide an environment for drying the grain to optimum moisture levels before sale on the market. Multiple LoRa sensors are deployed in each of the grain bins to measure temperature and humidity levels. One sensor is attached to the underside of each bin roof and a second identical sensor is suspended from a rope and deployed inside the grain about 2 meters from the top of the bin. The reason for this is to 1) compare the readings from the two sensors to see if there is an appreciable difference between them due to placement in the bins and 2) to investigate the amount of RF penetration loss of the grain in the bin.



**Figure 3 - Main farm complex with grain bins (right), machine shed (center), and farm house (left). Hog house #1 is in the background to the left of the grain bins.**

There are 4 hog houses which are not at the main farm complex. Hog house #1 can be seen in Figure 3 and is roughly 600 meters to the east. Hog house #2 is about 2 km northeast. Hog houses #3 and #4 are adjacent to one another and are 4.5 km to the west. In each of the 4 hog houses there are deployed 2 different types of LoRa environment monitors. The first type is the same type as the ones that are deployed in the grain bins to monitor temperature and humidity. The second type is a more capable device and monitors temperature, humidity, CO2 levels, as well as barometric pressure.

Door monitors are deployed on each entry door of the hog houses as well as the machine shed in the main farm complex. Using an activation magnet these sensors are triggered by door open and close events and simply send an event message to the LoRa network which is recorded by the cloud-based IoT application system. The IoT cloud system can be configured to generate alarms upon door events during certain hours of the day. The generated alarms can send notifications to alert designated individuals of the events.

There are 5 instances of equipment monitors – hog house fans, hog house window shades, hog feed trough, grain bin blower, and well pump. Relays are used in each instance to trigger the accompanying LoRa devices when power is activated and de-activated. In one case there is a well pump being monitored that is over 2 km away from the main farm complex.

In another instance a custom device was constructed using a programmable LoRa module with a GPS carrier board to create a device with vehicular tracking capability. This custom device was mounted on the farm's tractor and is powered by 12V from the vehicle. Once synced to GPS it transmits a position reading in a custom LoRa message every 12 seconds. Using the coordinates from the LoRa message a location dot is then plotted on a digital map to show the present and recent locations of the tractor.

WiFi network extensions have been set up between the hog houses and the farmhouse complex. Point-to-point 802.11ac devices at 5 GHz are used for the longer distance links (500m – 4.5 km), and point-to-point 802.11ad devices at 60 GHz are used for high bandwidth shorter distance links (< 100m).

Each of the 4 hog houses has 2 HD video cameras mounted in them. This allows for monitoring of the livestock from anywhere via the user interface. The required bandwidth for the HD RTSP video streams



(~4 Mbps) is too much for the LoRa system to accommodate so the video streams are carried by the WiFi links.



**Figure 4 - Image from hog house #1 video camera.**

In one of the hog houses there are an additional 2 FLIR thermal cameras. With additional computing resources it could be possible to recognize deceased animals and send notifications upon detection which would provide an increased benefit to the livestock manager. This detection capability was beyond the scope of this trial, but the cameras were installed to illustrate the potentials and determine possible system requirements for such a feature.


#### **4. C-Band Rural Broadband Connection**

The transport connection for the IoT data is a fixed wireless broadband link using C-Band spectrum. We obtained a special temporary authority (STA) permit from the FCC to use a 100 MHz channel in the C-Band spectrum, in which we deployed pre-commercial 5G NR and UE equipment for this trial.



**Figure 5 - Wireless field test trailer with extendable radio mast for 5G NR and EPC equipment**

In the 100 MHz C-Band channel we were able to demonstrate 429 Mbps of broadband downlink throughput at a distance of 4.5 km with a prototype radio software load. Higher throughput speeds will be achieved with production software loads as more features and capabilities are incorporated into the radio products. The 5G equipment vendor estimates that they should be able to provide at least 800 Mbps at a distance of 4.5 km using full production software. One thing to note is that this estimate is based on an NR antenna height of 10 meters which is the mast height of the trailer that was used in our testing.



**Figure 6 - C-Band rural broadband connection**

For shorter link distances, e.g. 1 - 2 km, the estimated achievable throughput is well over 1 Gbps. Again, this estimate is based on an NR antenna height of 10 meters. So for typical rural deployments where the NR antenna height could be 50 – 100 meters, we could expect throughput speeds in a 100 MHz C-Band channel to be in excess of 1 Gbps for a much greater distance than what we have seen in this trial.

A farm of the size that is used in our trial (250 acres) covers a total area of 1 sq km. A rural broadband cell site with a service radius of 5 km would cover more than 78 sq km. So with a single 100 MHz channel in mid-band spectrum, a service provider could offer a 1 Gbps service to ALL of the farms under that cell footprint with an oversubscription ratio of 78. Of course, other service offerings can be created as desired by the service provider to accommodate their customers' needs.

In addition to transporting IoT data, the C-Band rural broadband connection enables new applications and services such as remote desktop. This type of service can provide high-power, low-cost and low maintenance (for the user) computing resources to users and locations where full computing platforms may be neither desirable nor cost-effective.

## **5. AR/VR User Interface**

The main aspect of the user interface is an application that uses the IoT data from the IoT Application System and overlays it on multiple 360 degree visual environments.

The Unity Application was designed for PC-based VR HMD's as well as Android-based platforms. It is currently being ported to more lightweight standalone HMD's to decrease the initial investment.

The application consists of the following features:

- 1x 4K 360 degree livestream (35 – 50mbps)
- 1x Mavic Pro Drone livestream (5mbps)
- 2x 720p Thermal camera livestreams (3mbps)
- 6x 1080p Security camera livestreams (30mbps)

Since the number of 360 cameras and total bandwidth was limited, some of the 360 environments were pre-recorded. However, the pre-recorded environments still allow for an immersive sense of presence (Figure 7). In both versions of the application, the user can change their surroundings and “travel” between the main farm complex and the surrounding hog houses (Figure 8).


Security cameras give a live video feed of the hogs during the day, while Thermal cameras provide the ability to see them at night. Proper heat-sensing thermal cameras have the potential for predicting livestock health issues and death.

Another implemented user interface feature allows alerts to be set up to alert the farmer when parameter thresholds are triggered by sensor readings. Heaters and fans can be automatically adjusted to account for fluctuations. Two-way interaction with the sensors can also be implemented to allow for manual adjustment while the farmer is remote.


If the farmer is somehow unable to visit the surrounding hog houses or crop fields, drone interaction has been integrated into the application. The livestream from the drone's camera is viewable in the application (Figure 9). Currently, the livestream operates over LTE, resulting in an input latency that doesn't allow for precision remote control. Hence, predetermined waypoints were added to reduce the presence of latency and decrease the possibility of motion sickness.



The current iteration of the VR application interaction is agnostic of the 360 degree video feed; the visualization of your hands in virtual reality originates at the user end and is not streamed. While the 360 video feed itself may contain latency, it will never promote motion sickness as your body's representation in virtual reality is maintained in real-time. As HMD's become smaller and the processing power moves to the edge, latency issues will need to be dealt with in order to prevent motion sickness. The common accepted max value for latency in VR streaming is approximately 20ms.



**Figure 7 - Temperature and Moisture sensors in the grain bins reduce spoilage and increase yield.**



**Figure 8 - Security cameras show live footage of the hogs. All pertinent information is available virtually.**



**Figure 9 - Need to visit a site in real-time? Remotely fly a drone to a pre-determined waypoint.**

## Conclusion

Smart Agriculture is one example of a burgeoning multi-billion dollar industry that can dramatically improve cost savings while presenting an opportunity for a new MSO business vertical through rural broadband. The deployed technology behind the Charter R&D 5G-powered Smart Farm is, at its core, a true rural broadband connection with an immersive user environment for displaying mission critical IoT sensor data which enhances operational efficiency and productivity for a real working farm in two agricultural categories. By utilizing available smart devices, netcams, drones, and 360 videos, all with unlicensed spectrum, the platform can be adapted to any environment in a variety of different business verticals. This allows a cable MSO such as Charter to address new business opportunities in a cost efficient and scalable manner to enable new valuable service offerings while still leveraging their existing infrastructure and work force.

## Abbreviations

5G	5 <sup>th</sup> Generation mobile network
AP	access point
bps	bits per second
FLIR	Forward Looking InfraRed
GHz	GigaHertz
GPS	Global Positioning System
HD	high definition
IoT	Internet of Things
JSON	JavaScript Object Notation
LoRa	long range (wireless IoT standard)
MHz	MegaHertz



MSO	Multiple System Operator
HMD	head-mounted display
RF	radio frequency
RTSP	Real-time Streaming Protocol

## Bibliography & References

LoRa Alliance website including specifications, <https://lora-alliance.org/>, retrieved July 15, 2019.

Federal Communications Commission, “Radio Spectrum Allocation”, <https://www.fcc.gov/engineering-technology/policy-and-rules-division/general/radio-spectrum-allocation>, retrieved July 15, 2019.

Federal Communications Commission, “Bridging The Digital Divide For All Americans”, <https://www.fcc.gov/about-fcc/fcc-initiatives/bridging-digital-divide-all-americans>, retrieved July 15, 2019.

Charter Communications, “Farm of the Future”, <https://policy.charter.com/charters-farm-of-the-future/>, retrieved July 15, 2019.

# **ATSC 3.0: A Look at the Infrastructure and Possible Impact that Next-Gen TV Will Have on Cable Operations**

A Technical Paper prepared for SCTE•ISBE by

**Ralph Bachofen**

Vice President of Sales and Marketing

Triveni Digital

Princeton, NJ

1 609 716 3502

[rbachofen@trivenidigital.com](mailto:rbachofen@trivenidigital.com)

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Overview of the ATSC 3.0 Standard .....	3
Basic Architecture of ATSC 3.0 and Headend Scenarios .....	4
Conclusion .....	7
Abbreviations.....	7
Bibliography & References .....	7

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Headend scenario showing legacy QAM system with OTA reception .....	5
Figure 2 - Headend scenario showing legacy QAM with direct connection to station via fiber link .....	6
Figure 3 - Headend scenario showing advanced CATV system with IPTV capabilities.....	6

# Introduction

ATSC 3.0 is near completion in the broadcast television industry, and is on its way to full-scale deployment in 2020 and beyond. What are the implications for cable operators?

This paper will provide a general overview of the next-generation broadcast television standard, including new features and service enhancements made possible with the technology. It will cover what is involved with receiving ATSC 1.0/ATSC 3.0 within cable headends. While in the beginning since broadcasters will simulcast ATSC 1.0 and 3.0, there may be no needed change to cable headends, over time increasingly broadcasters and multichannel video programming distributors (MVPDs) can collaboratively provide enhanced services for cable customers.

Several different setups for receiving ATSC 3.0 programming from broadcasters will be examined for cable headends, including near- and long-term. In particular, the paper will discuss headend requirements for HEVC IP transport stream delivery, and system architectures for combination of off-air ATSC 3.0 service reception with enhancements delivered from the broadcaster over alternative connections.

## Overview of the ATSC 3.0 Standard

ATSC 3.0 promises to help broadcasters deliver advanced over-the-air (OTA) and over-the-top (OTT) services, including ad replacement, service guides, emergency communication support, addressable content delivery, interactive program enhancements, and data broadcast applications. While ATSC 3.0 was specifically developed for broadcast transmissions, cable operators can benefit from its video and audio improvements.

Currently, with ATSC 1.0, broadcasters deliver an MPEG-2 transport stream (TS) containing MPEG-2 Video and AC-3 Audio. MPEG-2 Video encoding is not as efficient as HEVC, the compression technology used to transmit ATSC 3.0 services. HEVC offers a significant improvement in bandwidth efficiency compared with the MPEG-2 Video Codec, which will allow broadcasters to meet the growing consumer demand for superior video quality, including 1080p, high frame rate (HFR e.g. 60 f/s), high dynamic range (HDR) and Ultra HD (UHD), at lower bitrates.

If cable operators have an HEVC infrastructure in place, they can receive HEVC encoded content from broadcasters and subsequently deliver better video quality to subscribers while minimizing bandwidth consumption on their network.

On the audio side, ATSC 1.0 supports 5.1 channel surround sound using Dolby Digital's AC-3 format. ATSC 3.0 will provide a much more immersive audio experience leveraging AC-4, the next-generation format from Dolby. This is another exciting capability that cable operators can take advantage of.

There is also potential for cable operators to leverage applications delivered as part of the ATSC 3.0 OTA broadcast to provide interactive enhancements to programs. These applications are implemented using standard W3C technologies including HTML5, JavaScript and CSS which could execute directly on the cable STB. This would require extensions to the operating environments on the STBs but environments such as the Comcast RDK already support similar applications. In any event, ATSC 3.0 broadcasts will be constructed such that they will still be viewable without the interactivity since not all ATSC 3.0-compatible TVs will necessarily support the interactive environment.

# Basic Architecture of ATSC 3.0 and Headend Scenarios

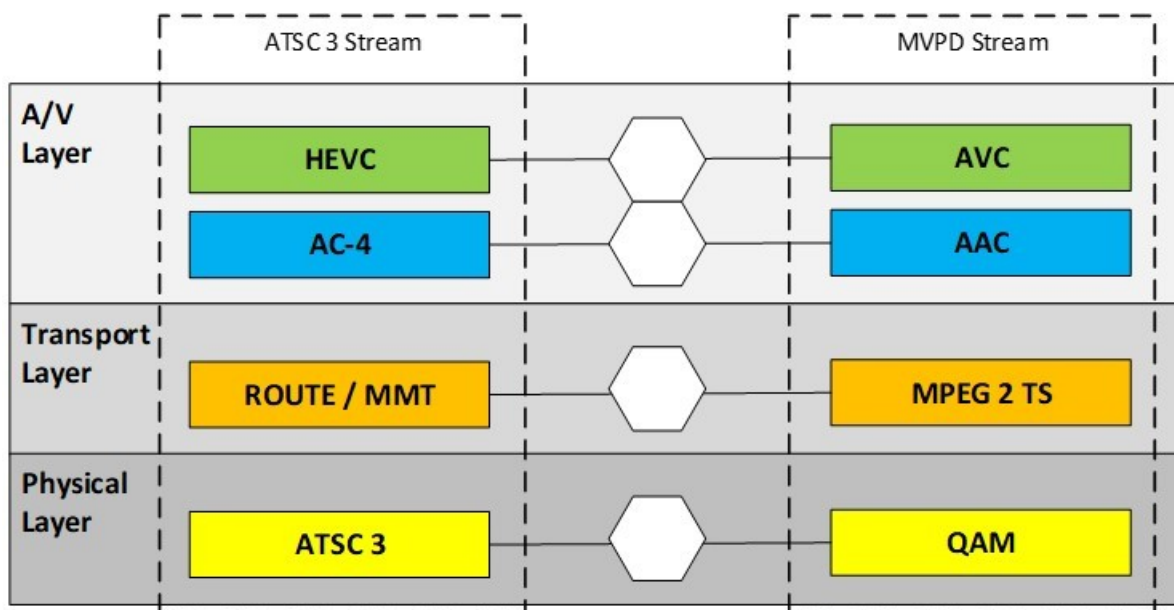
The basic architecture of ATSC 3.0 is different and more complex than ATSC 1.0. Rather than receiving an MPEG-2 TS, cable operators will begin receiving IP signals carrying HEVC and AC-4.

The FCC mandates that high-power and Class-A broadcasters simulcast ATSC 1.0 and 3.0 for at least five years; therefore, the immediate impact on cable headends will be minimal. If cable operators want to continue receiving MPEG-2 TS as usual, they can do so without making any infrastructure changes. However, this mandate does not apply to low-power broadcasters.

But, at some point, operators will need to make infrastructure changes, especially if they want to benefit from the improvements in video and audio quality provided by ATSC 3.0. Moreover, in the long term, the transition to ATSC 3.0 may impact how cable operators acquire signals from broadcasters. The ability to receive an ATSC 3.0 system will improve generally of the DMA because of the new capabilities of the ATSC 3.0 physical layer. However, since new transmission systems may be deployed in different locations, there may be the need to tune antennas to see the new ATSC 3.0 transmission.

There are two ways cable operators can receive ATSC 3.0 signals: Over the air or via direct connection over broadband. The direct connection can be either a legacy MPEG2 Transport Stream carrying MPEG2 Video and AC-3 Audio, an MPEG2 transport stream carrying HEVC Video and AC-4 audio, or a set of one or more ATSC 3.0 IP streams containing HEVC and AC-4. What is sent by the broadcaster will be dependent on the agreements in place.

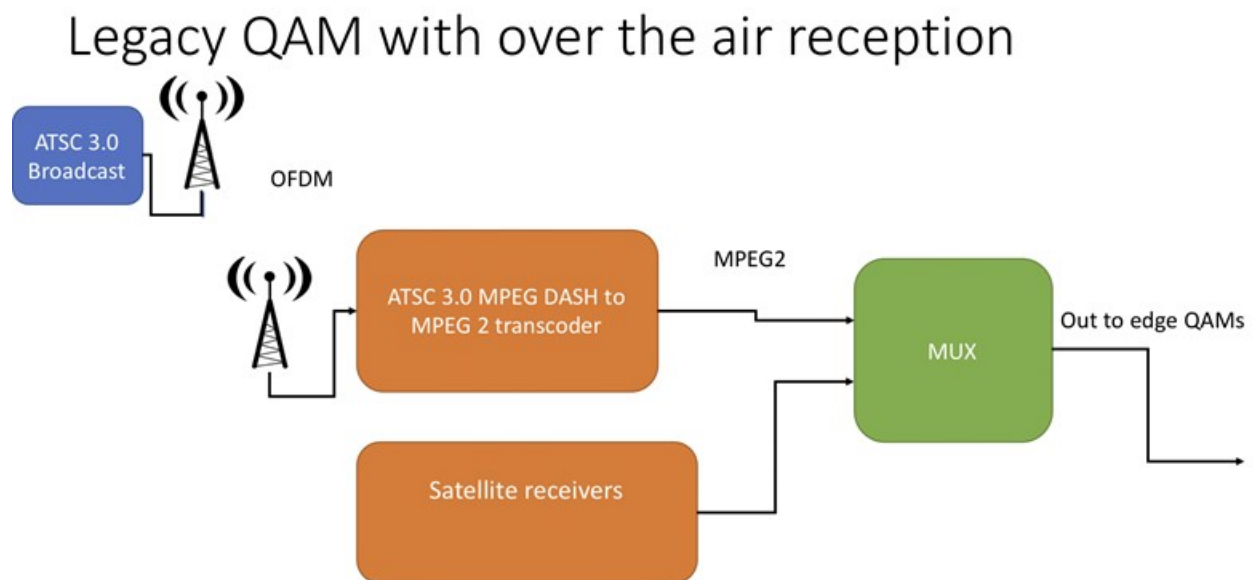
The ATSC 3.0 suite of standards includes an A/V layer with codec support, a transport layer, and physical layer. Within each layer, shown in the figure below, it's likely that one or more of the content of these layers needs to change. However, making changes to one layer does not necessarily impact another. The figure below shows the various layers and the aspects of each layer that need to be transformed.



Operators with legacy MPEG-2 over QAM systems will need to convert the broadcast signal from ATSC 3.0 to MPEG-2 TS. This can be done at the cable headend receive site with a transcoder that converts ATSC 3.0 to MPEG-2 TS. An operator that receives multiple feeds from broadcasters will need to ensure the signals are muxed together correctly. Alternatively, headends with a fiber connection to the station master control can receive the legacy MPEG-2 TS directly on the fiber link, similar to how they do in the ATSC 1.0 environment. Finally, CATV systems with advanced IPTV systems can use the higher order signals on their cable system without downconverting to MPEG-2 TS (case by case basis).

For IPTV, the system is typically not going to be based on ROUTE or MMT; therefore, a transformation needs to be done. If the IPTV content is delivered in HLS or another adaptive bitrate transport mechanism, buffering will need to occur to convert from the push broadcast transport scheme into the pull IPTV transport scheme.

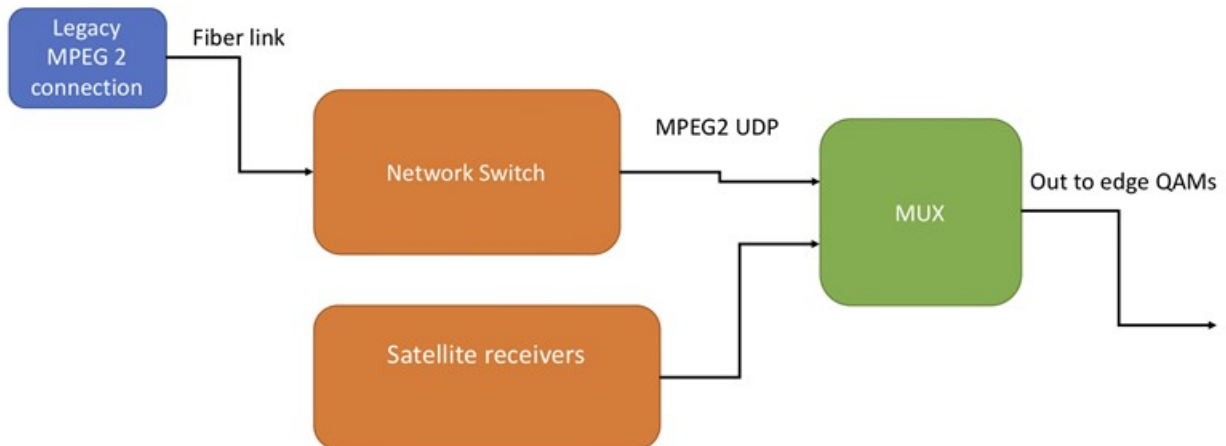
Figure 1 shows a headend scenario where an operator is outside of the ATSC 3.0 reception area and is only able to receive the local broadcast via ATSC 3.0. Here the operator can transcode the ATSC 3.0 signal into MPEG-2 TS using a receiver-converter. This does not require any significant infrastructure changes on behalf of the cable operator.



**Figure 1 - Headend scenario showing legacy QAM system with OTA reception**

Figure 2 demonstrates an MVPD that is operating a fiber link between the legacy MPEG-2 TS connection and the network switch. Cable operators with a fiber link will need to ensure there is an ATSC 3.0 compliant modulator at the broadcast television station or in the headend. Fiber signals are received in the cable headend by a demodulator. After the link is demodulated at the headend, it will go into the same ATSC 3.0 to MPEG-2 TS receiver and be muxed in with the other MPEG-2 services as HEVC and AAC streams.

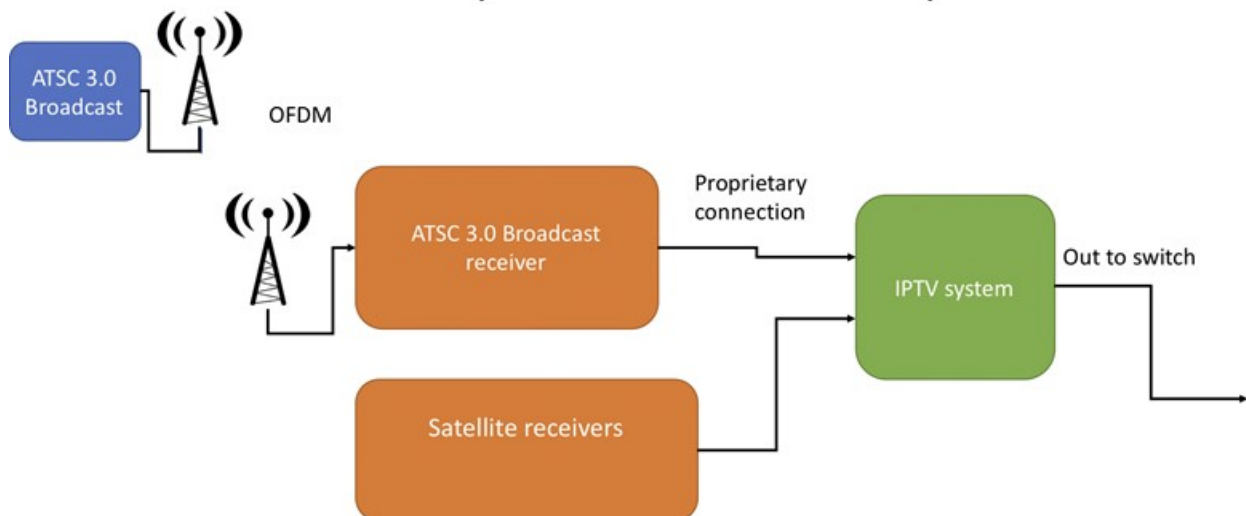
## Legacy QAM with direct connection to the station



**Figure 2 - Headend scenario showing legacy QAM with direct connection to station via fiber link**

The headend scenario presented in Figure 3 show the infrastructure changes that cable operators need to make in order to take advantage of the advanced audio and video capabilities of ATSC 3.0. Instead of transcoding from ATSC 3.0 to ATSC 1.0, or DASH to MPEG-2, the operator fully receives the ATSC 3.0 broadcast signal. In this case, the operator must upgrade its infrastructure to support HEVC and AC-4.

## Advanced CATV system with IPTV capabilities



**Figure 3 - Headend scenario showing advanced CATV system with IPTV capabilities**

## Conclusion

In the short term, cable operators will not need to make any changes to their infrastructure, as broadcasters will simulcast both ATSC 1.0 and ATSC 3.0 signals. Yet, at some point in the foreseeable future, changes will be required. Communicating directly with broadcasters about what changes they'll be making to their own infrastructure is essential. Operators will also want to learn how long broadcasters plan to simulcast and when they plan to fully transition to ATSC 3.0.

Down the line, operators may need a transcoder or CATV system with IPTV capabilities. Once operators are able to support the ATSC 3.0 suite of standards, they can start delivering some of the exciting improvements it offers, including better picture quality and more immersive audio.

## Abbreviations

HDR	high dynamic range
MVPD	multichannel video programming distributor
OTA	over the air
OTT	over the top
TS	transport stream
UHD	Ultra HD

## Bibliography & References

“ATSC 3.0: The Impact on Cable Operations;” *Broadband Technology Report*, R. Bachofen and C. Hamilton; March 2017; <https://www.broadbandtechreport.com/docsis/headend-hub/article/16447712/atsc-30-the-impact-on-cable-operations>

“ATSC Recommended Practice: Conversion of ATSC 3.0 Services for Redistribution (A/370)”, Doc. A/370:2019, Advanced Television Systems Committee, Washington, D.C., xx August 2016.



# **Broadcast and Digital Evolution**

## **The Evolution of Delivering to Any Screen**

A Technical Paper prepared for SCTE•ISBE by

**Stuart Kurkowski, PhD**  
Distinguished Engineer  
Comcast Incorporated  
1515 Wynkoop Street, Denver, CO 80202  
303-503-2680  
Stuart\_Kurkowski@comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Evolving the Definition of “Quality” .....	3
A Unique Path for Every Playback .....	4
A Case Study on Broadcast Market Change .....	5
As Files Inflate, Compression Gets Complicated.....	6
You’re Only Live Once: Building Redundancies Into Event Planning .....	6
“Simulated” Live Events and Virtual Channels: New Monetization Models that Feels Familiar.....	8
The Evolving Relationship Between Advertisers and Providers .....	8
Conclusion .....	9
Abbreviations.....	9
Bibliography & References .....	9

# Introduction

Ensuring that content is available and that ads can be placed on any screen is table stakes to content providers and advertisers alike. Now, new questions are surfacing like how to keep the quality standards of broadcast on all devices, how to simultaneously deliver to both broadcast and online destinations, how to reach targeted audiences across platforms, and how to see a holistic reporting view of all content?

Broadcast and digital destinations are no longer separate in the minds of consumers, therefore backend workflows have transformed to seamlessly serve both. The intelligent merging of broadcast and digital technology combines workflows, delivery, and reporting to get content and advertising to any screen.

Advancements in video quality and global delivery are changing the industry at a more rapid pace than ever before. In today's data-heavy world of multi-platform video, profitability and quality assurance are hard-won metrics that must be continually reassessed and updated. Providers of every ilk are contending with the squeeze on both sides of the equation – the need for faster operations, real-time advertising, and lower cost in order to divert more resources into what matters most, the creation or curation of compelling content.

More than ever, video service providers are ruled by three masters. At one corner, there's the fundamental requirement to deliver video at the highest quality, both in picture and playback consistency. At the next corner, dollar signs abound as businesses try to offset ballooning costs of content development and acquisition by streamlining and automating video delivery and commerce workflows. At the third corner is the need to bring learnings from an extremely fragmented media landscape into a holistic view that produces better performance measurements and, ultimately, better commercial performance of every piece of content, be it short-form ad, linear channel, or on-demand.

This paper doesn't intend to encompass all of the changes and challenges of multi-platform delivery. It does, however, aim to highlight some ways in which the convergence of broadcast and digital delivery is – with growing pains – evolving the way that providers deliver and monetize content, and how consumers consume it.

## Evolving the Definition of “Quality”

“Broadcast quality” is something that every viewing experience strives for, but in a world where 4K screens are racing towards ubiquity, the hallmark of broadcast quality is its consistency. Paired with the capabilities and image quality of today's devices, as consumers we've really got an embarrassment of riches to choose from. For video service providers, that translates into a desire to connect more meaningfully with viewers at every playback, with repeatable, reliable, more personalized experiences of the highest quality.

Last year, Comcast Technology Solutions published a research paper called *Survive and Thrive: the Changing Environment for Content and Video Providers*. The objective of the survey was to better understand how content and video providers are leveraging online resources and delivery mechanisms, and what they perceive as significant challenges in today's evolving market. Survey respondents listed “discoverability” as the top business challenge they're facing. Maintaining consumer attention in today's market speaks directly to the need for every playback experience to execute flawlessly, lest the viewer go elsewhere. And, depending on the type of program, the consumer might not even have to choose

something else to watch. Instead of watching the newest superhero flick on service A, for example, they can just switch over to service B where they can watch the same movie with less buffering. More and more individuals every year have access to multiple video services, whether served through an online subscription, or over a set-top box or game console – which means that one program might be accessible by one consumer in multiple ways (like a new movie that’s just been wide-released for download / streaming). If a show is buffering, lagging, or of otherwise poor quality on service A, but can be watched with reliable quality on service B, it has a definite impact on future viewing decisions. Consistent quality is a huge factor in maintaining audience relationships, especially as more premium services enter the market.

With the accelerating commercial adoption of 4K / UHD / HDR screens, the eventual implementation of the ATSC 3.0 broadcast standard (and the arrival of consumer devices that can take advantage of it), we are seeing [the beginning of] a massive increase in available 4K programming across the entire provider ecosystem. Though clearly a significant improvement in picture quality, the sheer size of ultra-high definition (UHD) video files brings new challenges to storage and delivery. As consumers acquire more of a taste for (and an ability to play) video and audio of increasingly high quality, every touchpoint in the delivery workflow is more heavily taxed by having to compress, store, transcode, and deliver files that are up to four times bigger than those previously handled. That makes for a lot of change that needs to be managed all at once, especially when the premium prices for premium content create premium expectations. It’s no wonder that 92% of our respondents in the same survey said that meeting or exceeding broadcast quality and reliability is still important.

IP-based delivery is light-years ahead of where it was just a few years ago, but as more and more consumers develop an appetite for “higher and higher quality” video, the widening data stream continues to push the performance limits of every platform. Advertisers and investors understand the value of reaching a gargantuan audience, but they also – rightly – seek assurances that viewers are getting an experience that supports their brand objectives.

## **A Unique Path for Every Playback**

Delivery to a disparate sea of screens remains a fascinating technical challenge; but the end-point device is not the only thing to solve for. Content providers now have to deliver content, not just to MVPDs, but also directly to consumers -and virtual MVPDs (vMVPD), using a solution that provides broadcast reliability. Business rules and playback rights are not just different from device to device, but quite literally from view to view – and everything needs to be solved for in real time.

So, for a company that aspires to deliver a winning experience anywhere, what does a successful new delivery model look like? Global operators and content providers require cost effective, high-quality methods to acquire, package, distribute, assemble content in order to reach more audiences, manage metadata, and provide more economical options to viewers. There are innovations across management and delivery workflows that can be leveraged:

- Content feeds from all partners can be acquired and aggregated in a more automated, cost-effective way that reduces degradation and time to market.
- A more intelligent multi-CDN architecture can employ predictive analytics for proactive decisioning and smarter routing and caching of content, as well as better performance measurement.
- A carrier-grade cloud infrastructure makes it easier than ever to define sources, upload assets, define policies, and create and deliver OTT streaming services of all types (linear playout, live, VOD, time-shift TV/cloud DVR), directly to viewers in hours instead of months.

- Live media necessitates buffering, error correction and synchronization schemes to overcome the inherent weakness of packet-based IP networks. Virtualized IP networks address this issue with control-layer decoupling from bespoke hardware to optimize the entire network for the demands of video.
- Ad insertion and automated ad commerce tools can be incorporated in order to improve monetization efforts, keep ad-supported content fresh and relevant to viewers, and in turn more valuable to advertisers.
- Metadata needs to be easily normalized and manipulated to SCTE 224 standard, so users can easily and accurately define sources, upload assets, define regional or global distribution policies, and manage the unique permissions and programming schedules for each screen.

## A Case Study on Broadcast Market Change

There's so much disruption coming from so many different sources that it's a constant challenge to try and identify the ramifications of change before they impact a viewing experience. It's not just a matter of digital destinations working to meet and exceed broadcast quality. Broadcasters need to provide the level of interactivity and ad insertion capabilities that digital brands offer. Going back just a few years, advertising within a broadcast feed was a relatively straightforward affair. Today, SCTE 35 and SCTE 104 information must be included so that ad insertion at the delivery end-points can occur. And not only occur, but occur at the right point in the video, giving the client a perfect user experience with frame accurate switching.

The service areas of regional broadcast stations are no longer simple, line-in-the-sand maps, as these broadcast constituents' content with content rights and restrictions that don't always line up. Here is an example of an issue we were faced with recently:

We conducted a test with a local television station that was a local affiliate for a national broadcaster. The television station was testing delivery of its content to a national virtual MVPD (vMVPD). We offer a way for organizations like this local station to manage their linear metadata rights, and they wanted to test the product's capability to control the rights restrictions for the station's content. This station normally broadcasts its signal via antenna to a roughly 30-zip code area around a major US city. This station's only blackout scenario happened during a national sporting event, where the territories of two teams overlapped in one subset of their 30-zip code audience area. Traditionally, this television station would black out the telecast of the sporting event to 9 zip codes in the northeastern section of their footprint. As such, their blackout rules had a restriction for these 9 zip codes, which our metadata management tool could easily manage.

Now it was time to share this video feed from the local station with the national vMVPD. When the vMVPD sent out the content during the blackout scenario, the 9 restricted zip codes were blacked out as expected; but roughly 42,000 other zip codes could see the content – essentially the rest of the United States. This was not the desired outcome, because obviously, a local affiliate should not be the station providing major sports content to a national audience.

As a result, this local station now had to get out of the mindset of thinking just about the 30-zip code area and parts of its restrictions to a national perspective. In reality, the blackout restrictions now ought to have been 39,976 zip codes – which includes the 9 restricted codes and the rest of the country. This subtle change in mindset can be rather large and intimidating for a small local television station, but this is just the start of what is required in this new digital/vMVPD ecosystem we have before us.

Additionally, if you are distributing internationally, then country boundary enforcement is another concern that must be addressed. And although some content providers are addressing this through spreadsheets, emails, and grids, these solutions don't scale and don't satisfy the requirements of the new vMVPDs that don't have large staffs and require machine-to-machine conveyance for their linear rights information.

## **As Files Inflate, Compression Gets Complicated**

For video compression, we're faced with a diversity of codecs serving a diversity of devices and file types, as storage and delivery work to keep up with the capabilities of today's screens. The codecs commonly used today are the result of a rapidly evolving and fragmented market – one full of constituents with different priorities beneath the fundamental stated goals around reliable quality, cost control, and profitability.

MPEG-4 Advanced Video Coding (AVC) has been utilized longer and is more prevalent than most, but it lacks the compression efficiency needed for ultra-high-definition (UHD) programming. Ultimately it will have to give way to codecs that can handle the never-ending quality march uphill – UHD, 4K, HDR, 8K, and whatever comes next. The codecs that are currently in use each have their own unique positives and negatives:

- High Efficiency Video Coding (HEVC), essentially the next-gen MPEG codec, is currently in use by over two billion devices. It solves the challenges presented by UHD, but is shackled to a challenging licensing structure. Complicating the picture are over 1500 HEVC patents in the U.S. across an ocean of companies, turning intellectual property (IP) rights into a potential quagmire.
- AV1, developed by the Alliance for Open Media, is a strong contender but has only been available for a few short years. It is royalty-free and might be an attractive option as the scale of devices deployed increases, but broadcasters largely have embraced HEVC over AV1, somewhat do to the fact that ATSC had to make decisions before AV1 was completely available, so we will see how this changes in the future.
- Essential Video Codec (EVC), otherwise known as MPEG-5, is in development but is essentially an unknown quantity from a licensing standpoint until 2022 at the earliest.
- Versatile Video Coding (VVC) is designed to handle immersive media applications and is essentially the best on paper, showing efficiencies of over 50% above HEVC. It's targeted for completion in Q4 2020. Much is still to be determined with VVC, especially the overall efficiencies that are gained, depending on tool selection, etc.

The bottom line is that video providers of all stripes have to contend with, and accommodate, compression technologies that are as fragmented as the market itself. As Streaming Media wrote in April 2019, "The market needs a codec that delivers high compression efficiency, reasonable encoder complexity, broad decoder support, and a clear licensing scheme. There are clear questions of scale and the market will need to move to an all-software base, but just as cinematographers now pick digital sensors as they would film stock for different applications, so broadcasters and service providers might one day be able to cherry pick, swap, and replace codecs with automated ease."

## **You're Only Live Once: Building Redundancies Into Event Planning**

Live programming is definitely an area where broadcast and digital are working together to improve experiences across the board. One of the biggest challenges with streaming live events over digital is that,

with so many potential failure points, solving for them all is an expensive proposition that most services don't account for. Backup plans and redundancies are expensive, so they're often victims to cost-cutting initiatives. Broadcast provides that straightforward, low-latency, consistent quality that digital delivery is benchmarked against. When millions of dollars (or more) are on the line, it can be an exceptionally costly corner to cut.

That said, online destinations can provide deeper interactivity that enhances the audience experience; and again, broadcast feeds need to provide advertisers with the same ad insertion capabilities as digital feeds. The best delivery architecture is one that leverages the strengths of broadcast and digital into a seamless multi-platform workflow – piloted by expert event engineers who plan every transmission and quarterback the communication from start to finish. Look at it this way:

- IP-only: lower cost, complex failsafes required, “best efforts” performance.
- Broadcast: more expensive, rock-solid best practices, high availability and reliability.
- Strategic use of both: trusted, cost-effective performance that supports both linear television and cloud experiences.

It's impossible to overstate the value that an experienced team brings to a top-tier live event; not just in overall quality, but also in cost savings. A complex live-streaming event, such as a high-profile awards show, is a great example. There are cameras everywhere: the red carpet and interview areas out front, around the audience, the main stage, backstage, etc.

A well-planned and executed event will map out all redundancy paths, but also will prioritize the use of broadcast signal. For example, the main stage cameras would be captured via broadcast to ensure quality, but lower-quality ancillary feeds could be leveraged to control costs. Satellite and digital captures are then converged into a master control so that the complete program can be compiled, encoded, and distributed across platforms.

The volume of live programming is going through the roof, with media brands of all stripes working to realize – and monetize – experiences that bring viewers closer to a true “front row experience” than ever before. That said, there are no do-overs in live video; so before the world tunes in, it's worth the extra planning to ensure that viewers on any platform get the kind of quality that keeps them in their seats.

## **“Simulated” Live Events and Virtualized Channels: New Monetization Models that Feel Familiar**

The hybridized broadcast / digital model is really a best-of-both worlds architecture that not only bakes in the video quality and process redundancy expected of a major live event, it also opens up opportunities to appeal to viewers in new ways. A program can be created that captures the immediacy and advertising allure of a live event, while also incorporating existing video assets into one seamless program. A live studio audience portion, for instance, could be woven into a playlist with sections of non-live material, encoded and then distributed just like a linear television feed. It's a cost-effective way of producing a special “one night only” event, or a way for a large organization – a coast-to-coast church broadcast, for instance – to maintain a live “feel” with the dependable quality that keeps viewers engaged.

Taking this idea into a 24x7 application results in a “virtualized channel” – a great way to turn existing on-demand assets into an always-on OTT experience that can operate as a standalone destination, or as a live feed that augments the merchandising of VOD content. It can be a particularly attractive way to build or expand niche audiences by compiling a programming schedule of curated content for a particular fan base or demographic. Content providers can extract more value out of their existing content with approaches like these; provided that they've got a system in place to manage content rights and restrictions with alternate content and electronic programming guides that accurately reflect the specific flavor of virtualized channel that's delivered to each screen. Monetization of content also improves – the new linear channel can be provided to multiple vMVPDs, increasing both viewership and revenue from consistent advertising opportunities.

## **The Evolving Relationship Between Advertisers and Providers**

Technology innovations aren't just opening up new opportunities for video advertising; it's changing the way that business is being done, and how providers and advertisers interact.

Programmatic advertising – the ability to automate media purchases based on viewer data – has not only given advertisers a way to improve ad performance at scale, it also provides broadcast and digital providers a more efficient commerce model with which to sell media space at scale. Addressable technology takes this model one step further, using advanced audience analysis and segmentation to focus ad delivery all the way down to the screen level. In other words, two people in one home could be watching the same program on different screens, but see two different ads.

In just the last couple of years, the advertising conversation has moved well past broadcast vs. digital delivery, and into deeper discussions about how to coordinate and optimize campaigns across the spectrum of video destinations, how to automate media commerce from both demand and supply sides more effectively, and how to further personalize experiences. Advertisers rely more heavily than ever on the data that they can get from providers so that each program, each destination, can be viewed holistically in order to more accurately understand how an ad's targeting or creative assets could be changed to improve performance. Innovations are changing industry norms from so many different directions that it has stretched the ability of companies to respond. The question “how do you get your content in front of your customers on the screen they want, when they want it?” is morphing into a much



deeper question, “how do you get your content in front of your customers on the screen they want, when they want it, in a way that makes it more relevant, valuable, and actionable?” There’s a lot to figure out, but it’s going to be an exciting ride:

- On the delivery side, new destinations and advertising opportunities are showing up regularly. When will providers and advertisers need to be ready to serve ads and content to new connected-home devices? When digital assistants like Alexa, Siri and Google start making purchase suggestions (on content or on anything else), how will that impact workflows?
- On the commerce side, automation and targeting are infusing operations with unprecedented levels of sophistication – and complexity. How much of your work will move towards programmatic and dynamic processes? How can dynamic creative optimization (DCO) be intelligently implemented to tailor ad service based on real-time viewer insights?

On both sides, manual and disparate processes are hampering business, draining time and resources. And they are inherently error-prone. How could we create an environment that brings everyone’s work into closer proximity, eliminates redundancies, and enables self-directed service and sophisticated real-time decisioning?

## Conclusion

The future for content providers and advertisers is sure to find them working more closely in concert, as the push for more personalized content experiences forces every organization to implement faster ways to manage complexity at scale, and more reliable ways to understand how each viewing experience impacts the viewer.

## Abbreviations

AVC	Advanced Video Coding
CDN	Content Delivery Network
DCO	dynamic creative optimization
EVC	Essential Video Codec
DVR	Digital Video Recording
HEVC	High Efficiency Video Coding
MVPD	Multichannel video programming distributor
OTT	Over the top
UHD	Ultra high definition
vMVPD	Virtual multichannel video programming distributor
VOD	Video On Demand
VVC	Versatile Video Coding

## Bibliography & References

Pennington, Adrian. April 4, 2019. *Video Codecs Today: Minefield, Muddle, or Multiple Choice?* StreamingMedia. Retrieved from <https://www.streamingmedia.com/Articles/ReadArticle.aspx?ArticleID=130999>

# **The Broadband Network Evolution Continues – How Do We Get To Cable 10g?**

A Technical Paper prepared for SCTE•ISBE by

**John Ulm**

Engineering Fellow

CommScope

22 Topside Rd, Moultonboro NH 03254

john.ulm@commscope.com

**Tom Cloonan**

CTO- Network Solutions

CommScope

2400 Ogden Ave.- Suite 180, Lisle, IL 60532

tom.cloonan@commscope.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Content .....	5
1. Cable 10G™ – A Journey, Not a Destination .....	5
2. Network Capacity Planning for 10G.....	7
2.1. Traffic Engineering for 10G Networks .....	7
2.1.1. The “Basic” Traffic Engineering Formula .....	7
2.1.2. Broadband Subscriber Traffic Consumption .....	8
2.1.3. Max Service Tiers in 10G Era.....	10
2.2. Spectrum Planning – What to Do with the Legacy Video Spectrum and Other Questions? .....	11
2.3. Network Capacity Modeling for 10G Downstream .....	11
2.4. Selective Subscriber Migration Strategies .....	14
3. Outside HFC Plant Considerations and Logistics.....	16
3.1. Technology Options to Enhance HFC Bandwidth Capacity .....	16
3.2. Cleaning Up the Plant to Maximize bps/Hz for D3.1 .....	18
3.3. Can We Achieve Cable 10G with Existing 1 GHz Taps? .....	18
4. Migrating from DS Only 10G to More Symmetrical 10G .....	21
4.1. 204 MHz Frequency Division Duplex (FDD) High-Splits .....	21
4.2. Full-Duplex DOCSIS (FDX) .....	21
4.3. Frequency Division Duplex (FDD) – 300 or 396 MHz Upstream Splits.....	23
4.4. Soft-FDX – Static or Dynamic .....	24
4.5. Addressing Cable 10G with Blended Fiber Deep and PON Systems .....	25
5. CPE Considerations.....	25
Conclusion .....	27
Acknowledgements .....	28
Abbreviations.....	29
Bibliography & References .....	31

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – 10G™ Logo .....	5
Figure 2 – Tavg, Average Subscriber Downstream Consumption.....	8
Figure 3 – Tavg, Average Subscriber Upstream Consumption .....	9
Figure 4 – Tavg, Downstream:Upstream Ratio.....	10
Figure 5 – 1218/204 MHz System – Spectrum Utilization .....	12
Figure 6 – 1218/204 MHz System – Subs per SG, DOCSIS Spectrum Needs.....	13
Figure 7 – 1218/204 MHz System – DOCSIS Usage: Tmax, Tavg, IP Video .....	13
Figure 8 – Innovation Adoption Lifecycle.....	14
Figure 9 – Downstream Growth with Multiple Service Tiers .....	15
Figure 10 – Sample Tap: Inter-amp Attenuation and Bit-Loading.....	19

Figure 11 – Tap Capacity for Various HFC Cascade Lengths.....	20
Figure 12 – Impact of Amp Spacings .....	20
Figure 13 – Full-Duplex DOCSIS (FDX) Spectrum Band Options .....	22
Figure 14 – 1218 MHz DS Capacity Impact from Ultra-High Upstream Splits.....	23
Figure 15 – Tap Capacity to 1410 MHz for Various HFC Cascade Lengths .....	24

# Introduction

On January 7, 2019 at the Consumer Electronics Show (CES), NCTA – The Internet & Television Association®, CableLabs®, and Cable Europe® introduced the cable industry’s vision for delivering 10 gigabit networks, or 10G™ – a powerful, capital-efficient technology platform that will ramp up from the 1 gigabits per second (Gbps) offerings of today to speeds of 10 Gbps and beyond – to consumers across the globe in the coming years.

10G is a goal, a lighthouse in the distance towards which all MSOs and vendors can use as a beacon in the night to steer their boats safely away from danger. It will take some time to get there. It is actually only a single point in a continuum of improvements that will occur in the future. It is not the end-point, it is an interim point. We will likely push on beyond that 10G point in the future. The focus of this paper is the migration needed over the next 7-10 years using existing technologies to achieve the 10G goals. Another paper [CLO\_2019] takes a deeper look into the future with some potential new technologies to see where the industry might be in 15-25 years beyond 10G.

To reach the 10G goals might require an aggregate of last-hop technologies, including: HFC, DOCSIS, PON, and Wireless which are all likely candidate technologies. It turns out that 5<sup>th</sup> generation (5G) mobile and 10G HFC are quite complementary and will help one another. This is discussed further in [ULM\_2019].

But in getting to 10G, what does it take for the DOCSIS/HFC system to actually deliver on the 10G Service Level Agreement (SLA) promise? Many MSOs do not really know what it takes to do 10 Gbps downstream, let alone 10 Gbps symmetrical.

The cable industry is going through unprecedented technology changes starting with the introduction of DOCSIS 3.1 (D3.1); then Distributed Access Architectures (DAA) such as Remote PHY (R-PHY); DOCSIS Full-Duplex (FDX); Low Latency DOCSIS (LLD) initiative; and now Extended Spectrum DOCSIs (ESD) efforts. These are all building blocks to help us reach Cable 10G. But what are the logistics to get on the right path?

This paper will discuss the importance of these technologies and address the following migration topics:

- 10G network capacity planning
- Outside plant considerations and logistics
- CPE considerations

Cable 10G is an ambitious initiative, but it is in keeping with the broadband heritage to which we’ve grown accustomed. We can get there with the right roadmap.

# Content

## 1. Cable 10G™ – A Journey, Not a Destination

A quick introduction to 10G™ can be found on the CableLabs website [www.cablelabs.com/10g](http://www.cablelabs.com/10g) and given below:



**Figure 1 – 10G™ Logo**

### **What is 10G™?**

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 up to 10 Gbps, but also lower latencies, enhanced reliability, and better security in a scalable manner.

### **Why Do We Need the 10G Platform?**

Our digital future will stall without a platform that can meet our needs. While we don't know what the next trend will be, we do know the internet will be central to it. By advancing device and network performance to stay ahead of consumer demand, 10G will provide a myriad of new immersive digital experiences and other emerging technologies that will revolutionize the way we live, work, learn, and play. Like the saffron in paella, or the milk in a latte, our industry's networks and innovations are the crucial ingredients in creating a better future for humanity.

In the downstream (DS) direction, 10G will be helpful in providing delivery of immersive video services (virtual reality & augmented reality for "Holodeck Experiences"). It will also be useful for providing more "snappy" services. For example, downloading:

- A two-hour HD movie in 3-4 seconds (vs. 5-7 min @ 100 Mbps)
- A two-hour 4K UHD movie in 12-15 seconds (vs. 20-25 min @ 100 Mbps)
- A large gaming program such as Call of Duty's Black Ops (~100 GBytes) in 90 seconds
  - instead of an 2½ hours @ 100 Mbps.

In the upstream (US) direction, it could be used for providing more "snappier" services again, but it could also prove to be very useful in enabling low-latency transport. The extra bandwidth (BW) helps enable Proactive Grant Services (PGS) to accelerate US delivery. That lower latency can permit 5G mobile backhaul and midhaul, and if the latencies drop low enough, it could even permit 5G fronthaul.

10G may also enable many different commercial applications such as remote medical procedures.

There are four key attributes to the 10G platform:

1. Speed
2. Latency
3. Security
4. Reliability

10G's promise of faster speeds, more capacity, lower latency and greater security will enable and help fully realize a wide variety of new services and applications that will change the way millions of consumers, educators, businesses and innovators interact with the world. Much of the underlying technology has already been specified or is a work in progress.

The speed attribute will leverage technologies such as :

- DOCSIS 3.1
- DOCSIS 4.0 - FDX, Extended Spectrum DOCSIS
- PON
- Coherent Optics
- Advanced WiFi including WiFi 6 (a.k.a. 802.11.ax)

Some applications driving the need for lower latency DOCSIS include: gaming; VR/AR (avoiding nausea); CoMP; and autonomous navigation. Latency is a function of packet processing times, queuing times, transmission times, and propagation times. We can improve all areas. CableLab's Low Latency DOCSIS (LLD) project includes ideas in all these areas. The existence of 10G BW capacity also helps, because higher bandwidth capacity leads to less congestion and also permit new techniques like Proactive Grant Service (PGS) to expedite BW grants in the upstream (US). Work is also being done in low latency mobile X-haul and low latency Wi-Fi.

Security is an integral part of 10G. Work continues at CableLabs in Micronets, secure downloads, & MACsec. This will become part of the new DOCSIS 4.0 specification. Vendors are also working to make more secure systems with separate, isolated processors & memory in chips.

With respect to reliability, new DAA node designs of the future will likely be adding in redundancy in processing and redundancy in NSI-Side links to Leaf Switches in CIN as Moore's Law improvements in gate density help. Reliability is being addressed by proactive network maintenance (PNM) and dual channel Wi-Fi. This will improve observability and redundancy and A/I to monitor (PNM, more data analytics in DAA nodes, redundancy in ring of DAA)

And for all four 10G attributes (i.e. speed, latency, security, reliability), it is equally important that they scale... on all markets.

The remainder of this paper will now focus on the speed and capacity required to achieve the 10G goals.

## 2. Network Capacity Planning for 10G

The “10G” in the announcement is for 10 Gigabits per sec (Gbps). But what exactly does that mean as there are different ways of measuring capacity? For example, Liberty Global’s Virgin Media division in the U.K. ran tests earlier this year over a 10G EPON network and demonstrated users getting 8.5 Gbps throughput. It turns out that this is extremely close to the theoretical maximum throughput for 10G PON technology. “10G” PON has a raw physical rate of 10 Gbps but there is ~15% overhead from the PHY and MAC layers. So the customer actually nets 8.5 Gbps from the 10G PON.

Our analysis first takes a look 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. Finally, some subscriber migration strategies are discussed that will help the transition to 10G platforms.

### 2.1. Traffic Engineering for 10G Networks

The CommScope (formerly ARRIS) team have been providing industry leading research in traffic engineering for many years which was most recently highlighted in [ULM\_2017]. Some additional references of note include [CLO\_2014], [EMM\_2014], [ULM\_2014], [CLO\_2016], [ULM\_2016] and [CLO\_2017].

#### 2.1.1. 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 busy-hour 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  $0 \leq K \leq \infty$ , 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. 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 less than 1.0.

### 2.1.2. Broadband Subscriber Traffic Consumption

ARRIS/CommScope 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 2, shows the average subscriber downstream consumption, DS Tavg, during peak busy hours for a number of MSOs over a ten year period. At the start of 2019, DS Tavg was approaching the 2 Mbps barrier.




**Figure 2 – Tavg, Average Subscriber Downstream Consumption**

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 4-5 years, this group of MSOs had an average downstream traffic growth that had been just under 40%. On a yearly basis, traffic growth can be very sporadic. It is not uncommon to see high growth in one year followed by little growth the next. This is exactly what happened in 2018 and 2019. The 35%-40% trend could be used as a longer term guideline on downstream traffic consumption,

but others feel the growth rate may continue to decline over time. DS Tavg could reach ~20 Mbps in roughly 7 to 8 years assuming growth rate remains constant, but it might only be ~15 Mbps in 10 years if growth rate trends continue to decline.


Interestingly, the upstream traffic is growing at a significantly slower rate as shown in Figure 3. During the same ten year period, the upstream Tavg generally grew at less than 20% compound annual growth rate (CAGR). Notice that there was a significant spike in 2018 followed by a dip in 2019, but the resulting long term CAGR still stayed under 20%.



**Figure 3 – Tavg, Average Subscriber Upstream Consumption**

Traffic has also become more asymmetric with video applications driving downstream consumption [EMMEN\_2014]. The DS:US ratios are shown in Figure 4. As of 2019, the average DS:US ratio is about 14:1, the MSO with the largest DS:US ratio seems to have stabilized around a 18:1 ratio.

The 10G platform will be opening up a world of new applications. Many of these will start driving upstream bandwidth consumption. This may come from a plethora of IOT devices in the home, or maybe inexpensive HD resolution video cameras pushing content to the cloud. For the purposes of this study, we will assume the upstream growth will roughly match the downstream growth over the next 7-8 years. This will push US Tavg up to ~1.5 Mbps in that timeframe.



**Figure 4 – Tav<sub>g</sub>, Downstream:Upstream Ratio**

### 2.1.3. Max Service Tiers in 10G Era

So, what kind of service tiers will subscribers enjoy in this new high 10G bandwidth era? The T<sub>max</sub> value from the basic formula helps define the Service Level Agreement (SLA) that the MSO can offer to their customers.

As previously discussed, the 10G PON provides a net downstream capacity of ~8.5 Gbps to the consumer. This capacity might support a downstream SLA of 8 Gbps. The service group (SG) utilization (i.e. N<sub>sub</sub> \* Tav<sub>g</sub>) for a 64 subscriber PON might grow to a bit over 1 Gbps in the 7-8 year window. That means a consumer with a 8 Gbps SLA will have a QoE coefficient of K=0.9 to 1.0 which is reasonable for this relatively small SG size. The next section will discuss what is needed on an HFC system to support this same DS SLA.

Getting to a true 10 Gbps downstream SLA that is equivalent to 10G Ethernet will mean providing slightly greater than 10 Gbps network capacity. This will push the PON networks into next generation PON technology (e.g. 20+ Gbps). Because HFC can incrementally add capacity with additional spectrum, there are certain downstream scenarios that will be discussed where existing 1218 MHz HFC might be able to hit this target. In general, future technologies such as 1.8 and 3.0 GHz HFC plants, are out of scope for this paper and are discussed further in [CLO\_2019].

Choosing the upstream SLA is a more complicated matter. As can be seen with the DS:US consumption ratio, there might be a 20:1 difference between the two. However, in the new 10G era, there may be a need for gigabit US SLA tiers with high burst rates, even if the US consumption might be much lower than downstream.

Looking at PON systems, they offer both symmetric and asymmetric data rates. GPON provides 2.5 Gbps downstream data rates with 1.2 Gbps upstream data rates for a 2:1 ratio. The IEEE 10G EPON

downstream might be paired with either a 1G or 10G upstream for 10:1 or 1:1 ratios. In the ITU world, XG-PON pairs 10 Gbps downstream with 2.5 Gbps upstream (i.e. 4:1 ratio) while XGS-PON provides a symmetric 10 Gbps in both directions for 1:1 ratio.

HFC systems have traditionally been extremely asymmetric, but these trends are changing. In the following sections, a range of upstream SLAs are considered to pair with the 8 Gbps DS SLA with a discussion on the technology trade-offs needed for each.

## **2.2. Spectrum Planning – What to Do with the Legacy Video Spectrum and Other Questions?**

In order to achieve 10G goals on HFC systems, it will require MSOs to get on a path to converting most or all of its legacy video spectrum over to DOCSIS high speed data (HSD) spectrum. Most MSOs have already taken the step of removing analog video channels from their plants. That is the first big step and biggest bang for the buck. Now is the time to start transitioning the rest of the digital video spectrum and there are a number of possible options.

The long term goal is to get to a true IP video infrastructure where every home has IP set-top boxes (STB) or IP capable devices for playing video. This provides a common infrastructure across all access technologies including PON, DOCSIS and wireless. However, the economic realities will mean that legacy STBs will be around for many years to come and need consideration.

For some MSOs with pre-dominantly MPEG-4 capable set-top boxes (STB), they can convert their MPEG-2 digital video content to MPEG-4 and roughly cut their legacy video spectrum usage in half. Most of that gain is in the HD content, so having all HD customers with MPEG-4 capable STBs is important.

If an MSO still has a significant amount of MPEG-2 only STBs, then they could relegate these to basic only video tiers and encourage customers who want advanced video tiers to migrate to newer technology (e.g. IP STB).


It turns out that an older technology, switched digital video (SDV), is having a bit of a renaissance and provides some very powerful capabilities to reclaim legacy video spectrum. This is detailed in [ULM\_2018]. Cloud-based SDV can be used to save as much as 400 MHz of spectrum. These savings will be critical as cable operators look to add multiple additional 192 MHz OFDM channels downstream and/or expand the upstream splits to higher thresholds. Pairing Cloud-based SDV with IP video migration can be very powerful.

## **2.3. Network Capacity Modeling for 10G Downstream**

Over recent years, there has been a slowing in the downstream usage growth rate (i.e.  $T_{avg}$ ) compared to the service tier growth rate (i.e.  $T_{max}$ ). This has a number of consequences including the network 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.

Below are some network capacity results from the CommScope network capacity modeling tools. This shows the potential capabilities for a 1218/204 MHz HFC plant. It begins with a 512 homes passed (HP) service group with 256 subs (i.e. 50% penetration). The max DS service tier starts at 1 Gbps and grows by 1 Gbps per year starting in 2022 until it finally reaches 8 Gbps DS SLA in the year 2028. For this case study, it is assumed the  $T_{avg}$  growth rate continues its gradual decline over the next decade. This will

leave Tav<sub>g</sub> at ~15 Mbps by the end of the decade. If the Tav<sub>g</sub> growth rate does not decline, then these dates might get pulled in by two to three years.



**Figure 5 – 1218/204 MHz System – Spectrum Utilization**

Figure 5 shows the spectrum utilization for the 1218/204 MHz plant. With the high-split, the downstream spectrum starts at 258 MHz. For this case study, a combination of IP video and SDV are used to reduce the legacy video spectrum requirements over a 5-6 year window. The figure shows how digital broadcast spectrum drops much faster than the corresponding growth in IP video. This spectrum savings allows the DOCSIS spectrum to grow with increasing Tav<sub>g</sub> and T<sub>max</sub> for a couple years without the need for a SG split.


By the year 2024, the 256 sub SG has maximized its downstream capacity. A larger SG will need a SG split by this time. However, a SG with 192 subs or less can survive for another couple years. Finally by 2027, all SGs need to be at 128 subs or less. The max subs per SG is shown in Figure 6. 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 7. It shows the amount of capacity needed for both DOCSIS 3.0 and DOCSIS 3.1. Each of these is then further broken out into separate components:


- QoE Delta (i.e.  $K \cdot T_{\max}$ ),  $N_{\text{sub}} \cdot T_{\text{avg}}$  and IP Video

As can be seen, the T<sub>max</sub> component dominates over time.

The lower red line in Figure 7 shows the amount of DOCSIS 3.0 capacity in the system while the upper red line shows the combined 3.0 + 3.1 total capacity for the system.



**Figure 6 – 1218/204 MHz System – Subs per SG, DOCSIS Spectrum Needs**



**Figure 7 – 1218/204 MHz System – DOCSIS Usage: Tmax, Tavg, IP Video**


Perhaps the key point of this case study is that a 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.

## 2.4. Selective Subscriber Migration Strategies

At first glance, Nielsen's Law and 10G are a scary proposition such that HFC networks might be obsolete in five to seven years while it may take decades to build out an FTTP infrastructure. However, this is not the full story. As was shown in [ULM\_2016, ULM\_2014], Nielsen's Law only applies to the top speed tiers which is a very small percentage of the entire subscriber base, perhaps less than 1%. So, the key question then becomes, "What happens to the vast majority of subscribers on HFC who are not in the top speed tiers (a.k.a. billboard tiers) and when?"

The [ULM\_2014] case study looked at service tier evolution at a few cable operators. It turns out that for many service providers, their mix of service tiers follows a distribution similar to the technology adoption life cycle that was pioneered by Everett Rogers' *Diffusion of Innovations* theory in the early 2000's. This is where the term 'Early Adopters' was coined. See Figure 8.

Looking at a typical cable operator today, they may have a top billboard tier of 1 Gbps that is being taken by a small percentage of innovators. A slightly larger group of early adopters may be taking a 400 Mbps performance tier while the majority of subscribers are taking a 50-100 Mbps basic or flagship tier. Finally, there may be the technology laggards who have older modems (e.g. 2.0, early 3.0) that still have 10-20 Mbps economy tiers.




**Figure 8 – Innovation Adoption Lifecycle**

Perhaps the key finding from the [ULM\_2014] study is that the different service tiers are growing at different rates. While the top billboard tier continues to follow Nielsen's Law 50%, each subsequent lower speed tier is growing at a slower rate. Hence, the lower the service tier rate, the lower its CAGR.

Figure 9 maps out a possible scenario from the case study of the various service tier growth over the next two decades. While less than 1% of subs in the top billboard tier hit 10 Gbps in ~2024, the 15% of subs in

the performance tier doesn't hit that mark until ~2032. Notice that 85% of subscribers in the flagship basic tier and economy tier stay below this mark for several decades.



**Figure 9 – Downstream Growth with Multiple Service Tiers**

With a selective subscriber migration strategy, only the heavy users and subscribers on the top tiers need to get migrated to the new technology. For example, a leading customer (i.e. Innovator) wanting 10 Gbps DS SLA might be given a next generation PON or DOCSIS extended spectrum network connection. Over several years time, early adopters will start to join in and more given the new technology. The vast majority of subs remain on the existing technology.

It is important to note that 99% of the subscribers can still comfortably use today's DOCSIS technology on HFC a decade from now. With a selective subscriber migration strategy, it is very important from a traffic engineering perspective to understand the behavior of the individual service tiers. But with this understanding in hand, selective subscriber migration can be used to extend the life of HFC for decades to come. When the time comes to migrate the premium tiers from traditional HFC to new technologies, it might be done with either FTTP technology or maybe extended spectrum HFC if it is viable by that time.

This is a crucial concept for our 10G migration strategy. Initially, 10G will only be for a fraction of a percent representing the innovators. This stage will then be followed by the early adopters that might represent 5% to 15% of the subscribers. Then finally 10G will be broadly consumed by the entire customer base. The 10G migration strategy can take advantage of this shifting over time which might take a decade.



### 3. Outside HFC Plant Considerations and Logistics

The cable industry has invested much into newer technology such as DOCSIS 3.1 and 1218 MHz capable components. How well will this support the 10G goals? It turns out that there are many variables involved that will be discussed here.

#### 3.1. Technology Options to Enhance HFC Bandwidth Capacity

The cable industry has many tools in their toolbox, each with its own strengths, advantages, and costs. A summary of possible options is shown in Table 1.

**Table 1 – Technology Options to Enhance Bandwidth Capacity**

Typical Options for last few years:	Ease & Cost	Tmax	Nsub*Tavg	Notes
Analog Reclamation	😊	😊😊😊 (DS only)	😊	Need DTAs
Node Segmentation, node splits	😊 to 😞		😊😊😊	HFC plant upgrades
DOCSIS 3.1	😊😊	😊😊		Need D3.1 CPE
<b>More Options for coming years:</b>				
Fiber deep – N+0, N+1; 1.2GHz DS, 85-204MHz US	😞	😊😊😊😊😊	😊😊😊😊😊	Long term strategic direction, often combined with DAA
MPEG-4 transition	😊😊	😊😊 (DS only)		If mostly MPEG-4 capable STB
IPTV transition	😊	😊 -> 😊😊😊 (DS only)	😊 (DS only)	Need IP capable CPE; gains depend on IPTV strategies
Cloud-based SDV	😊😊😊	😊😊😊 (DS only)	😊 (DS only)	Leverage existing legacy STB
<b>Future Options:</b>				
DOCSIS FDX	😞😞	😊😊😊 (US only)		Assumes DAA & N+0, too
FTTH, Ext Spectrum, FTTTap	😞😞😞	😊😊😊😊😊	😊😊😊😊😊	Next steps after fiber deep

Over the past decade or so, most operators have reclaimed analog spectrum and continued to do node segmentation and node splits as business as usual. It turns out that each addresses a different component in our traffic engineering formula. The analog reclamation frees up spectrum which is critical to offering

higher service tier SLA (i.e.,  $T_{max}$ ) while node segmentation and splits reduce SG size which directly addresses the  $N_{sub} \cdot T_{avg}$  component in our traffic engineering formula.

More recently, operators have started using DOCSIS 3.1 as they migrate to 1 Gbps downstream service tiers. D3.1 operates in today's existing HFC plants without any changes. It improves spectral efficiency (i.e., bps/Hz) enabling more capacity from a given amount of spectrum. D3.1 is also robust enough to operate in the roll-off region to provide additional bonus capacity. It is straightforward to implement without plant updates by deploying D3.1 cable modems and some SW/HW upgrades to the existing CMTS/CCAP platforms.

Over time, the D3.1 benefits will continue to grow as the plant is improved. The spectral efficiency improves significantly with fiber deep networks and it enables operation over wider frequencies (e.g., 1218 MHz downstream, 85/204 MHz upstream).

From a strategic perspective, all operators plan to push fiber deeper until it eventually becomes fiber to the home (FTTH). However, this is a multi-decade process. Some operators are considering a fiber deep HFC (with N+0, N+1 cascades) as the next major step along the way. But this can be a monumental task with pulling fiber deeper in the plant and increasing nodes counts by a factor of 10X to 20X. So, while this option gets top marks for increasing spectrum AND significantly reducing SG size for both upstream and downstream, its major costs and complexities will force this option to be done slowly over time.

If an operator is focused on more easily increasing DOCSIS downstream bandwidth capacity, there are several other options available in the near-term including:

- Migrating broadcast video to MPEG-4/H.264 on existing legacy STBs
- Migrating broadcast video to IP video delivery over IP STBs
- Migrating broadcast video to Cloud-based SDV over existing legacy STBs

The MPEG-4/H.264 option makes sense if all or the clear majority of the legacy STBs support MPEG-4 decoding. Overall, it provides a 2:1 gain in broadcast spectrum converted. Problems arise if there is still a substantial MPEG-2 only STB user base. That may limit the broadcast programs that can be converted (e.g., limited program tiers) or require those older STBs to be replaced.

The IP video migration includes a wide array of potential solutions. Some operators may only move video on demand (VOD) and select programming content to IP video while other operators may move aggressively to IP video everywhere. The potential bandwidth capacity gains from reducing legacy broadcast video can vary dramatically as well. In addition to an all new video infrastructure, the operator also needs to replace its CPE with IP video capable boxes. During the transition window, the operator must continue to support the legacy STBs, creating a bubble in their bandwidth needs.

The third option for reducing legacy broadcast spectrum is SDV. In the early days with large SGs and limited number of SDV EQAM modulators, the spectrum gains were somewhat limited (e.g., 84 MHz). However, with today's current technologies, Cloud-based SDV can be used more aggressively and achieve gains in excess of 400 MHz. A SDV case study is described in [ULM\_2018]. SDV also has the advantage over the other approaches in that it works on existing legacy STBs (MPEG-2 and/or MPEG-4) and requires minimal video infrastructure.

The final options for increased bandwidth capacity are looking much further into the future. Operators will be considering options such as FDX DOCSIS to enhance the available upstream spectrum and network architectures such as FTTH, fiber to the tap (FTTT) and DOCSIS extended spectrum (e.g., 1.8 and 3.0 GHz).

### 3.2. Cleaning Up the Plant to Maximize bps/Hz for D3.1

HFC networks have always been an evolving and changing infrastructure, repeatedly delivering bandwidth capacity increases to accommodate the needs of their various services (video, high speed data, and voice) in a “just-in-time” fashion. MSOs have long recognized that the HFC plant contains vast quantities of un-tapped bandwidth capacity that can usually be enabled in a gradual fashion using minor evolutionary transitions to various sub-systems within the plant. This low-cost evolutionary approach to network transition has long been preferred over more expensive revolutionary changes that attempt to change a large amount of the HFC plant equipment all at once.

The HFC plant still has an incredible amount of un-tapped bandwidth capacity that can be enabled in a very smooth and cost-effective fashion in the coming years without causing any unnecessary and premature plant augmentations. Let’s take a look at how existing technology can be exploited to meet some or all of the 10G goals.

DOCSIS 3.1 brings many benefits to the table, with vastly improved spectral efficiency being a key advantage over DOCSIS 3.0 technologies. A target configuration to support the max 10G downstream SLA of 8 Gbps would be 4x192 MHz OFDM channels bonded with 32 SC-QAM channels. This consumes 960 MHz of downstream spectrum with total capacity of ~8.6 Gbps at 4096-QAM modulation. The HFC system now provides basically the same total downstream capacity as 10G PON while providing backwards compatibility for DOCSIS 3.0 modems that will be around for a while.

The total amount of downstream capacity from the OFDM channels might vary by as much as 20%-30% depending on the communication channel quality as originally shown in [ALB\_2014]. There are several techniques that the operator can use to improve OFDM capacity. First, they can replace long analog fiber links (e.g. 40+ km) with digital optics of a Distributed Access Architecture (DAA) system. The 2<sup>nd</sup> item is to push fiber deeper and to shorten the cascade length (e.g. N+6/N+3 down to N+1/N+0).

If an operator has a 1218/85MHz HFC plant, they have enough spectrum to support the 8 Gbps DS SLA, provided they can squeeze the legacy video spectrum down to 150 MHz over time. This is very feasible using the techniques described earlier. With an 85 MHz upstream split, the operator may still be able to pair the 8 Gbps DS tier with a 400-500 Mbps US tier, which provides a reasonable 15:1 to 20:1 DS:US ratio.


However, some operators may want to support a 1 to 1.5 Gbps US tier to go with the 8 Gbps DS tier. At this point, the operator will need to consider a 204 MHz upstream split. Note that there is still 960 MHz of downstream spectrum available, but this requires that the legacy video spectrum has been migrated to all-IP video delivery over the DOCSIS channels.

### 3.3. Can We Achieve Cable 10G with Existing 1 GHz Taps?

As mentioned above, cable has been very successful at growing incrementally as capacity is needed. As the HFC system electronics are pushed to 1218 MHz, a question arises to what are the capabilities of the passives in the networks? It turns out that most HFC plants today have had 1 GHz taps installed over the last couple decades. This represents an extremely large installed base. Can we achieve the Cable 10G goals with these 1 GHz taps? That would make this transition much more cost effective.

The ARRIS/CommScope team analyzed a number of the most common taps in our research labs to get a better understanding of how they behave above their specified 1 GHz limit. Figure 10 shows the behavior for one of the ‘middle of the road’ taps.

As can be seen, it can support 12 bps/Hz (i.e. 4096-QAM modulation) up to and slightly beyond 1000 MHz. It dips slightly to 11 bps/Hz (i.e. 2048-QAM modulation) for another ~100 MHz. The roll-off starts to accelerate above 1200 MHz, but it can still support 9 bps/Hz (i.e. 512-QAM modulation). The bottom line for all the taps is that the 192 MHz OFDM channel above 1000 MHz had anywhere from 75% to 90% of the capacity of OFDM channels below 1000 MHz.



**Figure 10 – Sample Tap: Inter-amp Attenuation and Bit-Loading**


In looking at the impact of various 1 GHz taps, the team also analyzed how tap capacity varied as the cascade length increases. As can be seen in Figure 11, all five tap types were able to achieve 8.5 Gbps capacity for a N+0 plant. Introducing a single amplifier in a N+1 system causes a slight degradation but overall capacity is still above 8 Gbps.

As the cascade length is increased to N+3 and N+6, the capacity degradation continues. The N+3 capacity is roughly around 8 Gbps while N+6 drops even further close to 7 to 8 Gbps. In both cases, there is a ~500 Mbps difference between the best and worst tap.


All of these capacities were calculated for an amp spacing of 1050 feet (e.g. 6 taps at 175' each). The team then analyzed the impact of closer amps spacings of 700' and 875' (e.g. 4 or 5 taps at 175' each). This is shown in Figure 12.

Note that the capacity for 700' and 875' spacing is very similar. This is due to cable attenuation on shorter cables not being sufficient to pull the signal close to the thermal noise floor. So, the SNR of the node (plus distortion added by the amplifiers) dominates and determines the received SNR. Therefore, a 100' cable would give similar performance to a 500' cable.

For longer cable spacing, the attenuation is sufficient that the EOL (End of Line) SNR is influenced by the thermal noise floor. One you reach this point, any further increase in length reduces throughput. So there is a soft threshold effect in operation. In this example, the longer spacing of 1050' (e.g. six taps @ 175') results in about a 10% drop in capacity.



**Figure 11 – Tap Capacity for Various HFC Cascade Lengths**



**Figure 12 – Impact of Amp Spacings**

The conclusion is that operators can still get decent performance up to 1218 MHz using the existing 1 GHz taps. However, the capacity of a particular plant is going to vary based a number of key variables including:

- Cascade length (e.g. N+0 to N+6)
- Amplifier spacing
- Optical link (i.e. DAA or Analog Fiber links)
- Tap type

To reach the 10G DS goals, the operator will need to look across all these variables to see how they might reach 8 Gbps DS SLA. Depending on plant variables, it may be necessary to slightly lower customer tiers to 5 or 6 Gbps DS SLA until the plant can be upgraded.

## **4. Migrating from DS Only 10G to More Symmetrical 10G**

To support cable 10G DS SLAs, the operator will need a minimum of an 85 MHz upstream split. This supports a 400-500 Mbps US SLA and a DS:US ratios of ~15:1 to ~20:1. Becoming more symmetrical and supporting Gbps US SLAs will require much more upstream capacity and newer technologies.

### **4.1. 204 MHz Frequency Division Duplex (FDD) High-Splits**

The original DOCSIS 3.1 specification supports a 204 MHz Frequency Division Duplex (FDD) upstream split, resulting in a usable spectral range that is up to five times larger than current 42 to 65 MHz widths. The 204 MHz split supports two 96 MHz OFDMA channels that can enable 1 to 1.5 Gbps US tiers. This gets the 8 Gbps DS SLA to a DS:US ratios of about 5:1.

This approach continues to use Frequency Division Duplex technologies that separate upstream spectrum and downstream spectrum (with a guard-band in between). The change to 204 MHz FDD high-split operation typically requires changes to both the existing nodes and the existing amplifiers on the HFC plant. Sometimes, plug-in modules for new diplex filters can be utilized to offer this upgrade path. The diplex filters typically permit the downstream spectrum to pass signals beginning at 258 MHz, so the resulting guard-band creates a diplex filter “spectral penalty” of 54 MHz (i.e. 54 MHz of spectrum is unusable from 204-258 MHz). The downstream spectrum has also been reduced by 150 MHz, which is close to ~1.5 Gbps of downstream OFDM capacity.


One of the potential side issues that must be dealt with when 204 MHz FDD high-split is utilized is the passing of downstream out-of-band (OOB) signals to existing set-top boxes. To support this OOB capability wherever it is required, vendors are exploring options that would permit these downstream OOB signals to be passed through the high-split nodes and amplifier in the HFC network (even though it is in the upstream portion of the spectrum).

### **4.2. Full-Duplex DOCSIS (FDX)**

Getting beyond a 1 to 1.5 Gbps US tier will require additional technologies beyond the 204 MHz FDD upstream. Some recent work at CableLabs has focused on a new technology called Full Duplex DOCSIS (FDX) and recently moved into the DOCSIS 4.0 domain. FDX leverages echo canceller technology to allow simultaneous upstream and downstream operation in the FDX band. FDX is targeted at a fiber deep Node+0 DAA environment.

The FDX capability offers a fundamental benefit that permits upstream spectrum expansions to occur without causing reductions in downstream spectrum. One of the key FDX technology enablers, Echo Cancellation, is required in both the optical node and potentially in the cable modem (CM).

Echo cancelling is a well-known technology. FDX was originally an addendum to the DOCSIS 3.1 specification but has since become part of the new DOCSIS 4.0 specification [FDX\_PHY]. The technology proposes to have downstream and upstream D3.1 transmissions occurring in the same frequency band at the same time at the CMTS. In the FDX specification, the overlapping frequency bands can be in any of the following ranges starting at 108 MHz and ending at: 204 MHz, 300 MHz, 396 MHz, 492 MHz, 588 MHz, or 684 MHz as shown in Figure 13. These FDX bands shown in Figure 13 are in addition to the standard 5-85 MHz upstream that can be utilized as well.



**Figure 13 – Full-Duplex DOCSIS (FDX) Spectrum Band Options**

On a fiber deep Node+0 plant, the upstream OFDMA channel might net capacity as much as 10 Mbps per MHz. This means that a 108-300 MHz FDX system might support a 2 Gbps US SLA while the full spectrum 108-684 MHz FDX system might support a 5-6 Gbps US SLA. Using the full FDX band would enable the operator to offer such DS/US service tiers as 8 Gbps x 2.5 Gbps or a fully symmetric 5 Gbps x 5 Gbps SLA.

With the Node+0 architecture, the 108-1218 MHz of downstream spectrum might actually net over 10 Gbps of downstream capacity which means the 1218 MHz FDX system could be pushed to a true 10 Gbps DS SLA with 5 Gbps US SLA. The downstream is now on par with 10G Ethernet while supporting a 2:1 DS:US ratio.

Current FDX work is moving along well. Initial field trials were in 2018 and continued in 2019 with very promising results. Real-world deployments are targeted to take place in 2020.


FDX should work fine in Node+0 environments, but its ability to perform in Node+X environments is still under study. A number of operators are reluctant to jump to N+0 but are still interested in achieving more symmetrical upstream service tiers. The issues with FDX in this environment and possible alternatives are explored in [ALB\_2019]. Because of the FDX challenges in N+X HFC plant, other technologies are under consideration for those scenarios.

### 4.3. Frequency Division Duplex (FDD) – 300 or 396 MHz Upstream Splits

Perhaps the simplest way to increase upstream capacity is to just raise the upstream split even higher than the 204 MHz defined in the current DOCSIS 3.1 specification. Pushing it up to 300 or 396 MHz will add one or two more 96 MHz OFDMA channels respectively. Each OFDMA channel will add almost 1 Gbps to the upstream capacity. This would let an operator support 2, 2.5 or 3 Gbps US SLA.

However all of this comes at the expense of downstream spectrum. As the upstream split is pushed higher, the guard band region also grows larger. So every increase of 96 MHz to upstream spectrum results in the loss of ~120 MHz of downstream spectrum, which is why this option is often paired with an extended spectrum downstream (e.g. 1.8+ GHz). But what can be managed from a 1218 MHz system? Figure 14 shows the impact to the 1218 MHz downstream capacity for various cascade lengths as the upstream split is increased.

If the cascade length is short (e.g. N+0, N+1), then a 300 MHz split might still be able to squeeze out a 7.5 Gbps DS x 2.5 Gbps US SLA. Or maybe an operator would rather have a 396 MHz split to offer a more symmetric 6 Gbps DS x 3 Gbps US SLA.




**Figure 14 – 1218 MHz DS Capacity Impact from Ultra-High Upstream Splits**

With the reduction in downstream capacity due to the higher upstream splits, is there any additional way to squeeze out extra DS capacity? In Figure 10, some 1 GHz taps actually continue to provide bandwidth above 1218 MHz. In the example above, the tap actually can support 4 bps/Hz (i.e. 16-QAM modulation) up around 1300 MHz. Our analysis for all of the taps was extended to include an additional OFDM channel from 1218 to 1410 MHz, which would also require next generation modems too.



A couple of the better taps might provide an additional 1 Gbps of capacity above 1218 MHz in an N+0 plant, but this plant is FDX capable and the extra spectrum may not be needed. Those capacity gains might be cut in half for N+3 and be fairly negligible at N+6. The lesser taps have almost no additional capacity above 1218 MHz, even at N+0. So extra capacity is very dependent on tap type.



**Figure 15 – Tap Capacity to 1410 MHz for Various HFC Cascade Lengths**

#### 4.4. Soft-FDX – Static or Dynamic

Although promising research has been conducted at CommScope and other vendors on FDX capable amplifiers in the past year, a problem called Interference Group Elongation has been identified, see [ALB\_2019]. This problem causes serious issues with this traditional FDX in a Node+X proposal. In the end, it causes large Interference Groups to be created that span most of the length of each RF Leg on a node. If we still want to share upstream and downstream spectrum, this drives us to time division duplex (TDD) operation. We will discuss two variants of this called static soft-FDX and dynamic soft-FDX operation.

Soft-FDX can be considered a special mode of FDX where the interference group includes an entire RF leg that might span multiple amplifiers. The ‘soft’ adjective refers to the ability to change the location of the US/DS split easily (potentially via software). Soft-FDX helps in supporting high US speeds, which are occasionally demanded by users, without permanently locking the spectrum to the US which can severely affect the valuable DS spectrum that is used to offer many services including video and high DS speeds which are demanded more frequently than the US.

Soft-FDX can be either static or dynamic. Static soft-FDX refers to the case where the US/DS split location does not change frequently (e.g., on the order of months or years). On the other hand, dynamic soft-FDX refers to the case where the US/DS split location changes in real-time based on traffic demand

(on the order of milliseconds or seconds). For instance, in the dynamic soft-FDX case, when there is a need for more US spectrum to run an US speed test or upload as an example, the split changes to accommodate that and when the need for the added US spectrum goes away, the split changes back to reclaim the valuable DS spectrum. Both static and dynamic soft-FDX can be implemented using special assignment of the FDX RBA messages.

From a burst speed perspective, soft-FDX has the same burst capabilities as traditional FDX. So they can support the same SLAs as FDX. However, this is fine as long as both upstream and downstream are not bursting for sustained periods of time. Based on current upstream utilizations, this should not be a problem. Traffic engineering studies are currently underway to determine if soft-FDX provides adequate bandwidth capacity. It may require that nodes, amplifiers and CMs be able to quickly switch the directionality of the frequency band between the upstream and downstream directions (via rapid RBA switching).

In reality, traditional FDX uses TDD concepts within the HFC plant. The creation of Transmission Groups within FDX defines groups of neighboring modems that are “noisy neighbors,” and they are required to transmit using TDD approaches. Soft-FDX simply extends the size of a Transmission Group to be an entire RF leg on the node instead of a sub-section of the RF leg.

This soft-FDX approach may permit the use of soft-FDX amplifiers in Node+X systems ( $X>0$ ), and it may permit sharing of the upstream and downstream frequency bands inside of Node+X systems ( $X>0$ ).

In conclusion, the traditional FDX Node+X solution does not have a clear path to success in the 2020-2028 time-frame for Node+X ( $X>0$ ) architectures. On the other hand, if its traffic engineering performance is deemed to be acceptable, soft-FDX is probably a technology that could be implementable by the mid-2020s if bandwidth capacities require it.

#### **4.5. Addressing Cable 10G with Blended Fiber Deep and PON Systems**

With all of the previous HFC options (FDX, Soft-FDX), the cable system has become much more symmetric, but still falls short of a true 8 Gbps x 8 Gbps SLA. In addition, questions start to be raised on how to finally get to a true 10G Ethernet equivalent of 10 Gbps x 10 Gbps SLA? Even 10G PONs can not meet that goal.

One option that an operator can consider is to use a Selective Subscriber Migration strategy and start moving to a blended system with HFC fiber deep plus PON system. As discussed earlier, only a small percentage of early adopters may go for these fully symmetric 10G systems. Rather than switching their entire subscriber based to fiber to the home (FTTH), an operator might choose to only provide FTTH to these innovators and early adopters.

To accomplish this on demand (i.e. install FTTH within a few days or weeks of customer requesting that service), the operator will need to have the fiber extremely close to every customer’s home. It turns out that a fiber Deep HFC upgrade is an excellent stepping stone for this strategy. Not only does it provide cable 10G for its vast majority of subscribers, but the operator can now deliver beyond 10G via FTTH in a timely and cost effective manner.

### **5. CPE Considerations**

First generation DOCSIS 3.1 modems supported downstream capabilities of 2x192 MHz OFDM bonded with 32 SC-QAM (i.e. 3.0 channels) along with 2x96 OFDMA upstream channels. These D3.1 modems

have five times the capacity of 32x8 D3.0 modems. These might provide subscribers with up to a 4 Gbps DS x 1.5 Gbps US SLA.

However, our network capacity modeling results show that it is still very important to upgrade the majority of D3.0 modems to D3.1 over the next decade. This will allow the operator to convert SC-QAM spectrum to more efficient, higher capacity OFDM channels. If an operator chooses to keep a large percentage of their subscriber base on D3.0 modems, then they may lose 1-2 Gbps from their total downstream capacity.

Only customers needing more than 4 Gbps DS SLA will require a new 2<sup>nd</sup> generation D3.1 modem that has at least four 192 MHz OFDM channels in the downstream. At this point in time, it is likely that the next generation of modems may also be FDX capable with up to 684 MHz upstream support.

One key tenet in all of the proposed upstream extensions above is that every approach can leverage an FDX capable modem. So, ultra-high splits (300/396 MHz) and soft-FDX will all be relying upon and using the same modem technology as traditional FDX. This should help the industry drive modem volumes and economies of scale.

There was a brief discussion regarding use of the roll-off range from 1218 to 1410 MHz. If that path is chosen, it would require a future generation modem with Extended Spectrum capabilities (e.g. 1.8/3.0 GHz) for the innovators and early adopters who need that extra bandwidth burst.

As operators ramp up cable 10G services, the home networking piece will also become extremely important to be able to move up to 10 Gbps throughout the home. Recent advances with Wi-Fi 6 (a.k.a. 802.11ax) should be up to the task. These wireless data rates should match up well with cable 10G targets.

There are also advances with home wireless routers. The newest ones are tri-band where they use one band for dedicated mesh backhaul inside the home. This will become even more powerful as additional spectrum becomes available in the 6 GHz band.

# Conclusion

This paper takes a look at how cable operators can reach the 10G goals with 10G PON and 1218 MHz HFC technologies within the next decade. [CLO\_2019] takes a longer term look over the next 25 years past 10G. After accounting for all the different overheads (e.g. PHY, MAC, IP layers), the subscriber is actually getting a 8 Gbps SLA in a 10G world. Table 2 summarizes the various options 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 2 – Summary of 10G Access Network Options**

<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

As can be seen by these options, an operator can choose how symmetric they want their system to be. This will be driven by competitive market forces as well as new yet unknown upstream applications that may appear in the future.

These SLAs will be quite adequate for the vast majority of subscribers over the next decade. There may be a small number of innovators and early adopters that want to go beyond these service tiers latter in the

decade, but that can be handled with a Selective Subscriber Migration strategy that moves this small percentage to a next generation PON (e.g. 20+ Gbps) or to an extended spectrum HFC (e.g. 1.8 or 3.0 GHz).

Our traffic engineering and network capacity analysis shows that 1218/204 MHz technology meets the needs through the end of the next decade. While getting to fiber deep N+0 is a good long term strategic goal, a 500 HP node size, N+X system is still reasonable as long as it can be segmented.

If more symmetric upstream services are needed or desired (i.e. greater than 1.5 Gbps), then a migration to traditional FDX for N+0 or Soft-FDX for N+X is a reasonable path. These also restore some downstream spectrum (i.e. 108-258 MHz) that may actually enable a true 10 Gbps DS SLA.

Our investigations into 1 GHz tap technology show that operators can achieve the 10G goals with the existing installed base of taps. This will buy the operator more time before they need to pull the trigger and replace them. Hopefully the 1.8/3.0 GHz future taps will be cost effective by that time.

Finally, fiber deep and DAA become more important technologies at helping operators to achieve the 10G goals.

## **Acknowledgements**

The authors would like to gratefully acknowledge the assistance of Frank O'Keefe for his PHY Layer modeling and analysis, in particular for the 1 GHz tap data. The data was priceless.

# Abbreviations

BAU	Business as Usual
Bcast	Broadcast
Bps	Bits Per Second
CAA	Centralized Access Architecture
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
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
DCA	Distributed CCAP Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DWDM	Dense Wave Division Multiplexing
E2E	End to end
EOL	
EPON	Ethernet Passive Optical Network (aka GE-PON)
EQAM	Edge Quadrature Amplitude Modulator
FD	Fiber Deep
FDX	Full Duplex (i.e. DOCSIS)
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
HP	Homes Passed
HSD	High Speed Data
I-CCAP	Integrated Converged Cable Access Platform
IEEE	Institute of Electrical and Electronics Engineers
IEQ	Integrated Edge QAM
LDPC	Low Density Parity Check FEC Code
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
Ncast	Narrowcast

NFV	Network Function Virtualization
NSI	Network Side Interface
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing Access (Upstream)
OLT	Optical Line Termination
ONU	Optical Network Unit
OOB	Out of Band
OPEX	Operating Expense
OTT	Over the Top
PHY	Physical interface
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-MACPHY	Remote MAC-PHY
R-PHY	Remote PHY
RX	Receive
SDN	Software Defined Network
SG	Service Group
SCTE	Society of Cable Telecommunications Engineers
SNR	Signal to Noise Ratio
TaFDM	Time and Frequency Division Multiplexing
Tavg	Average bandwidth per subscriber
Tmax	Maximum Sustained Traffic Rate – DOCSIS Service Flow parameter
TX	Transmit
US	Upstream
VOD	Video on demand
WDM	Wavelength Division Multiplexing

# Bibliography & References

[ALB\_2014] A. Al-Banna et. al., “The Spectral Efficiency of DOCSIS® 3.1 Systems,” SCTE Cable-Tec 2014, SCTE

[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

[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

[ULM\_2019] J. Ulm, Z. Maricevic, “Cable 10G vs. Wireless 5G/CBRS – Foe or Friend? A Survey of Next Gen Network Directions”, 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



# **A General-Purpose Operations Cost Model to Support Proactive Network Maintenance and More**

## **Operations Value Model**

A Technical Paper prepared for SCTE•ISBE by

**Jason Rupe, Ph.D.**  
Principal Architect  
CableLabs®  
858 Coal Creek Circle  
303.661.3332  
j.rupe@cablelabs.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Model Description.....	4
1. More About the 11 State Repair Model.....	7
2. More about the Degradation Model .....	8
3. Parameters .....	8
3.1. Definitions .....	8
3.2. Transition Rates .....	10
3.3. States.....	11
3.4. Constants .....	12
4. Finding Parameters to Feed the Models .....	12
Model Use and Use Cases .....	14
1. Aligning the Repair Model to the Degradation Model .....	15
2. Modeling Plant Degradation as it Relates to Repair .....	16
3. Reactive and Proactive Relationship with Failure Rate .....	17
Model Study Results.....	17
1. Reactive Only versus Proactive and Reactive Mixed .....	18
2. Reducing Repeat Rates .....	21
3. Degradation Over Time and Optimizing Cable Plant Rehabilitation.....	24
Conclusions.....	31
1. Generalizing Results .....	31
2. Operator Uses .....	31
3. Enhancements.....	31
Abbreviations.....	33
Bibliography & References .....	33

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – A depiction of the 11 state repair model with proactive (green) and reactive (brown) states and transition rates shown.....	5
Figure 2 – A depiction of a birth-death process diagram as is applied to the cable plant degradation model.....	6
Figure 3 – A depiction of the 11 state repair model combined with the birth-death cable plant degradation model, showing the repair states at each degradation level, and how repair impacts the degradation.....	6
Figure 4 – A depiction of the first two degradation levels with the repair states and transition rates shown. ....	7
Figure 5 – A verbose depiction of the reactive lobe of the repair model, with the common state at the bottom, and reactive repair states above. ....	9
Figure 6 – A verbose depiction of the proactive lobe of the repair model, with the common state at the bottom, and proactive repair states above. ....	9

Figure 7 – Service availability over 100 Monte Carlo runs for a proactive and reactive maintenance system, and reactive only where the proactive work generates 50% more to 100% more reactive work when not handled proactively.....	19
Figure 8 – Cost rate over 100 Monte Carlo runs for a proactive and reactive maintenance system, and reactive only where the proactive work generates 50% more to 100% more reactive work when not handled proactively. ....	20
Figure 9 – Cost rate difference between the best reactive case to the proactive and reactive case for add degradation states, over 100 Monte Carlo runs.....	21
Figure 10 – Service availability over 100 Monte Carlo runs for a proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost to 110% more time and cost.....	22
Figure 11 – Cost rate over 100 Monte Carlo runs for a proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost to 110% more time and cost.....	23
Figure 12 – Cost rate difference between the proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 110% more time and cost, over 100 Monte Carlo runs. ....	23
Figure 13 – Cost rate difference between the proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost, over 100 Monte Carlo runs. ....	24
Figure 14 – Service availability over the degradation states, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.....	25
Figure 15 – Maintenance cost rate over the degradation states, with constant degradation of 20 per year, and 25% of repairs not fixing the problem. ....	25
Figure 16 – Cost over time per month starting in degradation state 38, with constant degradation of 20 per year, and 25% of repairs not fixing the problem. ....	26
Figure 17 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 20 per year, and 25% of repairs not fixing the problem. ....	26
Figure 18 – Cost over time per month starting in degradation state 38, with constant degradation of 10 per year, and 25% of repairs not fixing the problem. ....	27
Figure 19 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 10 per year, and 25% of repairs not fixing the problem. ....	27
Figure 20 – Cost over time per month starting in degradation state 38, with constant degradation of 20 per year, and 12.5% of repairs not fixing the problem. ....	28
Figure 21 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 20 per year, and 12.5% of repairs not fixing the problem.....	28
Figure 22 – Cost over time per month starting in degradation state 38, with constant degradation of 10 per year, and 12.5% of repairs not fixing the problem. ....	29
Figure 23 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 10 per year, and 12.5% of repairs not fixing the problem.....	29
Figure 24 – Probability distribution function at 12 months, starting in state 1, with constant degradation of 20 per year, and 25% of repairs not fixing the problem. ....	30
Figure 25 – Probability distribution function at 12 months, starting in state 20, with constant degradation of 20 per year, and 25% of repairs not fixing the problem. ....	31
Figure 26 – An adjustment to the 11 state repair model to allow for two types of reactive repair handling, as recommended by an operator.....	32

# Introduction

This paper presents a flexible but basic set of Markov and math models that allow estimation of availability, reliability, and operations costs for repair operations under various levels of plant degradation. As a result, these models can address several use cases that can help operators make decisions about how and whether to address certain operations problems.

We are motivated by several observations.

- Proactive network maintenance (PNM) is an advantage given to cable operators for managing operations expenses to target service reliability which results from the underlying technology, usually DOCSIS® networking. But many operators do not take full advantage of it.
- Addressing the root cause of a problem takes much more technician time, so comes at a cost; but there is a benefit to the overall plant condition as well. It is important to examine this tradeoff, like so many other tradeoffs related to operations costs.
- Operators can be under pressure to improve services, which drives deployment of new technologies; but it is difficult to know where and when to make these transitions, and operations impacts are difficult to quantify.

In this paper, we explain the models in depth, define a few use cases for them, run these models using reference parameters to explore some possible generalizations useful to all operators, explain some questions that operators can answer with these models, and relate some maintenance optimization knowledge to the results.

Hopefully, this paper will entice operators into examining more operations improvement possibilities and enlist CableLabs to help explore some of them.

## Model Description

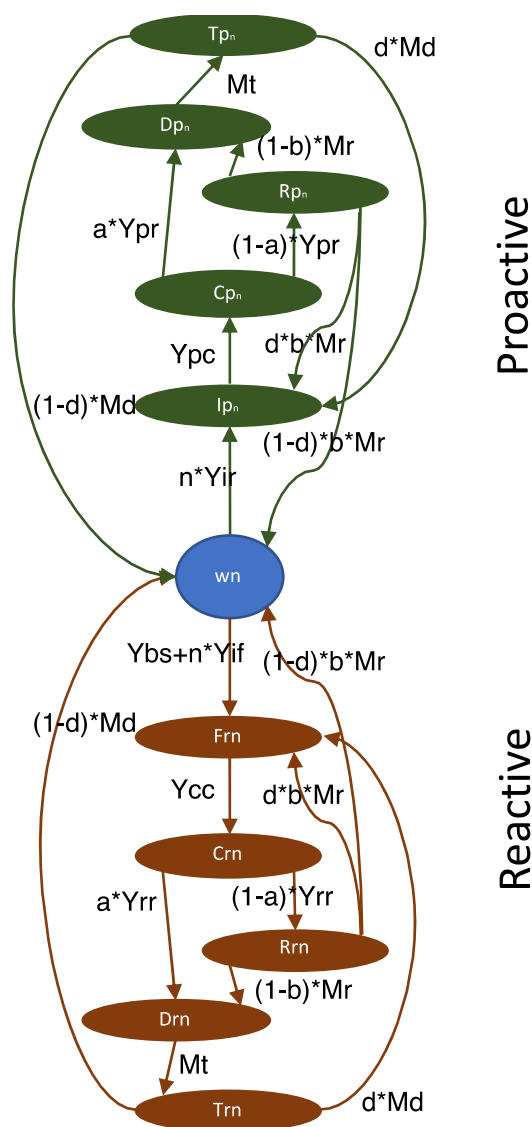
There are two major models contributed in this paper: an 11-state repair model, and a degradation birth-death model. Both are Markov models, which requires transition time distributions to be exponential (though this can be relaxed through extensions). We add a small amount of mathematical modeling and software implementation of these models to combine them to address many use cases in operations. For readers not familiar with Markov modeling methods, there are many excellent references including Taylor and Karlan's "An Introduction to Stochastic Modeling" [1].

The 11-state repair model describes the repair state for a section of cable plant that is in a static condition. With these 11 states, we can model both proactive and reactive repair, just proactive or reactive repair, or two parallel repair processes of any type at once. By exploiting features of this model, we can estimate the cost per unit time of an operations process it models. Therefore, we can use this model to compare operations solutions based on their effect on repair processing times, repair effectiveness, impact on customers, changes to costs, work handling, and other operations features. The 11-state model is shown in Figure 1.


Decoupled from the repair model, we have a degradation model based on a simple birth-death process. While the parameters that control the degradation and improvement of the condition of the plant can be modeled independently, we apply parameters to the repair model to better link them so that repairs that

find and remove impairments will improve the plant condition, and causes of degradation that arrive externally will result in more repair work. A simple birth-death model is depicted in Figure 2.


When combined, the models can be viewed as multiple 11-state repair models, one for each number of impairments in the cable plant being modeled. Degradations shift the model to a set of repair states associated with a higher rate of failure; likewise, a repair that removes an impairment (failure cause) will improve the condition of the cable plant, and be reflected in the model with a move to a set of repair states associated with a lower failure rate. A view of how the 11-state repair model combines with the birth-death degradation model is shown in Figure 3. Note that above state  $w_0$ , where there is no degradation of the plant yet, there are no proactive repair steps shown (in green) because there is no proactive work to do. Also, Figure 4 shows a close up of the first two degradation levels with repair states and transition rates. The proactive states above state  $w_0$  are shown translucent as they will not exist in most implementations.



**Figure 1 – A depiction of the 11-state repair model with proactive (green) and reactive (brown) states and transition rates shown.**



**Figure 2 – A depiction of a birth-death process diagram as applied to the cable plant degradation model.**



**Figure 3 – A depiction of the 11-state repair model combined with the birth-death cable plant degradation model, showing the repair states at each degradation level, and how repair impacts the degradation.**

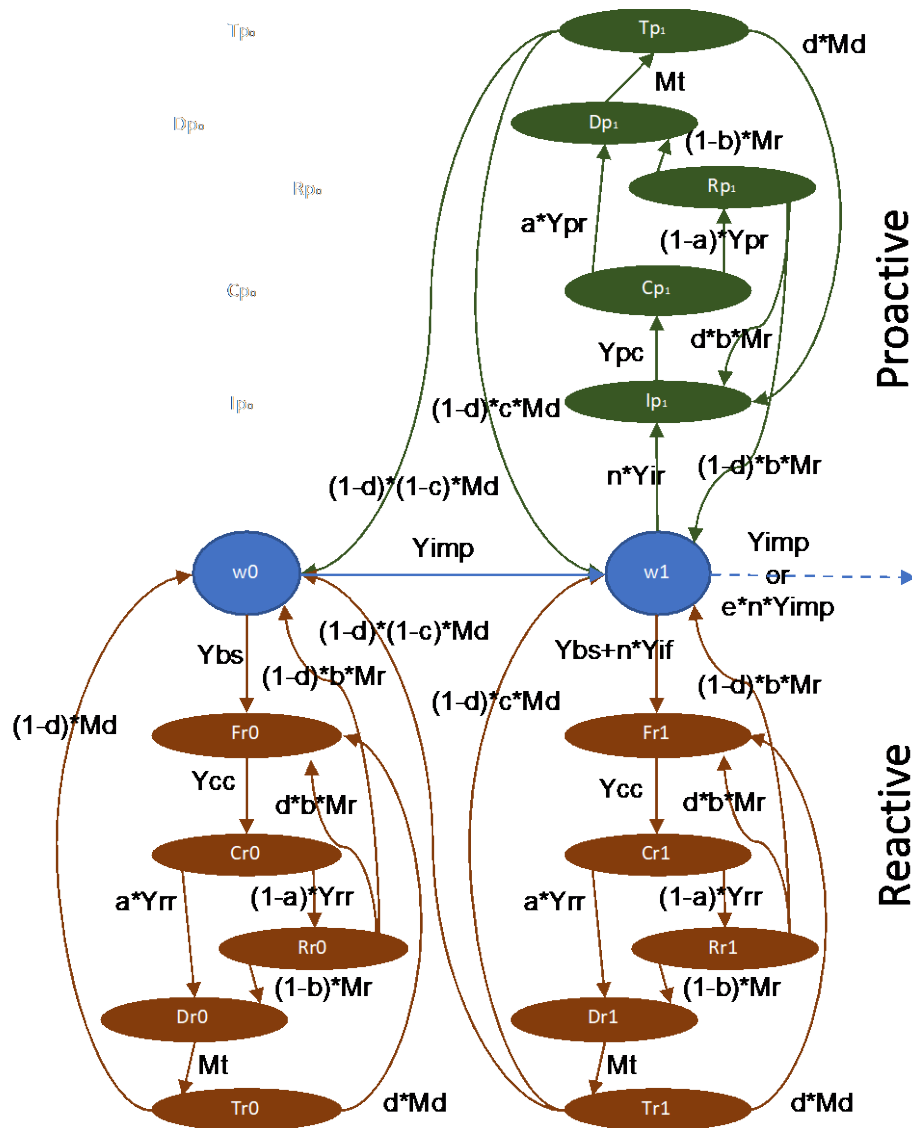


Figure 4 – A depiction of the first two degradation levels with the repair states and transition rates shown.

## 1. More About the 11-State Repair Model

The 11-state model is generally described as follows. It models the repair state of a section of plant serving a large enough number of customers to capture all customers potentially impacted by the largest such failure group, and small enough that it is unlikely that two failures would be addressed at the same time. About 1000 customers might be an appropriate group, for example. A key assumption to point out here is that the group of customers can experience one repair at a time, either due to proactive response or reactive response (depending on the model conditions), so we assume that failures and plant problems are rare enough that they get repaired before a second one happens. This seems reasonable for most cases, as it is for most maintenance models.

The model can be used to describe all failures for the section of plant, or a subset of the failure types. For example, only a set of failures that can be captured with a proactive system might be modeled for easy

comparison. Likewise, a model of all possible failure modes would be created so that a degradation over time could be studied.

Because the repair model is small, and we expect to examine it over longer periods of time, we are mostly interested in the steady-state behavior of the repair system; we will often solve for the steady-state distribution of the 11-state repair model.

## 2. More About the Degradation Model

For the degradation model, if we assume there is an upper limit to the degradation that a reasonable network can sustain, finite state birth-death processes will work for the larger degradation model.

We can connect the repair model by determining the state probabilities and rates of not correcting the causes of the failures. Impairments represent degradation and can arrive according to an independent Markov process. Degradation should be slow, and not likely to reach a steady-state condition, so we need to study the transitive states of the model, not the steady-state. We will need to use a method called uniformization to solve this degradation model for various times over the lifetime of the cable plant. Because the model is a birth-death model, it solves sufficiently fast for reasonable time horizons.

A key assumption to point out here is that the group of customers can experience one repair at a time, either due to proactive response or reactive response (depending on the model conditions), so we assume that failures and plant problems are rare enough that they get repaired before a second one happens. This seems reasonable for most cases, as it is for most maintenance models.

## 3. Parameters

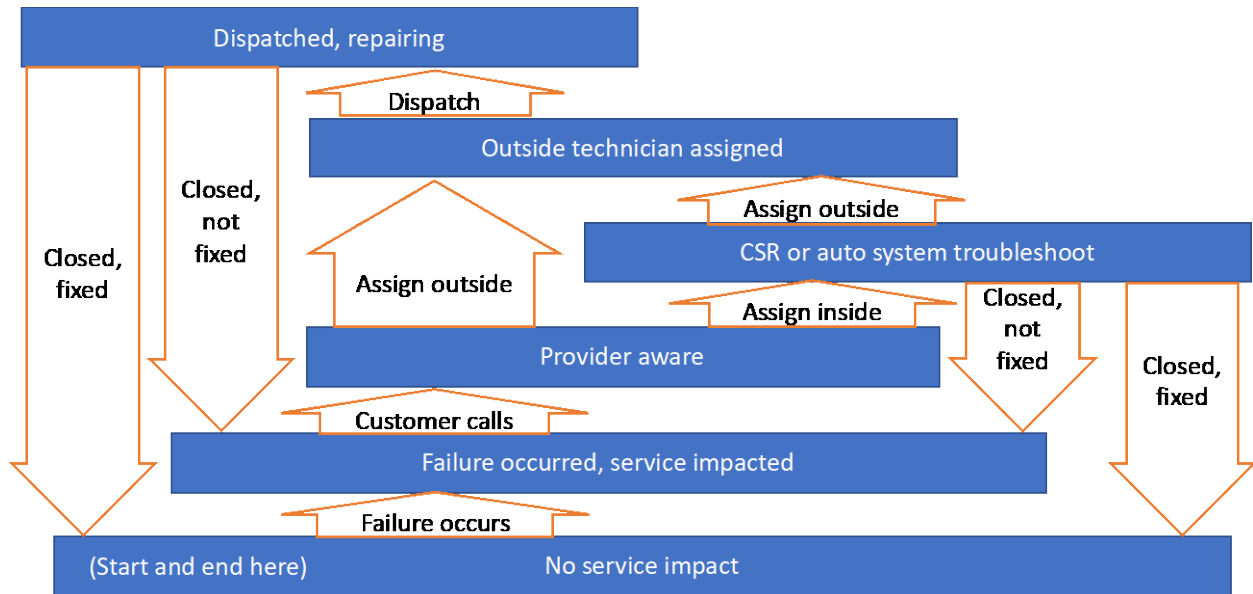
In the models described here, the following notation is necessary.

### 3.1. Definitions

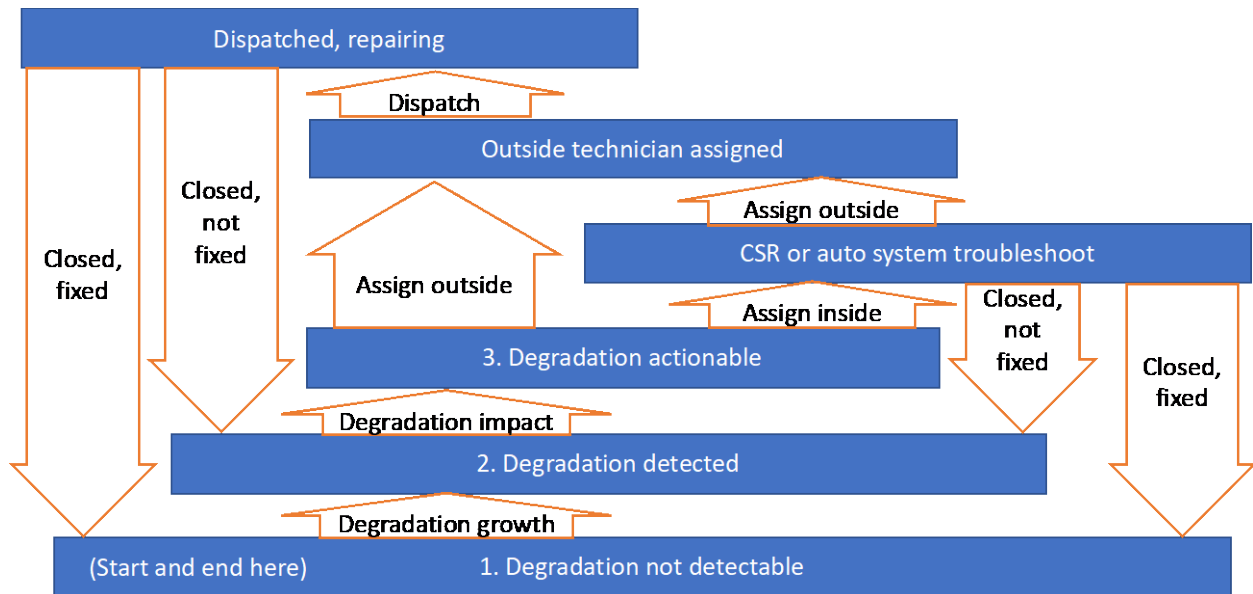
- **Impairment:** This is a problem in the plant that will, after time and degradation, result in a failure that impacts one or more customers. These include loose connectors, nicks in the cable sheath, or a cable that swings in the wind, all of which eventually will result in a problem that worsens and needs to be corrected. Ideally, when corrected, the plant will be in better condition; but often times the plant gets minimally repaired, and so the impairment can result in another failure later.
- **Shock:** This is a type of event that results in immediate failure, with no proactive opportunity. Causes include storms, a shovel cutting through a drop, lightning, sabotage, or accidental damage. In these cases, the plant is damaged, and the damage immediately results in one or more customers having their service adversely impacted, so a repair is necessary.
- **Remote repair:** Without dispatching a technician, the problem impacting service is resolved. Perhaps this is done by a setting change in the system, or instructions with a customer. A remote repair attempt could be made but result in the need for a dispatch; this would be an unsuccessful remote repair attempt but is the same as a triage.
- **Dispatch repair:** This is a truck roll to send a technician out to locate and troubleshoot a problem to repair it. The repair may result in a minimal repair, where service is restored, but the plant is not improved. Sometimes this minimal repair actually fixes nothing, but removes the symptom that the customer experiences, such as placing a filter on the line to block certain frequencies. Or the repair may result in the cause of the problem being removed and the plant condition improved as a result. Improving the condition of the plant reduces the number of impairments in the plant, and thus the net rate of failures over time.



In the subsections that follow, we explain the model details further in terms of transition rates, states, and additional parameters. Figures 5 and 6 show verbose states and transitions for reactive and proactive repair process, respectively.



**Figure 5 – A verbose depiction of the reactive lobe of the repair model, with the common state at the bottom, and reactive repair states above.**



**Figure 6 – A verbose depiction of the proactive lobe of the repair model, with the common state at the bottom, and proactive repair states above.**

### 3.2. Transition Rates

Each of these parameters is described as a rate (events per unit time) because the model requires rates for input. But we can estimate the expected time each event type takes and estimate the rate as the reciprocal of the expected time.

- $Y_{imp}$  = impairment arrival rate. This rate is not a part of the 11-state repair model but is used in the degradation model. The cable plant is exposed to the elements, and subject to damage. When damaged or degraded, there may not be immediate impact to service. But after time, the degradation becomes severe enough to lead to service impact. This measurement is the number of plant impairments (not service impairments) that arrive to a unit of plant in a unit of time. For example, for a section of plant, say a neighborhood, or all customers served downstream from a given amplifier, impairments might arrive due to damage or storms or erosion in the elements at a rate of one per month. While we know this is nearly impossible to tell, we can get at it another way as described next.

Instead of the impairment arrival rate and impairment failure rate, we ask for the failure rate of a unit of plant, and the breakdown of the types of causes to the best of our knowledge. If we can determine the failure rate for a unit of plant (number of failures per unit of time) over a reasonably short period of time (say a year or less), and we can further determine the number of these failures that occur due to shocks versus degradations (by inspection and reporting from the technicians or center persons who work and close the ticket, or from the disposition codes or closure codes or comments), then we can approximate the factors needed for the model.

- $Y_{bs}$  = base failure rate for shocks. The cable plant is subject to shocks that lead to failures. A cable modem (CM) can fail, an amplifier can fail, a cable can be cut, and so forth. Such plant failures are immediate, and different from degradations which can be detected proactively. The base failure rate from shocks can't be reduced without impacting external causes or changing the underlying technology such that you change the failure modes and failure rates of the external (extrinsic) shocks.
- $Y_{if}$  = impairment failure rate. Plant impairments eventually result in service failures that impact customers. It is difficult to know in many cases how long an impairment has existed in the cable plant before it results in a failure. Therefore, we ask for other information that allows us to model this rate instead.
- $Y_{cc}$  = customer call rate. This is the rate at which a customer, whose service is impaired, will call in the problem. The rate is events per unit time, so usually this rate is the reciprocal of the average time between when service is impaired to when they notice that, and then call in for service support.
- $Y_{rr}$  = response to reactive call rate. This is the rate at which, once a customer calls in for help, the operator responds for action, either through an immediate dispatch, or by addressing the problem in a call or service center. The average time from when a customer calls to when the operator acts to resolve the problem is the reciprocal, and sufficient for this parameter.
- $M_l$  = dispatch rate. Once a problem is decided to need action, if not resolved in a call or service center, and it is decided a dispatch or truck roll is required, this is the rate at which the dispatch is done. Sufficiently, we can use the time between when a dispatch is decided to when the dispatch is initiated.
- $M_d$  = dispatch repair rate. Once dispatched, the average time to drive to the area, locate the trouble, and return the services to fully functional is sufficient for this rate parameter.
- $M_r$  = remote repair rate. Once a customer calls in a problem and is connected to a call or service center, if the center resolves the problem, this is the rate at which they resolve the problem. Again, the average time to resolve a problem from the service center, without a dispatch, is

sufficient for this parameter. This estimate can include the center handling time for dispatched repairs too. Queue time should be included in any case.

- $Y_{ir}$  = impairment reveal rate. Once an impairment is present in the network, it takes time before it impacts the radio frequency (RF) signal or other behavior of service, but well before the customer will notice it. This is where the sensitivity of your proactive monitoring and operations shows its impact. The average time between when an impairment exists in the network to when it can be detected using whatever methods you have in place for proactive detection will suffice for this parameter.
- $Y_{pc}$  = proactive alarm rate. Once an impairment is detected by a proactive method, it will exist in the network and impact service until it rises to the level requiring action. This is where the operations method, culture, and practices are revealed best. Provide the average time from when a proactive issue is detected to when it gets bad enough to be addressed by your program, and the rate will be the reciprocal in most cases.
- $Y_{pr}$  = response to proactive rate. Once action is warranted on a proactive issue, this is the rate at which action is taken. Provide the average time it takes from when a proactive opportunity is sufficient to warrant action and when appropriate action is taken. The rate will be considered to be the reciprocal of this average.

A note about  $Y_{ir}$  versus  $Y_{pr}$ : Look at the number of times a degradation that could be detected from a proactive program are reacted to versus proacted against. The proportion of the time that a degradation problem is proacted against should be close to  $Y_{ir}/(Y_{ir} + Y_{pr})$  if the time to address known proactive problems is roughly the same as the time to address reactive ones. Use this approximation for determining a target or for creating targets for future PNM programs or projects, or for estimating whether there is sufficient opportunity for a particular PNM solution in an area.

### 3.3. States

For the 11-state model that models the repair state of a section of plant, consider the following state definitions.

- $w_n$  = working state at degradation level  $n$ . This is when there are no failures in the scope of the network being modeled, so service is working.
- $Ip_n$  = indication of a proactive action. This is the state of the system in which an impairment has become severe enough that it is detected by your proactive system. As the model (in this formulation) assumes that the next action is to address the proactive issue, precluding a reactive one from happening meanwhile, the time to leave this state should be small, so it is likely best to assume the detection is severe enough to warrant action, too.
- $Cp_n$  = customer impacted or triggering a threshold for proactive action. This is the state of the system in which the impairment that was detected proactively has now gotten severe enough that it will definitely be addressed now, either because a threshold was triggered, or a customer was impacted and indicated so.
- $Rp_n$  = remote proactive fix attempt. This is the state in which a service center is working a proactive ticket.
- $Dp_n$  = dispatched proactive fix attempt. This is the state in which the proactive ticket is generating a technician dispatch.
- $Tp_n$  = technician working the proactive ticket. This is the state in which a dispatched technician is working the proactive ticket.
- $Fr_n$  = failure, reactive state. This is the state under which a reactive service failure has occurred for the system being modeled.
- $Cr_n$  = customer called in on a reactive failure. This is the state under which the customer has called in a problem, or an alarm has been indicated that requires repair.

- $Rr_n$  = remote reactive fix attempt. This is the state of the system in which a call or service center is attempting to remotely fix the problem.
- $Dr_n$  = dispatched reactive fix attempt. This is the state of the system in which a technician has been dispatched to fix the failure and service disruption.
- $Tr_n$  = technician working the reactive ticket. This is the state under which the technician is actively working on the trouble ticket to fix the problem, after being dispatched.

### 3.4. Constants

Several constants need to be defined as well, because after a time in a state in the model, there may be choices for the next state, and a constant proportion of the exits from a state may enter a given state next, which must therefore be defined. For each proportion defined, there is a complement (one minus the proportion) that describes the probability that the next state is some complementary state.

- $a$  = proportion of reactive tickets that are truck rolls or dispatched.
- $b$  = proportion of remote repairs that close the ticket, not passing it off to a dispatch.
- $c$  = proportion of repairs that don't resolve an impairment but close the ticket and complete a service repair (minimal repair, doesn't improve the plant).
- $d$  = proportion of dispatch repairs that have to try again, either becoming a repeat, or perhaps don't have access to the location of the trouble.

Note we can add a sub-notation of  $p$  for proactive or  $r$  for reactive to each of these constants to have different proportions for reactive and proactive work. In cases where either  $p$  or  $r$  exist in the notation but not both, the assumption is that the notation without the qualifying subscript is represented by the factors missing the subscript, to simplify the equation noise.

There are additional constants that are not proportions but are important to the model.

- $e$  = constant factor for the variable part of the degradation rate, a function of the state. This constant factor is not a proportion, and applies to the degradation model.
- $n$  = working state number most closely linked, i.e.,  $w_1$ ,  $w_2$ , etc. This is the level of degradation in the plant:  $n=0$  when the plant is brand new, pristine, and without impairments; as it ages,  $n$  increases to a high number representing a highly damaged and degraded plant which generates lots of impairments that need to be fixed.

## 4. Finding Parameters to Feed the Models

To obtain the parameters for the model, it is anticipated that operations measures can be translated to the model parameters suitably through mathematical transformation: Rates can be determined from frequency averages, and proportions can estimate probabilities. As such, these expected operations measurements, which should be available, can provide the basic information for the model. Additional parameters can be set as goals for programs to be evaluated or searched for tipping points.

Several operations measurements that are relevant to the model inputs include the following, though some cable operators may have different sources for estimates than those assumed here.

- Number of trouble calls per customer over a period of time.
- Number of system alarms per customer over a period of time.
- Number of proactive service impairments discovered per customer over a period of time.
- Number of other categories of sources of maintenance work per customer over a period of time, if any, and what makes them distinct from the categories above. We expect this is zero but need confirmation. If it is non-zero, we will need additional information about this category of maintenance work in parallel to the other categories, as below.

- For each of these above sources of work, the proportion of the work done to correct a shock failure versus a failure due to degradation.
- For all the sources of work except proactive, the proportion of these that were indicated by proactive systems but not addressed on time.

Several measures are for reactive work specifically.

- The average or expected time between when service is impacted to when a customer calls in to request help.
- The average or expected time between when service fails to when the alarm indicates there is a problem.
- The average or expected time between when the customer calls in to when someone in a call center or service center responds (include dispatch if it is possible to go right to dispatch or truck roll without touching the center).
- The average or expected time between when the alarm sounds to when someone in a call center or service center responds (include dispatch if it is possible to go right to dispatch or truck roll without touching the center).
- The average or expected time between when a call center or service center responds to a customer call to when it is resolved or decided to dispatch.
- The average or expected time between when a call center or service center responds to an alarm to when it is resolved or decided to dispatch.
- The average or expected time between when a customer call is dispatched to when repair begins.
- The average or expected time between when an alarm is dispatched to when repair begins.
- The average or expected time between when repair begins for a customer call event and when service is restored.
- The average or expected time between when repair begins for an alarm event and when service is restored.
- The proportion of customer calls that go to a call or service center versus right to dispatch (if any).
- The proportion of calls to a service center or call center that are then dispatched.
- The proportion of calls to a service center that are resolved in the center and do not repeat.
- The proportion of calls to a service center that are resolved in the center but end up generating a repeat problem.
- The proportion of dispatched repairs that generate a repeat problem.
- The proportion of dispatched repairs that do not generate a repeat problem.
- For reactive work that is dispatched and does not generate a repeat problem, the proportion of repairs that correct the source of the problem versus minimally repair the cable plant just enough to get service restored.

Likewise, some of these equivalent parameters are needed for proactive work too.

While the above parameters are likely available from operations management systems, we need to translate them into parameters for the model still. For averages, we convert those to rates. If there is information about the variance of these averages, then we can use that as further verification as to the reasonableness of the assumption of exponential distribution for the durations, or make adjustments accordingly. The proportions are directly usable in the model. However, some of the first few parameters on this list need translation as follows.

The first group of six estimates will tell us the parameters  $Y_{bs}$ ,  $n*Y_{if}$ , and  $n*Y_{ir}$ . The fifth estimate(s) gives us a breakdown of the first four so that we can estimate  $Y_{bs}$ . Subtracting  $Y_{bs}$  from the sum of the rates for issues to address, we find the sum of the other two rates. Then the sixth estimate helps us estimate how

much of this sum is attributed to  $Y_{if}$  and how much to  $Y_{ir}$ . The fact that these are averages only tells us an average value for the parameters multiplied by  $n$ . So, we rely on the estimate of the proportion of work that minimally repairs the cable plant versus improves it as a way to estimate the impact on plant impairments. That helps us estimate the value of  $n=1$ . From the proportion of work that is based on shocks versus degradations, we can estimate the number of impairments in the plant utilizing methods from software reliability. Then, we can estimate the  $Y_{if}$  and  $Y_{ir}$  from our estimate of  $n$  and the sum of the rates we estimated above.

The age of the plant being modeled is somewhat important as we estimate the time to an impairment entering the plant and being discoverable. The distinction of those two steps is not likely important to separate, but we do want reasonable estimates for our model.

## Model Use and Use Cases

Many use cases can be addressed with just a steady-state analysis of the 11-state repair model by comparing two or more sets of parameters. Very quickly, the future model state no longer depends on the starting condition, so a steady-state result is sufficient, and easy to solve using traditional methods. The overall service or plant availability can be determined by examining the steady-state probabilities of the model. A current state and future state model pair can be compared, and more complicated combinations can be calculated too.

Decisions around changes to operations are usually assessed by cost-benefit analysis, so a cost model is very important. Costs often have a fixed component, and a variable component that is often a function of time. So, we need to know the time spent in each model state, and the number of visits over time to each state, as well as the costs involved.

To keep the analysis simple, we will consider the average costs of visiting each state, so that the cost includes the fixed and variable costs together.

In the models we encoded and ran, we broke out the costs to the visits to each of the states so that the cost of repeating cycles can be handled in detail. The basic set of cost parameters we use are as follows.

- 10 = cost of goodwill from a reactive failure
- 20 = cost of goodwill when a customer calls to complain
- 30 = cost of a remote fix
- 50 = cost of a truck roll fix
- 50 = cost of the actual repair, separate from truck roll
- 10 = cost of a proactive failure, likely system cost
- 5 = cost of a proactive alarm, likely system cost
- 20 = cost of a remote fix to a proactive problem
- 30 = cost of a dispatched technician fix to a proactive problem
- 50 = cost of the actual proactive repair, separate from the truck roll
- 0 = cost of being in a fully working state, likely 0

Note we could add a negative cost to the fully working state if we wanted to reflect a benefit from being in that state.

For the models we run later, we asked for input from operators, and a few shared some numbers we were able to use. We selected baseline transition rates and constants to reflect what we found from the operators who responded.

The state transition rates we used are as follows.

- $Y_{imp} = 20/(365*24)$  = impairment arrival rate (ignored in 11-state)
- $Y_{bs} = 200/(365*24)$  = base failure rate for shocks
- $Y_{if} = 50/(365*24)$  = impairment failure rate, multiply by n for degradation state
- $Y_{cc} = 4/24$  = customer call rate
- $Y_{rr} = 4$  = response to reactive call rate
- $M_t = 1$  = dispatch rate
- $M_d = 1/2$  = dispatch repair rate
- $M_r = 4$  = remote repair rate
- $Y_{ir} = 2/365$  = impairment reveal rate, multiply by n for degradation state
- $Y_{pc} = 1/24$  = proactive alarm rate
- $Y_{pr} = 4$  = response to proactive rate

The baseline constants we use are as follows.

- $a = 0.0$  = proportion of reactive tickets that are immediate truck rolls
- $b = 0.75$  = proportion of remote repairs that close the ticket
- $c = 0.25$  = proportion of dispatch repairs that don't resolve an impairment but close the ticket
- $d = 0.17$  = proportion of dispatch repairs that have to try again
- $e = 0.1$  = acceleration factor for impairment arrival rate for third order effect
- $ap = 0.0$  = proportion of proactive tickets that are immediate truck rolls
- $bp = 0.75$  = proportion of proactive remote repairs that close the ticket
- $cp = 0.25$  = proportion of proactive dispatch repairs that don't resolve an impairment but close the ticket
- $dp = 0.17$  = proportion of proactive dispatch repairs that have to try again

## 1. Aligning the Repair Model to the Degradation Model

A simple approach for using these models for a long-term cost analysis is as follows. See the top and bottom halves of the 11-state model in Figure 1. The top and bottom lobes (in green and brown) each represent a repair cycle. The top lobe (green) is for proactive work, and the bottom (brown) is for reactive.

When starting in the working state, the probability of entering the proactive lobe next (before entering the reactive lobe) is  $(n * Y_{ir}) / (n * Y_{ir} + Y_{bs} + n * Y_{if})$  so the complement probability, of entering the reactive node next, is  $(Y_{bs} + n * Y_{if}) / (n * Y_{ir} + Y_{bs} + n * Y_{if})$ . Each lobe though can cycle within, and follow different paths in each lobe, so obtaining the costs requires more work.

For the 11-state repair model, we can calculate the expected time in each state from the transition rates, and solve the model for steady-state probabilities to get the percent of time in each state. By taking the probability of being in a state, dividing by the expected time in that state when visiting it, and multiplying by the cost of being in each state, we can estimate the expected cost of a cycle of the model. Dividing by the expected time for a cycle, we get the cost per unit of time of being in that 11-state model. Because we can do this for all degradation states, and thus the 11-state model that aligns with each degradation state, we can find the cost per unit time as the degradation proceeds.

Let  $Q$  be the transition rate matrix, where  $Q_{i,j}$  = transition rate from state  $i$  to state  $j$  when  $i \neq j$ , and the negative of the row sum of the off-diagonal values otherwise for  $i=j$  so that the row sums to 0. This is the standard transition rate matrix for stochastic models.

$PI_a$  is the steady-state probability results from the  $Q$  matrix, representing the probability of being in each state  $i$  over the long run, meaning regardless of the starting state.

$-Q_{i,i}$  is then the transition rate out of a state  $i$ . Thus  $-1/Q_{i,i}$  is the expected time in a visit to state  $i$ , or the sojourn time in state  $i$ .

Thus, we can get  $1/(-Q_{i,i} * PI_{a,i})$  to be the circle back time, or the time between visits to a state  $i$ . Also, this is 1 over the number of visits per unit of time. Therefore,  $-Q_{i,i} * PI_{a,i} * C_i$  is the cost per visit to a state multiplied by the visits per unit of time, so is the cost per unit of time for each state, which we can sum to get an estimate of the cost per unit of time for the whole model.

This lets us calculate the 11-state model for every level of degradation we want to model, then apply a disjoint model of degradation level to the cost estimates to get an overall cost of the plant. Further, we can separately model degradation to model how the cost changes over time, and then plan for optimal maintenance of this cost profile. Likewise, this allows us to use a cost per unit time as a way of comparing different model settings which may represent two use cases to compare, as we will show later.

## 2. Modeling Plant Degradation as it Relates to Repair

The 11-state repair model can be connected to a degradation model as we describe here.

First, we model the degradation level of the section of plant as a birth-death model, meaning that the state increases or decreases by 1 at each state transition. If we let the first state represent the pristine cable plant, then a birth represents a degradation, and a death represents removing the degradation from the plant to return it back to good as new for the part of the network that was affected.

To use a degradation birth-death process model, we need to determine the degradation rate (birth) and improvement (death) rate based on the state  $n$ . Assume the degradation birth-death model has 100 states numbered  $n=1$  to 100. This assumption is arbitrary and can be easily relaxed; but for our purposes, 100 states should be enough to model reasonable levels of cable plant degradation.

For degradation, we define a constant rate of degradation  $Y_{imp}$  added to a rate that depends on the state  $n$ ,  $e*(n-1)*Y_{imp}$ . Note we use  $n-1$  instead of  $n$  because the first state 1 will then have a degradation equal to the base constant rate. Also, note that this degradation rate applies to states  $n=1$  to 99 as state 100 can't degrade any more.

For the improvement rate, we have to determine the probability of being in a repair state that allows improvement, times the rate at which an improvement happens. There are two such states in the 11-state model. This rate, dependent on the current state  $n$ , is then

$$PI_{a,Tr,n} * (n * Y_{ir} / (n * Y_{ir} + Y_{bs})) * Md_r * (1-d) * (1-c) + PI_{a,tp,n} * Md_p * (1-dp) * (1-cp)$$

Note the improvement rate applies for  $n=2$  to 100. Also note we allow  $M_d$  to be different for proactive and reactive in this equation, so adjusted the notation respectively to  $M_{d,p}$  and  $M_{d,r}$ .



### 3. Reactive and Proactive Relationship with Failure Rate

The impairment failure rate and impairment reveal rate are related in a potentially complicated way, depending mostly on the operations data feeding the model. Some of the impairments will lead to failures, and some will be revealed and acted on before becoming a failure due to a proactive capability or program. This relationship requires a reasonable model to trade off one for another when comparing reactive only to proactive and reactive programs mixed.

First, we assume that the failure rate is a combination of multiple failure causes each occurring with its own rate, and the net failure rate is approximated by the sum of rates. Incidentally, this summation simplicity is one reason we prefer to work with rates rather than time between failures; one is the reciprocal of the other. Due to this assumption, we can discuss one part of this failure rate for a single failure cause. Note that a failure cause in this case can specify a network component type and failure mode, or a group of either, but is specific to the scope of the proactive maintenance that is possible.

Say the rate at which a given problem appears in the network being modeled is  $Y_1$ . For simplicity, this rate can include the  $n$  factor for degradation within it. If the proportion of these problems that get caught proactively is  $p_p$  and the reactive proportion is  $p_r$  such that the two sum to 1, then a simple way to split them is to let the proactive part be  $p_p * Y_1$  and the reactive portion being the complement. As long as the proportion is applied randomly, the results are two thinned Poisson processes which are each Poisson processes. They also can add back together, which makes the modeling simple to manage and adjust.

However, it is far more likely that we have the all-reactive operational data and are looking to examine the impact of changing some reactive work to proactive. Doing this requires we reverse the process we just described earlier by reducing the reactive rates by the amount that converts to proactive (thinned Poisson process), and then pushing the proactive work to the proactive part of the 11-state model. If we believe that the shift of work would change the other reactive rates in the model, then adjustments to those are in order. However, for our scenarios in this paper, and likely for many operations cases, the secondary effects are likely negligible. Therefore, for comparing reactive only to proactive-reactive mix environments, we can simply divert some of the reactive work to the proactive side, and reduce the rates as thinned Poisson processes so that the all reactive case has a rate  $Y_{ir}$  that is equal to the proactive-reactive rates  $Y_{ir} + Y_{ir}$  which is a simple comparison.

We can write this relationship as  $Y_{ir,r} = Y_{ir,p} + Y_{ir,p}$  adding the  $r$  for the reactive-only model parameter, and the  $p$  for proactive-reactive model parameters. We will apply this method in the models in this paper unless otherwise stated.

## Model Study Results

We will now use these models to study several use cases for reasonable parameter estimates obtained from discussions with operators willing to share their knowledge.

- Advantages of PNM – We will run a sensitivity analysis on the conversion of proactive and reactive to reactive only in the 11-state model and determine if PNM is a strong advantage.
- Reducing Repeat Rates – We will start with a simple repeat rate model where we model a high repeat rate and then a much lower rate but with different repair times and costs, all using the 11-state model.
- Degradation Over Time and Optimizing Cable Plant Rehabilitation – We will model degradation over time according to the degradation model outlined before, but with different settings for degradation and removing the degradation causes encountered in repair. The analysis will reveal


the behavior over time of the degradation model, which we can consider against best practices for maintenance.

## 1. Reactive-Only versus Proactive and Reactive Mixed

In our first use case study using the 11-state repair model, we compare the baseline parameters for a mix of proactive and reactive, then adjust the proactive work to be reactive as described earlier but allowing for some portion of the proactive work to turn into reactive work, from 50% to 100% in increments of 5%. This allows for situations where the conditions generating proactive work may not always be conditions that turn into reactive work. As little as 50% is considered, but as much as all of it is in the extreme.


We allowed the baseline rates and constants to vary so that we could generate Monte Carlo runs of different parameters for comparison. Rates were allowed to vary from 1/3 to 3 times the base rates. The constants were varied by 1/3 of the proportion less than 0.5, on either side. For example, if a constant was 0.75, then it was allowed to vary by  $1/3 \cdot (1 - 0.75) = 0.25/3$ , so can be  $0.75 \pm 0.25/3$ . We then generated uniform random numbers in the range that each was allowed to change and created 100 different sets of settings. We then applied each model condition to the 100 setting sets to generate 100 model results for each of the 11 conditions. We then used these results for comparisons as follows.

See Figure 7 for the availability results of the 11 different model results. We generated box plots of the 100 Monte Carlo availability results for each model condition, assuming the starting state was in degradation state  $n=38$ , as that was a state close to what one operator described as their average condition. See that almost all of the availability results for a proactive and reactive repair system were better than any of the results for the reactive-only results, even when the reactive fixes of proactive problems were half what the proactive work would have been. In other words, from a service availability standpoint, proactive repair is much better in general.




**Figure 7 – Service availability over 100 Monte Carlo runs for a proactive and reactive maintenance system, and reactive-only where the proactive work generates 50% more to 100% more reactive work when not handled proactively.**

Now examine the cost rate results in Figure 8. The data in these box plots are cost rates from the same runs as in Figure 7, so consider a degradation rate at 38. Note that the repair system that has proactive and reactive work has a much lower cost rate as indicated by the box plots, compared to the various reactive-only repair systems. Once again, we have some evidence that in general a mix of proactive and reactive repair will cost less than a reactive-only repair system.



**Figure 8 – Cost rate over 100 Monte Carlo runs for a proactive and reactive maintenance system, and reactive-only where the proactive work generates 50% more to 100% more reactive work when not handled proactively.**

To check whether the degradation level of 38 is somehow special or not, we provide box plots of all degradation levels, comparing the best case reactive-only repair to the mixed proactive and reactive repair solutions, in Figure 9. Because we take every model run and calculate the difference in the cost rates of the two model settings, we can say that any negative values would be due to a reactive only solution being better than the proactive and reactive mixed system. But look at Figure 9 and see that there are no box plots with even extremes in negative cost differences. This indicates that, at least for the settings run in the Monte Carlo model runs, having a PNM element to a repair process is cheaper than not having one.



**Figure 9 – Cost rate difference between the best reactive case to the proactive and reactive case for add degradation states, over 100 Monte Carlo runs.**

## 2. Reducing Repeat Rates


One operator expressed a concern about their high repeat rates, so we decided to demonstrate how to use these models to assess whether it makes sense to let technicians take longer to troubleshoot and repair problems so as to reduce repeat rates, or not. The baseline repeat rate we used was 0.17 (3xR), and we examined competing cases where the repeat rate was 0.17/3, but at a cost of increasing the time and cost of repairs by 10% (LR1) to as much as 110% (LR11). Once again, we generated 100 Monte Carlo runs of the model, letting the rates and constants vary as before, except for the repeat constant.

See Figure 10 which has box plots of the results for the high repeat rate first, then the lower repeat rates with longer repair times as indicated, all for degradation level  $n=38$  as before. As might be expected, the impact to service availability is very minimal, and decreases as repair times take longer. However, a good repair may mitigate this impact by restoring service quickly, then spending more time to find and remove the root cause of the problem.


Now look at Figure 11, which shows the costs per unit of time for the same runs. The high repeat rate has the cost shown in the first box plot, but notice that the first few lower repeat rate cases have lower costs per unit of time, suggesting that spending up to 30% or so extra time (cost) on the repair will still result in lower cost overall. More time spent in the field correcting problems will pay off in the long run.

To show that the degradation state of  $n=38$  is not so special, we calculated the difference in cost per unit time between the high repeat rate case and the low repeat rate but spending an additional 110% of time


troubleshooting. While the latter mostly has a much higher cost per unit of time, it isn't always the case. There may be times where it is still cheaper to spend twice the time to find and fix troubles.



**Figure 10 – Service availability over 100 Monte Carlo runs for a proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost to 110% more time and cost.**




**Figure 11 – Cost rate over 100 Monte Carlo runs for a proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost to 110% more time and cost.**



**Figure 12 – Cost rate difference between the proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 110% more time and cost, over 100 Monte Carlo runs.**

Last, we took the cost difference between the high repeat rate case and the low repeat rate case that requires spending only 10% more time and cost troubleshooting. That result is in Figure 13. Notice that the high repeat rate case is always more expensive, and in some cases significantly so.



**Figure 13 – Cost rate difference between the proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost, over 100 Monte Carlo runs.**


### 3. Degradation Over Time and Optimizing Cable Plant Rehabilitation

We can model the probability distribution of a future state or estimate the past costs from an assumed start condition and age and parameters of current cable plant to estimate the lifetime cost functions for a section of plant. There are multiple uses for these results, but a major one is to use these results to set the optimal replacement times for the cable plant [2]. It is known that for systems that degrade, when the total cost of the system divided by its age (so the cost per unit of time of the system) is lower than (overtaken by) the cost per unit of time of the repair, then the overall cost per unit of time of the system begins to increase, so it is the optimal time to replace or fully repair it.


We run the 11-state repair model over all 100 degradation levels, then plot the availability and cost rates for each in Figure 14, and Figure 15 respectively. Note as should be expected that the availability decreases with degradation, and cost increases.

But as we start in a specific degradation state but let the birth-death model transition over time, we get probability distributions for the degradation state in the future. We show the cost per unit time, then the probability distribution functions for the degradation levels, starting in degradation state  $n=38$ , for all four combinations of settings of 20 or 10 degradations on average per year, and 25% or 12.5% of repairs not fixing the root cause of the problem (not to be confused with repeat rate, but correcting a degradation 75% or 87.5% of the time), at an age of 60 months, in Figure 16 through Figure 23.







**Figure 14 – Service availability over the degradation states, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.**




**Figure 15 – Maintenance cost rate over the degradation states, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.**




**Figure 16 – Cost over time per month starting in degradation state 38, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.**




**Figure 17 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.**




**Figure 18 – Cost over time per month starting in degradation state 38, with constant degradation of 10 per year, and 25% of repairs not fixing the problem.**




**Figure 19 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 10 per year, and 25% of repairs not fixing the problem.**




**Figure 20 – Cost over time per month starting in degradation state 38, with constant degradation of 20 per year, and 12.5% of repairs not fixing the problem.**



**Figure 21 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 20 per year, and 12.5% of repairs not fixing the problem.**



**Figure 22 – Cost over time per month starting in degradation state 38, with constant degradation of 10 per year, and 12.5% of repairs not fixing the problem.**




**Figure 23 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 10 per year, and 12.5% of repairs not fixing the problem.**


Examining the probability distribution functions, we see a system that gravitates toward a low degradation level or a high degradation level, depending on the parameters. The system does not stabilize to a steady-state distribution of degradation states, but rather is likely to move toward the extremes. But there are cases where it could be at either extreme under the same settings. See Figure 21 for an example of such settings, where a low degradation rate is highly likely, as is a high degradation rate, but not so for intermediate states. Also, the costs per unit of time seem to stabilize. But this is likely due to the system stabilizing at an extreme. If that extreme is a high degradation level, then it could be that the degradation is stopped only by the finite states of the model, and that a real system would continue to degrade to alarming levels.

However, this movement toward a high degradation level can be avoided by PNM practices. Another well-known repair optimization result for systems like this is that it is almost always better to do condition-based maintenance over time-based maintenance. If we can examine that the rate of trouble tickets is increasing, and many of them are from degradation causes, then it might be time for a plant sweep or a plant rehabilitation (even if not proven to be the optimal time for it).

See Figure 24, and Figure 25 for the probability distribution function of states at 12 months, starting in degradation state 1, and state 20, respectively. These plots are both under the worst case of high degradation rate and lowest chance of correcting the degradations at repair opportunities. Yet after 12 months the systems are still likely in low degradation levels, because they started in low degradation levels. This suggests that removing degradations from the plant, and keeping them low, can keep costs low. An annual sweep might be sufficient, but monitoring for conditions and triggering a cleanup effort when degradations are appearing might work even better.



**Figure 24 – Probability distribution function at 12 months, starting in state 1, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.**



**Figure 25 – Probability distribution function at 12 months, starting in state 20, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.**

## Conclusions

While the models presented in this paper are general examples intended to be tuned to specific operator needs, and the analyses are examples to consider for applications, the results are still of use in a general sense.

### 1. Generalizing Results

These general models suggest that proactive maintenance can reduce maintenance cost overall and keep service availability high. Further, it can be cost effective to spend more time and technician cost correcting problems to avoid repeat troubles. Also, there is evidence that cable system repair can shift from maintaining cable plant at high quality to having accelerating degradation, but that monitoring or annual sweeps can keep operations costs lower.

### 2. Operator Uses

In this paper, we demonstrated several use cases for these models, but we expect operators will find new ones. We look forward to working with operators to try specific analyses to support their operations improvements.

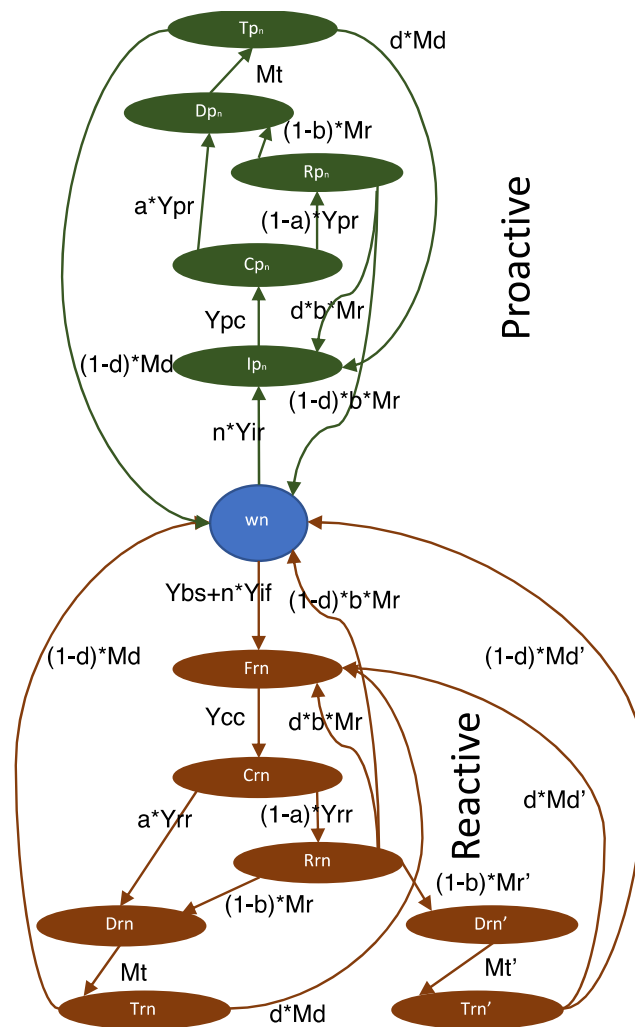
### 3. Enhancements

The models presented in this paper can be extended for other operations assumptions simply, and the methods provided will work for them as well. Some envisioned enhancements are as follows.

- Add the ability to fail while in the first proactive state.
- Decouple the proactive and reactive models so they are separate, can overlap, and can be solved generally; then combine them with the degradation model.
- Add the ability to model two types of repairs after center triage, so that two different repair response processes can be modeled accurately; this adjustment is depicted in Figure 26.

These models were built without working with specific operators' parameters, so we expect them to be a starting point only, with changes needed to better model specific operator situations. Certainly, the parameters may need to be updated, but there may be operations processes used by some operators that are significantly different and need further changes to properly model the situation.

Nonetheless, the results presented may apply and inform many existing operators already, making it easier to justify trying proactive methods in their operations with greater confidence.



**Figure 26 – An adjustment to the 11 state repair model to allow for two types of reactive repair handling, as recommended by an operator.**



## Abbreviations

CableLabs	Cable Television Laboratories
CM	cable modem
DOCSIS	Data-Over-Cable Service Interface Specifications
ISBE	International Society of Broadband Experts
PNM	proactive network maintenance
RF	radio frequency
SCTE	Society of Cable Telecommunications Engineers

## Bibliography & References

- [1] Taylor and Karlan “An Introduction to Stochastic Modeling,” Academic Press, 1984, ISBN-10: 0126848807.
- [2] Jason Rupe, “Optimal Maintenance Modeling on Finite Time with Technology Replacement and Changing Repair Costs,” Annual Reliability and Maintainability Symposium. 2000 Proceedings. International Symposium on Product Quality and Integrity (Cat. No.00CH37055)

# **Rethinking Customer Support**

## **Proactive Customer Engagement: Experimentation in Real-Time Data**

A Technical Paper prepared for SCTE•ISBE by

**Joe Keller**

Executive Director Analytics  
Cox Communications  
6305 Peachtree Dunwoody Road, Atlanta, GA 30328  
404-269-5938  
Joseph.Keller@cox.com

**Sam Plant**

Analytics Engineer  
Cox Communications  
6305 Peachtree Dunwoody Road, Atlanta, GA 30328  
404-269-1743  
Sam.Plant2@cox.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
1. What is Proactive Customer Engagement? .....	3
2. Going Proactive .....	4
3. Improved Outage Module Detection and Staging.....	5
3.1. The Opportunity Factory .....	5
3.1.1. Hypothesis Testing.....	6
3.2. The Experimentation Factory .....	6
3.3. Pilot Architecture .....	7
3.4. The Model Process.....	8
3.5. Tuning Model Triggers.....	9
3.6. Production Shadow Results .....	11
4. What's Next .....	11
4.1. Post-Outage Offline Device Resolution .....	12
Conclusion .....	12
Abbreviations.....	13

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Cross Industry Standard Practice for Data Management .....	4
Figure 2 - CMTS Trap Example .....	6
Figure 3 - Pilot Design Process.....	7
Figure 4 - Pilot Architecture .....	8
Figure 5 - Main Loop Process .....	9
Figure 6 - Upper Control Limit.....	10
Figure 7 - Second Potential Limit .....	10
Figure 8 - Basement .....	11

# Introduction

Multiple Service Operators (MSOs) traditionally interfaced with customers reactively. When experiencing service degradation, seeking education, or requesting new services, the customer's only real recourse is to contact their service provider and attempt to describe their needs to agents who often struggle to provide resolutions sourced from a multitude of applications, datasets, and data sources. In the past, attempts to proactively interface with customers were stymied by a comprehensive lack of data understanding, by low data velocity, and by the cost-prohibitive nature of the operational and technological capabilities required to identify issues before they impact customer.

...But the industry and technology have changed.

Proactive Customer Engagement (PCE) represents a cultural shift in how Cox Communications interacts with our customers. Leveraging probabilistic models, higher velocity data, and cloud-based technologies, Cox Communications seeks to shift customer interactions from a reactive to a proactive stance.

## 1. What is Proactive Customer Engagement?

Proactive Customer Engagement (PCE) seeks to address our industry's emerging number one challenge: improving customer experience while simultaneously driving down operational costs. Cox Analytics posited a solution: Predict and address customer needs **prior to** a customer contact, thereby saving customers the arduous task of explaining their issues to an agent, and simultaneously reducing the cost of serving these needs through traditional high-cost channels.

There are three engagement strategies for PCE:

- **Deflection:** Predicting a customer's intent and correctly addressing that intent upon receipt of a customer contact (e.g. improved outage detection, identifying remote pairing issues based on set top box errors, identifying the wrong HDMI input based on tuning and set top box error data, etc.).

This initial phase of PCE serves as the cross-over point from reactive management of customer needs to proactive management by leveraging real-time information while interacting with customers via more traditional means.

- **No Outbound Contact:** Predicting a service degradation or out of service scenario that can be fixed without customer contact (e.g. device reboot, device re-authorization, device re-provision, firmware push, etc.).

The second phase of PCE corrects issues that might otherwise manifest in a customer contact without the customer's knowledge. Fixes are only applied when the customer is not actively using their services.

- **Outbound Contact:** Predicting a service degradation or out of service scenario issue that cannot be fixed without customer contact (e.g. home technician visit, network technician visit, device swap, pairing remote, etc.).

This final phase of PCE combines customer behavioral data and customer service usage data to interact with customers according to their channel of preference (e.g. SMS, email, outbound dialer).

When engaging customers proactively, thorough consideration must be made for establishing the correct balance between ‘being caring’ and ‘being creepy’. Initial experiments are designed for deflection scenarios with no outbound contact while ongoing analysis focuses on classifying the customer by preferred channel of interaction.


## 2. Going Proactive

The cable industry continues to strengthen its ability to react to customer needs by leveraging the once latent data produced by the network itself. Customer Premise Equipment (CPE) telemetry, network telemetry, and guided trouble-shooting platforms arm agents with the ability to collect data, trouble-shoot, and identify the best fix agent to address a customer’s need.

To realize the vision of going proactive, PCE requires streaming/real-time data assets and advanced analytics capabilities to predict and address customer needs before the customer initiates contact. A non-trivial task.

To this end, it is necessary to determine the feasibility of proactive engagement. This is done leveraging a repeatable methodology: **The Analytics Lifecycle**.

Cox Analytics team employs the Cross-Industry Standard Process for Data Management (CRISP-DM) to manage PCE use cases through the **Analytics Lifecycle**. Hypotheses are developed and prioritized via the **Opportunity Factory**, promoted to pilot via the **Experimentation Factory**, and finally operationalized at scale via the **Deployment Factory**.



**Figure 1 - Cross Industry Standard Practice for Data Management**

### Opportunity Factory

- **Business Understanding:** The process of developing an understanding of project objectives and requirements, translating this understanding into required data sets, and the performing preliminary data discovery and preparation.

- **Data Understanding:** Performing Initial data collection, insight development, and hypothesis development. (What data informs the objectives and outcome of each use case scenario? What is the quality, velocity, and availability of that data set?)
- **Data Preparation:** Definition of the logical data model inclusive of table, record, and attribute source selection and transformation (if applicable) and aggregation for consumption by operational models.
- **Modeling:** Design, build, and iterate on models while focusing on model recall, precision, and confidence. This includes shadow model training and the introduction/removal of features based on their ability to improve on model effectiveness.

### **Experimentation Factory**

- **Evaluation:** The execution of a model in Pilot with continued iteration over model development and continued evaluation of model effectiveness. This will lead to a go/no go decision for enterprise deployment.

### **Deployment Factory**

- **Deployment:** The enterprise-wide deployment of predictive/probabilistic models beginning with an initial pilot and extending to enterprise availability.
- **Continuous Improvement:** Continuous evaluation of the predictive/probabilistic models for effectiveness which includes reassessment of feature selection, further A/B testing, and channel-specific effectiveness.

The balance of this paper tracks our initial foray into proactive engagement through the Analytics Lifecycle, elaborating upon our initial use case: improved outage module detection and staging.

## **3. Improved Outage Module Detection and Staging**

Project Off Ramp is a code-name for Cox Communication's improved outage detection and call deflection program. The program represents an initial foray into experimentation with proactive engagement, leveraging real-time streaming trap data to detect HFC network outages and stage IVR outage messages.

Historically, outage detection was based upon a combination of 1) the number customer calls received by the IVR within a given time frame for a common node, and 2) polling-based telemetry. These outage module detection algorithms lagged outages by 15 minutes or more. With the introduction of Off Ramp, we can use real-time, streaming trap data to rapidly detect an outage condition, allowing for the deflection of more calls, eliminating unnecessary truck rolls to the customer premise, expediting service restoration, and improving overall customer experience.

### **3.1. The Opportunity Factory**

The first step on the path to proactive engagement begins with the discovery of available datasets that meet the dual requirements of being both available in real-time and containing leading indicators of customer issues. One such data source is real-time streaming CMTS Modem Online/Offline trap data that can be used to identify the status of a DOCSIS device.

We hypothesized that the streaming trap data could be used in combination with IVR messaging and work order blocking to reduce support calls and eliminate unnecessary truck rolls during an outage more quickly than the system that was in place.

```
"Identifier": "172.30.63.99 cdxCmtsCmOnOffNotification d4 04 cd d7 f4 47 0 3",
"Serial": "372440111",
"Node": "DT1XCAPC06",
"NodeAlias": "DT1XCAPC06",
"Manager": "MITRAPD",
"Agent": "CISCO-DOCS-EXT-MIB",
"AlertGroup": "cdxCmtsCmOnOffNotification",
"AlertKey": "d4 04 cd d7 f4 47",
"Severity": 3,
"Summary": "Cable Modem Online(CMStatus: online; DownChannelIfIndex: 68617; UpChannelIfIndex: 396748)",
"FirstOccurrence": "2019-07-05T06:56:52",
"LastOccurrence": "2019-07-05T06:58:36",
"Type": 0,
"Tally": 2,
"ProbeHost": "FED1ISMV01",
"ProbeType": "MITRAPD",
"Hostname": "
"MACAddress":
"IPAddress":
"Condition": "ONLINE",
"NEVendor": "CISCO",
"NEModel": "cBR8",
"NEType": "CCAP",
"SystemName": "SAN DIEGO"
```

**Figure 2 - CMTS Trap Example**

### **3.1.1. Hypothesis Testing**

Can CMTS Online/Offline Traps Outage Detection Enable Us to Outperform Current In-Place Outage Detection?

This hypothesis above was tested using historical call and outage data combined with a POC model utilizing the real-time CMTS trap messages. Calls and truck rolls that occurred between the start of an outage detected using the new model but before the time the same outage was detected using the in-place model were counted toward the model's effectiveness.

This preliminary analysis confirmed the viability of the hypothesis. Even with allocations made for the subset of our customer base passing through the IVR during an outage that still want to speak with an agent, we saw a significant opportunity for additional call and truck deferrals. Table 1 represents annualized reductions in calls and trucks calculated during use case validation.

**Table 1 – Preliminary Business Case**

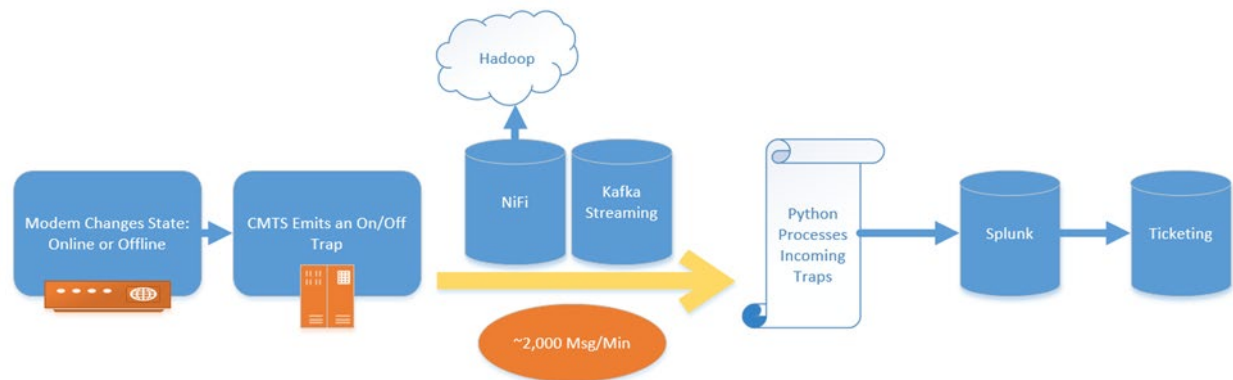
	<b>Transaction Reduction Opportunity</b>
Technical Support Calls	233,169
Scheduled Truck Rolls	38,815

With the business case validated, the use case proceeds onto model development and experimentation.

## **3.2. The Experimentation Factory**

Experimentation begins with defining the conceptual architecture and process for tuning the outage detection model. In this case the streaming data source originates at the CMTS, NetCool collects these

messages and forwards them on to a Kafka topic to minimize the latency of this pipeline. A VM running the outage detection model then consumes the messages.



**Figure 3 - Pilot Design Process**

Within the model, messages are enriched with additional attributes that allow for the identification of a customer's node using the Upstream and Downstream interface card ID associated with the message. A state table is then updated with the online/offline state of the device.

The approach is analogous to methods of counting children in a classroom.

If you have counted the number of children in a classroom once, you can keep an accurate count of the children by increasing or decreasing your count based on how many pass in and out of the room rather than recounting everyone as you would in a polling system.

In much the same way, we populate an initial snapshot of online devices in the device state table and then maintain synchronization with any changes for added or removed devices. As an added precaution, daily polling updates remove devices not present for > 3 days from the device state table.

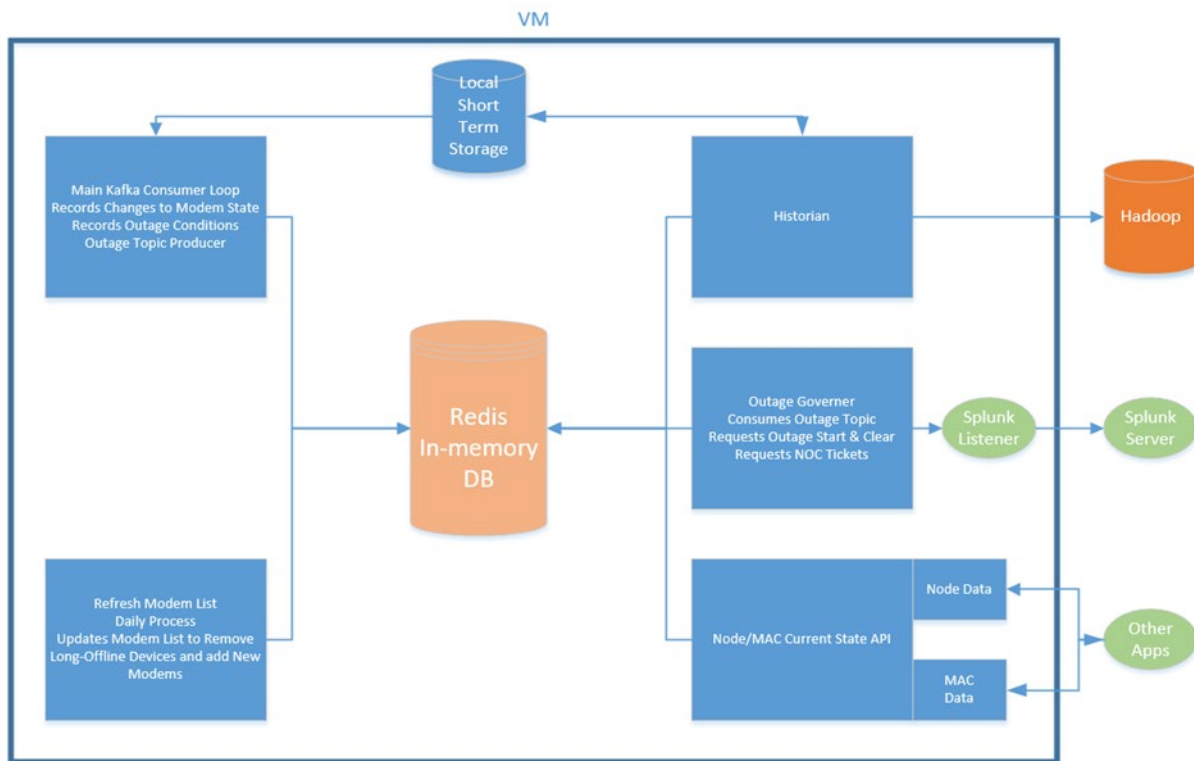
### 3.3. Pilot Architecture

A modular approach was chosen for the CMTS Traps outage detection model. This was done for three main reasons. First, we wanted to be able to expand the capabilities of the model without being forced to change the main model loop. Second, we needed the ability to make changes to data elements while the main outage detection model continued to function. This was specifically to enable maintenance on the device lists and to give access to node and device information in real time through an API. Lastly, we wanted to design a model that could easily be migrated to the cloud.

Our model relies on an in-memory database tool to manage a number of key data objects: modem mac address key value pairs to manage modem level state and location information, node key value pairs to manage aggregated state information for our HFC nodes, and a number of streams to manage communication between the main consumer loop and our governor process. Since we had a variety of data type needs, we selected Redis as our in-memory DB.

We designed this model to be easily migrated to the cloud in the future. The main module is essentially a message consumer. Using AWS Kinesis, ElastiCache, API Gateway, and Lambda, we could replicate the architecture below while increasing the overall scalability of the model.





**Figure 4 - Pilot Architecture**

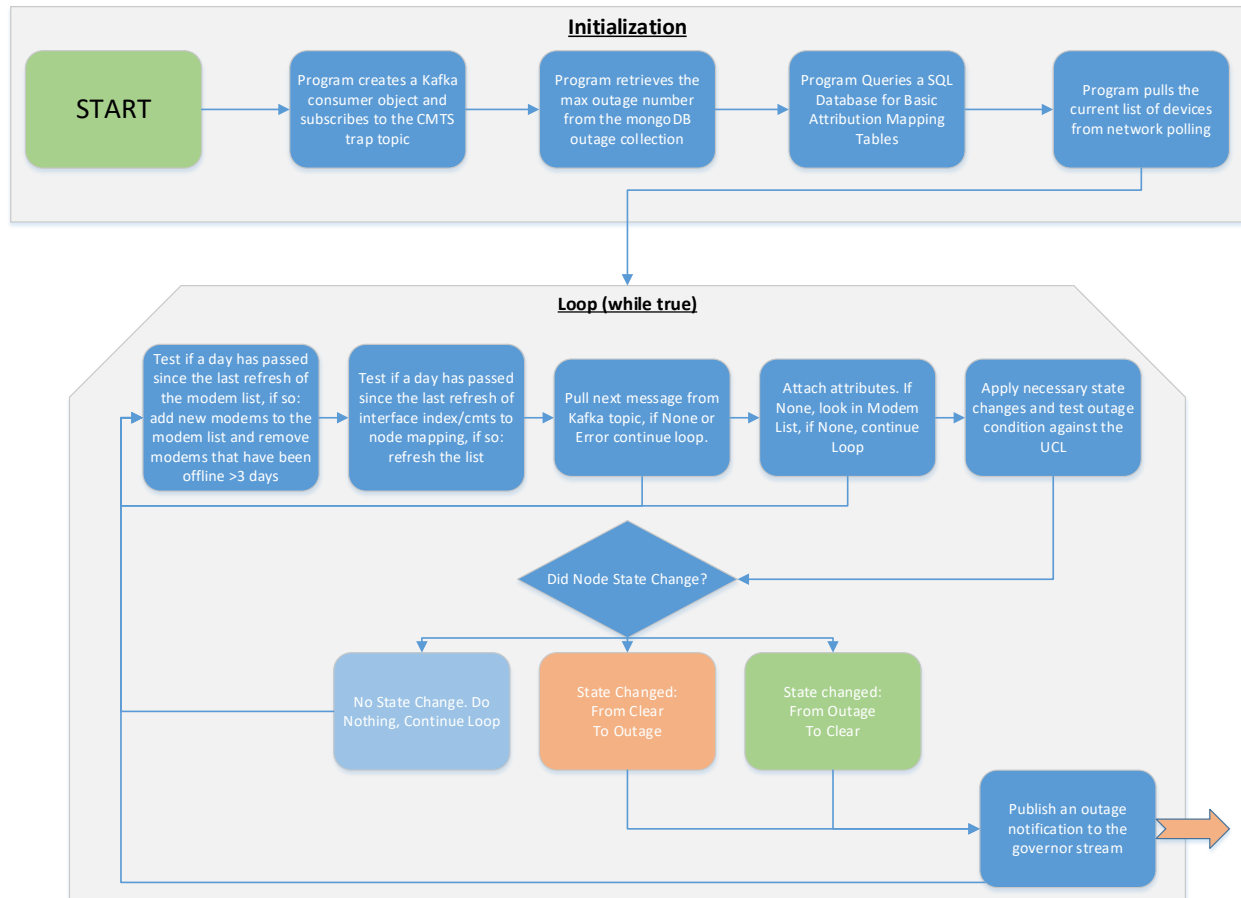
### 3.4. The Model Process

There are 4 processes that comprise the Off-Ramp model:

1. **Main Loop Process:** Core process to the model. The Main Loop process receives streaming trap data, applies attribution for mapping to billing system Site ID and node, updates the device state table, and checks for the Upper Control Limit for offline devices, thereby not taking any action, declaring an outage, or clearing an outage.
2. **Device List Update Process:** Process for maintaining an up-to-date device list. The Device List Update process leverages device polling data, comparing polled device lists to the device state table. Devices that have not been present in polling results for > 3 days are removed from the table. Devices that have not been present are added.
3. **Outage Trigger Governance Process:** Process that consumes outage messages from the Main Loop process. The Outage Trigger Governance process adds each message to a list of pending outage start/clear events per node.
  - If a clear is received after an outage and before 3 minutes, both the outage and the clear are deleted. (No outage is triggered.)
  - If an outage message is received after a clear and before 10 minutes, both the clear and the outage are deleted. (The outage remains active.)
  - If an outage is received and no clear is received before three minutes, the outage is triggered.
4. **Historian:** Leveraged for historical outage profiling and for weekly calibration of the Upper Control Limit logic. The Historian reads the outage stream and copies the data to a long-term

database while removing the messages from the stream. A database is used given that it makes the process of recalculating upper control limits easier.

Figure 6 below provides a detailed explanation of the Main Loop process.



**Figure 5 - Main Loop Process**

### 3.5. Tuning Model Triggers


The most critical step in developing any model is calibrating that model for intended outcome. In the instance of the improved outage staging model, calibration focuses on the triggering logic for staging and clearing outages. Model tuning must account for small increases in offline modems due to customers restarting their devices without triggering an outage.

The upper control limit is the maximum of three potential values used to trigger the staging or clearing of an outage. If the number of offline devices exceeds the UCL for > 3 minutes, then an outage is triggered. If the number of offline devices for an active outage returns below UCL levels for > 10 minutes, then an outage is cleared.

1. The first limit is calculated using Median Absolute Moving Range (MAMR). This method utilizes the median absolute deviation methodology which is much more resilient when your data has outliers. Outlier events in terms of offline devices at an HFC node level represent a significant departure from the normal state of a node. While a healthy node usually only has a


few devices offline at any given time, and changes tend to occur in small single digit increments and decrements, outage events tend to occur suddenly and include dozens of users. This ruled out the use of standard deviation which tends to exaggerate outlier events and could result in an overly high upper control limit (especially if there is a history of whole-node outages).

As indicated in figure below, the center line represents the median number of offline devices by node. We add to this our sigma value multiplied by 3. Sigma is calculated as 1.0483 times the median of the absolute value of the differences (over time series) in offline devices between measurements by distinct node.




**Figure 6 - Upper Control Limit**

2. The second potential limit is calculated as a percent of the total number of devices on a node. As indicated in figure n.n below, if the number of offline devices meets or exceeds 2.5% of population of devices on a node, an outage will be triggered. This was chosen to avoid over sensitivity on lower activity nodes.



**Figure 7 - Second Potential Limit**

3. The ‘basement’ or minimum number of offline devices required to trigger an outage is established leveraging the actual value distribution of offline devices on a node. As indicated in figure n.n. below, 5 offline devices is the minimum number of devices allowed to trigger an outage from any node.



**Figure 8 - Basement**

### 3.6. Production Shadow Results

Production Shadow Models leverage production data to validate the effectiveness of a model without invoking action. In this instance, no outages were declared; however, the times that the model triggered outages are logged and compared against instances of observed production outages. The corresponding volume of calls and trucks that could have been avoided is calculated based upon the outage start time according to the Off-Ramp model vs. the outage declaration time according to the legacy model.

The results represent a significant improvement in outage deflection based upon the improved outage module's detection capabilities leveraging streaming trap data.

**Table 2 – Production Shadow Results**

	Shadow Run	Transaction Reduction Opportunity
Technical Support Calls	1	140,950
Scheduled Truck Rolls		27,314
Technical Support Calls	2	176,428
Scheduled Truck Rolls		32,219

Based upon these results the Off-Ramp outage detection and staging model will graduate to production pilot over the course of Q3 2019.

## 4. What's Next

As mentioned above, the Off-Ramp outage module detection and staging experiment represents an initial foray into proactive experimentation. That said, there are several use cases currently undergoing

feasibility assessment that represent the evolution of our vision for proactive customer engagement. Each use case considered for PCE must meet the minimum qualifications of being enabled via real-time datasets that serve as a leading indicator for customer issues.

The following is an example of a use case that is currently undergoing discovery and validation as next priority for PCE.

#### **4.1. Post-Outage Offline Device Resolution**

As a fast-follower to improved outage module detection and staging, teams are evaluating the feasibility of leveraging the same streaming CMTS trap data to identify customer devices that have not returned to an active online state after an outage. The post-outage offline device resolution experiment focuses on the 75th percentile of devices that remain offline following an all clear.

By detecting post-outage offline devices, we can proactively engage customers to restore their service to a healthy state, eliminating calls and improving customer experience.

## **Conclusion**

Proactive Customer Engagement represents a new frontier in customer service. By leveraging increasingly available real-time datasets and harnessing the burstable compute power of next generation analytics frameworks, the cable industry has an opportunity to realize a cultural shift in customer engagement, pivoting from a reactive to a proactive stance. The crawl, walk, run approach to realizing PCE begins at familiarization with real-time data-sets and applying them to cross-over use cases such as outage deflection.

As our results demonstrate, there are considerable cost savings to be realized by harnessing the power of these real-time datasets.

1. We observed an improvement of ~ 160,000 calls and ~ 30,000 truck rolls through early outage detection in combination with IVR messaging and truck roll work order blocking.
2. We validated the application of a real-time dataset to a high-performance model leveraging next generation technology that is readily transferrable to the cloud.
3. We identified a fast follower use case for identifying those devices which have not returned to an online status after an outage is cleared for proactive remediation and customer engagement.

It is through the practical application of real-time datasets to achievable use cases that we will realize our evolution into proactive customer engagement.

## Abbreviations

API	application programming interface
AWS	amazon web services
CMTS	cable modem termination system
CPE	customer premise equipment
DOCSIS	data over cable service interface specification
HDMI	High definition multimedia interface
HFC	hybrid fiber coax
IVR	Interactive voice response system
MAMR	median absolute moving range
PCE	proactive customer engagement
UCL	upper control limit

# **The Evolution of Network Virtualization In The Home**

## **Improving User Experience And Manageability**

A Technical Paper prepared for SCTE•ISBE by

**Patrick Goemaere**

Chief Architect Cloud Services Connected Home CTO office

Technicolor

1626 Craig PI 90732 San Pedro US

+1 (818) 442 7183

patrick.goemaere@technicolor.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Content .....	5
1. Historical overview of SDN Technology .....	5
1.1. Introduction .....	5
1.2. Active Networking .....	6
1.3. OpenFlow .....	7
1.4. SDN Evolution .....	9
2. A survey of SDN solutions in the Home .....	11
2.1. Taxonomy of surveyed works .....	12
2.2. Generic theme around Home Networking Management .....	13
2.3. Specialized SDN in the Home themes .....	13
2.4. Some basic observations .....	15
2.5. Conclusions .....	16
3. Network function virtualization .....	17
3.1. Evolution of NFV .....	17
3.2. NFV use cases .....	22
3.3. Challenges .....	25
3.4. Future Evolution around NFV .....	26
3.5. Observations and Conclusion .....	28
4. Virtualization on residential CPE .....	30
4.1. The first wave residential vCPE .....	30
4.2. Current residential CPE landscape .....	32
4.3. Containerization and Edge compute .....	34
4.4. Software Life Cycle Management and Orchestration .....	36
5. Conclusion, The future of residential CPE .....	37
Abbreviations .....	42
Bibliography & References .....	44

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Active Networking Architecture .....	6
Figure 2 - OpenFlow Architecture .....	8
Figure 3 - SDN Full stack architecture .....	9
Figure 4 - OpenVSwitch (OVS) architecture .....	10
Figure 5 - Taxonomy of SDN use cases in the home .....	13
Figure 6 - Statistics on SDN in the Home related research papers .....	15
Figure 7 - Statistics on Cloud solutions versus CPE local solutions .....	16
Figure 8 - Openflow usage for in home SDN scenarios .....	16
Figure 9 - Network Function Virtualization Approach .....	18
Figure 10 - ETSI NFV Architectural Framework .....	19
Figure 11 - Major NFV Open Source Projects .....	20



Figure 12 - OPNFV NFV Stack .....	21
Figure 13 - CORD Technology Stack .....	22
Figure 14 - Some NFV use cases .....	22
Figure 15 - Universal CPE (uCPE).....	24
Figure 16 - Software Defined Wide Area Networking.....	24
Figure 17 - 5G Network Slicing .....	25
Figure 18 - Current NFV Industry Reality .....	26
Figure 19 - The three waves of NFV .....	28
Figure 20 - vCPE with SDN capabilities on CPE side .....	30
Figure 21 - vCPE with tunneling on CPE side .....	31
Figure 22 - Plume's OpenSync Architecture.....	32
Figure 23 - Cujo CPE Agent.....	33
Figure 24 - Glasgow Network Function High Level Architecture .....	35
Figure 25 - LeanNFV Key Value store integration .....	37
Figure 26 - Residential CPE architecture Layers.....	39

# Introduction

With the rise of Internet of things, the success of mobile computing and the consumption of more and more video related content, the home network has increasingly become a more complex environment, containing a rising amount of heterogeneous network devices that need to satisfy the demands of their inhabitants.

Compared with traditional enterprise networks, users typically don't have the technical skills, or want to be burdened with complex tasks as home network management, troubleshooting or configuration issues. The home infrastructure is predominantly cooperative and self-managed, with different type of devices often owned and controlled by different household members.

*“The technical know-how required to set up a network and run music or video across cables or wi-fi, is the elephant in the room that no-one wants to talk about.”*

In this respect, home networks face 3 systematic challenges:

1. Hard to Manage: Since home users lack the technical savviness to configure their home networks, they typically operate with default settings, and as a result are poorly secured, difficult to extend with new devices and services, resulting in a lack of functionality and experience desired by users.
2. Hard to customize: The home network cannot be customized for the needs of specific applications (Ex Home WiFi coverage, Video streaming) and lacks the means to rapidly introduce new functionality that spurs innovation.
3. Hard to share: In most cases the infrastructure is managed by a single provider with bespoke integration of other siloed solutions, which limits the user's choice and prevents other solutions to share the same infrastructure to keep cost low.

The challenges mentioned above are structural and cannot simply be solved by just putting a nicer user interface on top of the current network architecture. They arise from a mismatch between the stable end-to-end nature of historical Internet protocols and the fast-evolving nature of the home environment. Home networks today are still using the same architecture and internet protocols as defined in the 1970s for the whole internet and carry many of the assumptions made at that time. These protocols were designed with the assumption of devices operating in a trusted environment, the availability of skilled network and system administrators, and tried to accomplish a set of goals that simply don't apply for a home network.

The evolution of CPE equipment design has been extremely slow over the last decades, mainly driven by the fact that they are manufactured by different vendors and combine the forwarding of IP packets with proprietary control software and API's to control and configure these network functions. As a result, different management protocols are used to control individual devices, which has led to a very fragmented CPE landscape.

In the last decade, we have seen a powerful paradigm shift in the networking domain towards Software Defined Networking (SDN) and network virtualization functions (NFV), which has become mainstream for optimizing the complex and dynamic interactions around networking in the datacenter and cloud infrastructure. This technology shift started around 2008 in research as a response on the difficulties to manage today's networks, cost of equipment and interoperability issues.

Thanks to this evolution, and the agility that this new approach delivered, a lot of focus and innovation have been achieved around SDN and NFV technology, applicable in the context of the datacenter. Nevertheless, this technological approach can bring a lot of innovation in other domains as well, and since 2012, we have seen a lot of research work and publications around applying SDN/NFV technology in the Home networking context, which faces similar complexities which cannot be handled by today's traditional networking technology.

This paper will cover a general introduction of the SDN architecture and its evolution with a focus on OpenFlow as the enabling technology and will focus on the applicability and different use cases in the context of home networking scenarios found in various research work done in the last 5 years in this domain. Also, the evolution around NFV in the datacenter will be covered as its applicability for consumer CPE devices.

Since these techniques are now becoming a commercial reality, with different vendors that provide solutions for CPE equipment, the paper will also cover the gaps and issues that still exist.

Finally, the paper will discuss a way forward in the industry to evolve current CPE design to allow these new technologies to complement the current network design in a more open way, leveraging a fully virtualized CPE architecture, complying with standards, and ultimately addressing the quest for more collaboration in the industry around open source and open standardization.

## **Content**

### **1. Historical overview of SDN Technology**

#### **1.1. Introduction**

Software defined Networking techniques have a long history that started more than 30 years ago in a pursuit to make networks more programmable (Feamster, Rexford, & Zegura). Today SDN is considered as a key enabler, enabling innovation in how we design and manage networks. Albeit the term SDN is relatively new, it became popular as a technology paradigm driven by the rise of cloud and the datacenter for solving virtualization of the network infrastructure.

Originally, networks were designed using network protocols that went through years of standardization efforts and interoperability testing. Network administrators traditionally configured individual network devices using configuration interfaces that varied across vendors and even between different products from the same vendors. Today Customer Premise Equipment (CPE), like residential routers and gateways, still operate at the level of different protocols, management and configuration interfaces.

By nature, this process of standardization and deployments is slow in convergence and has frustrated many researchers and network service providers with the timescales, which were necessary to develop and deploy new network services, let alone to experiment and innovate in the networking domain.

SDN technology today fundamentally changed the way on how we design and manage networks, and typically has three profound characteristics.

SDN separates the control plane, that defines how we handle traffic, from the data plane which forwards traffic based on decisions made in the control plane.

SDN consolidates the control plane so that a single software program can control multiple data-plane elements remotely.

The SDN control plane has direct control over the state of the network data plane elements, routers, middle boxes, switches or servers, via a well-defined Application Programming Interface (API).

OpenFlow is a prominent example of such an API, while there is a multitude of controller platforms and frameworks that have emerged (NOX, POX, Onix, ONOS, ODL, OpenContrail, FloodLight). Today programmers have used these platforms to create new network applications such as network virtualization, dynamic access control, load balancers, etc.

## 1.2. Active Networking

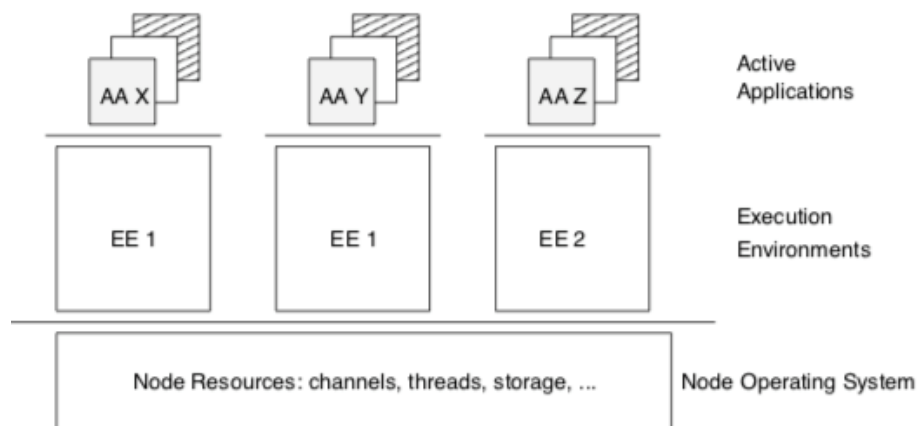
SDN borrow a lot of concepts and research from the area of active networking research (1996 – 2002), which was the first attempt to articulate a vision for programmable network infrastructure (Calvert). In contrast with the current SDN approach, active networking focused on making the data plane programmable and flexible, focusing on two distinct models to distribute networking code to intelligent nodes.

- Capsule model, where the code to execute on network nodes was carried in-band in data packets.
- Programmable router/switch model, where the code to execute at the network nodes was distributed by out of band mechanisms.

The following picture visualizes the architecture of an active node, where each node in an active network runs one or more execution environments (EE) and where each of these EE defines a virtual environment that operates on packets. An example of such an environment was a JAVA virtual machine extended to parse byte code programs carried in packets and send running code as packets.

Users invoke Active Applications (AAs), which provide code to program an EE, to implement end-to-end services.

An execution environment has access to node resources like computing, storage, hardware queues, etc, via a Node Operating System (NodeOS), which was responsible for managing and sharing these resources among EEs residing at that Node.



**Figure 1 - Active Networking Architecture**

The conception that packets would have to carry code written by end users created a lot of opposition in those days and made it possible to dismiss this evolution as inherently unsafe and too far removed from reality. This was further augmented by the lack of a clear business case or killer application for this approach.

Due to these arguments, the technology did not see widespread deployment, although it was the first technology that articulated a vision around a programmable network.

As a consequence, the next phase of research focused on a better demarcation between the functionality of the control and data plane, to enable a better focus on innovations in the control plane that presented a lower barrier for innovation than the data plane that also needed hardware evolution.

### 1.3. OpenFlow

Due to the initial failure of the first active networking research to address real business needs and the lack of acceptance, OpenFlow designs started from a more narrow and modest scope of problems to address, and initially started to focus more on routing and configuration management. The real innovation in these new approaches, was that there was a far cleaner separation between the functions in the data plane versus the control plane.

These first projects focused on trying to solve real problems in the network management plane, with a special attention on:

- Innovation by and for network administrators, rather than focusing on improvements for end users, or delivering an infrastructure for researchers.
- Programmability in the control plane, rather than in the data plane.
- Network wide visibility and control, compared to active networking, which was focused only at the device level.
- The concept of separation between control plane and data plane has been the basic architecture in all-further evolutions of SDN designs.


The key elements in this design and its success has been entered around:

- A logical centralized control plane using a well-defined and standardized open interface (OpenFlow API) towards the data plane.
- Distributed state management in the control plane, that could cope with resilience and scaling.
- Embrace of the Openflow standard in hardware switch design resulting in the increased availability of merchant-silicon chipsets providing commodity inexpensive hardware.

Due to the more pragmatic approach, Openflow provided a better balance between fully programmable networks and coarse grain control of the data plane, which could be enabled in hardware while still addressing real-world deployment use cases. Despite this somewhat limited flexibility in the data plane, Openflow was almost immediately deployable, thanks to the availability of existing hardware switches, which contributed greatly to its success. This success was accelerated due to the evolution of the modern data center and its needs towards network virtualization.

By design an OpenFlow capable switch has a table of packet handling rules (flows). Each rule includes a list of actions, drop, flood, forward to specific interface, modify the header field or send the packet to the controller and let the controller decide what to do with the packet. Each rule has a set of counters to track the number of bytes and packets and a priority to disambiguate between rules, which have overlapping

patterns. Upon receiving a packet, an OpenFlow switch identifies the highest-priority matching rule, performs the associated actions and increment the counters.




**Figure 2 - OpenFlow Architecture**

The decoupling of the control plane offers a better way to expose API's from the network to applications and middleware, and such API's are commonly defined as Northbound Api's. These Northbound API's define a central place in the infrastructure where the global application and network policies can be defined. These API's are independent from the southbound API's (E.g. Openflow), as they define a different level of abstraction, although the functionality offered by the Northbound Api's need to be translated and supported by the Southbound Api's.

For reliability, performance and scalability Controllers are typically designed in a distributed fashion. Depending on the implementation and deployment requirements, different solutions make different tradeoffs between the distribution granularity, the function partitioning, the data replication scheme and the consistency choices.

Since controllers provide a high level of abstraction, they are typically referred to as the network operating system. Due to the wide variety of SDN applications, different operating environments (Datacenter, Operators Core network, access network and on premises infrastructure), a lot of different implementation exists with different architectures, written in different programming languages (java, C/C++, Python, etc) both in the open source domain as well as in commercial solutions. It is beyond the scope to detail all these different frameworks, but two prominent and popular open source examples in this space are Open Network Operating System (ONOS) and OpenDaylight. They run cross-platform, are full featured and present high modularity. Besides support for OpenFlow they also support a multitude of different alternative Southbound API's like OVSDB, SNMP, NETCONF, COAP, etc.

The picture hereunder shows more or less a full stack SDN architecture and references some of the most popular technologies used in the space, like OpenvSwitch (OVS) which is the most common open source implementation of the SDN data plane.

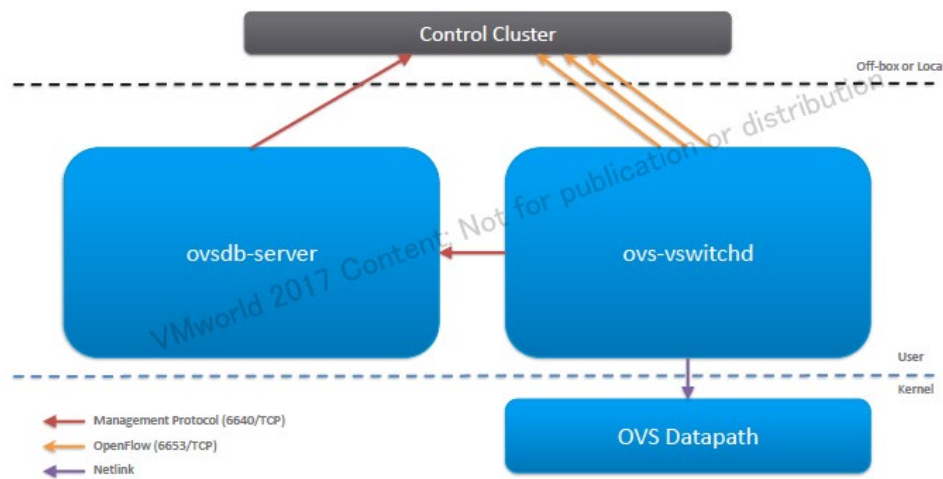


**Figure 3 - SDN Full stack architecture**

#### 1.4. SDN Evolution

OpenFlow as defined in 2008 is one of the most prominent technologies and has been used by most of the hyper scale cloud providers. It is also the most popular networking backend for Openstack deployments. One of the most mature open source implementations is OpenVswitch (OVS). OVS can operate both as a soft switch running within the hypervisor, and as the control stack for switching silicon.

One of the key components that differentiate OVS from other software switches is the native support of the OpenFlow protocol with multiple extensions. For this reason, OVS can operate either as an L2/L3 switch, a hybrid OpenFlow switch or a pure OpenFlow switch. See Figure.



**Figure 4 - OpenVSwitch (OVS) architecture**

Google has been one of the most active users of OVS in their public cloud, where OVS runs on all their servers in a distributed way to manage their network overlays and configure the networking between the local VM's or container's running on that host. For scaling purposes this platform, which is called Andromeda, has a in house developed distributed control plane that enables them to scale out their global network.

Azure is using their own virtual switch implementation called Virtual Filtering Platform (VFP) which is very close to the design of OpenFlow and follows a similar Match-Action Table (MAT) programming model, but using connections as a base primitive rather than packets and stateful rules as objects to allow use cases like network address translation (NAT) and better security with stateful access control list.

Today, OVS is not the only open source SDN capable software switch. Thanks to the evolution around network virtualization, the experience around different use cases and requirements that provides a better understanding related to cost and scaling issues and the evolution towards white boxes and move to Linux user space for data processing (DPDK), several new innovative opensource implementation's have come into existence.

Some of the most notable examples are FD.io (VPP), OpenContrail (now Tungsten Fabric), Distributed Data Path (DDP) and the P4 programming language.

OpenContrail originally developed by Juniper and later rebranded as Tungsten Fabric focuses on cloud networking and NVF use cases. It consists of two components, the controller and the data plane, called vRouter, which is conceptually similar to OVS but also provides routing and higher layer services. Opencontrails control plane addresses better scalability due to the distributed and horizontal scale-out design.

FD.io which stands for the Fast Data project, primarily focusses on drastic performance improvements. It uses Vector Packeting Processing (VPP) originally contributed by CISCO, which conceptually works on a graph structure (DAG), similar to many big data frameworks like Hadoop and Spark, to allow better parallelization of network data plane code, to exploit better utilization of caching and the parallel execution of code either by different CPU cores or hardware assisted logic.



One of the most promising innovations of the last years is the introduction of the P4 language, a programming language specifically designed and optimized for SDN network programming. the P4 Language Consortium has devised a solution to specifically handle forwarding instructions. This means it can exploit any advantages built into merchant silicon including everything from common ASICs and CPUs to advanced FPGAs and GPUs giving it the ability to perform at the same level as proprietary switches and routers, and even interoperate with them using the appropriate plug-in. At the same time, it can utilize standard SDN formats like OpenFlow while also supporting hardware-optimized legacy devices, providing a nifty bridge between old-style infrastructure and advanced virtual environments.

Today we are witnessing a fast adoption of the technology by switch manufactures and the white box ecosystem with recent announcements from Google to incorporate P4 in their edge infrastructure called Stratum and inclusion in AT&T's DANOS operating system.

With the ever-increasing network bandwidth to 25GB then to 50 GB and then to 100GB or even 400GB, hyperscalers like Amazon, Google and Microsoft realize that normal CPU's can't cope cost effectively with this evolution. Most application performance gains will no longer come from the linear Moore's Law increases in CPU power but from network accelerators (Bianchi & Bonola).

Network Interface Cards (NIC)s have been used for more than 30 years to connect servers and other computers to networks. Over the last 10 to 15 years NICs have become more capable, supporting higher-speed network interfaces, offloading basic network functions, such as TCP/IP, and more recently offloading virtualization. Smart NICs take this development one stage further by integrating a programmable resource that can be configured to provide additional CPU offload functions for different applications. Smart NIC's today can offer a mix of general processors' (CPU, FPGA or ASIC) specific capabilities to provide the right balance between good price performance, easy programmability and highly flexible.

The OVS project has followed many of these evolutions and today supports different configurations, either using DPDK for acceleration or/and using hardware offloading capabilities for their fast path. Using the Linux kernel traffic classification (TC) subsystem, Openflow provides an offloading mechanism called TC Flower, to allow exploitation of hardware capabilities. SmartNic's even allow it to run the complete fast path of OVS in hardware. Experimental support for the P4 language has been added in OVS as well.

It is clear that there is still quite some evolution in the SDN area, but the industry has finally agreed that there is not one solution that fits all, and that different use cases require different implementation approaches.

## **2. A survey of SDN solutions in the Home**

Although the concepts of SDN have been successfully applied over the last decade to solve the management complexities around Datacenter and cloud deployments, these concepts are so powerful that they can be applied in other context and environments as well.

In this chapter we explore the state of the art of research work found on using SDN concepts in the context of home networking scenario's.

The main driver for doing so, is the fact that the current networking approach expects low-level knowledge which goes far beyond the expertise and willingness of the average home user to configure

their home environment and the failure to improve this task by means of adapting the current networking stack due to the lack of flexible open interfaces (APIs) and programmability of the forwarding path.

The complexity of current home environments on the other hand keeps on increasing, since more and more heterogeneous devices are installed in the home environment, while also more innovative applications are getting introduced in the home besides the more traditional initial internet browsing activities of home users. Integrating these new devices and applications is challenging since there are two contradicting requirements that are hindering the acceptance of these evolutions: Ease of use versus tight control of these information flows to enforce user privacy and user preference; meaning controlling who is using these devices and where the information that is collected is sent.

In this respect, the main driver that made SDN successful in the datacenter context, the simplification of networking management, makes it interesting to be applied in the home networking context as well by virtualizing the home networking infrastructure. As explained in the previous, chapter SDN separates the control plane from the data plane which provides an abstraction from the lower networking layer into a logical network view that can be understood and programmed more easily by network programmers. By allowing access to low level network configuration by software programs, users can manage their home network by more user centric high-level applications (Mortier & Rodden, 2012). Thanks to the standardization and openness of the SDN approach these applications can be developed by third parties, allowing users more freedom and choice so that they can outsource network configuration and management to trusted third parties. To the extreme, the user's network could be partitioned or sliced to allow each application to work in full isolation but still have programmatic access to core networking functions.

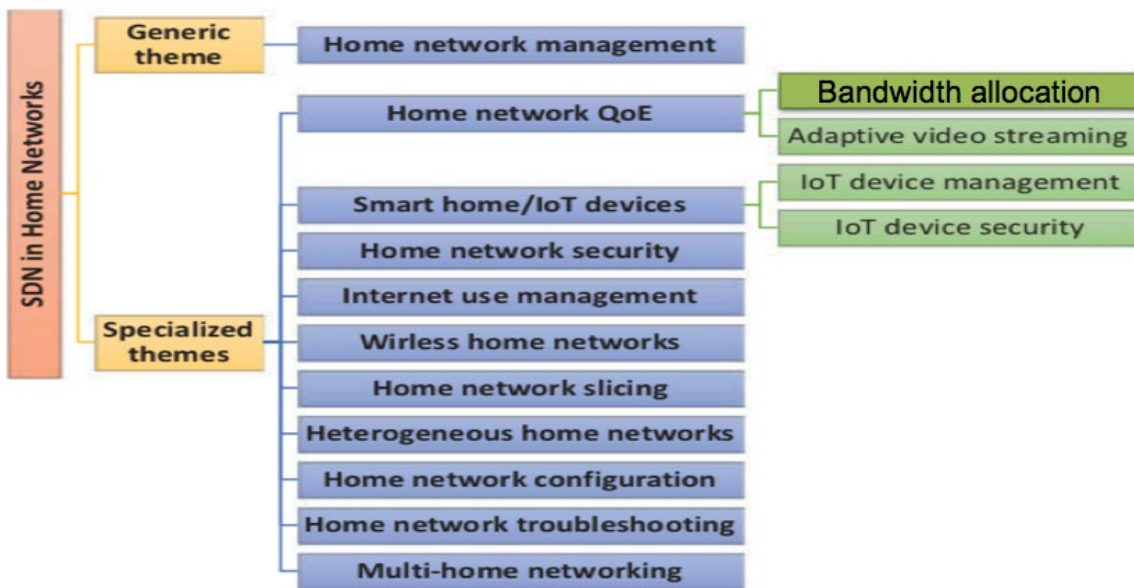
The need for a more open ecosystem for home devices, more in particular for the home gateways is not new, and has been previously addressed by trying to align on standardized execution environments like OSGI. Nevertheless, these environments do not offer the flexibility to easily change network specific functions, are written for a specific programming language and have to rely on a fragmented landscape of network management protocols and APIs which are hardwired in the legacy networking control functions. As a result` most of the research of the last years have focused more on an SDN approach and technologies as OpenFlow to address the problem of programmability of these type devices.

## **2.1. Taxonomy of surveyed works**

Based on a survey on research topics on applying SDN techniques in the home (Alshnta, Mohd, & Al-Haiqi, 2018), the authors have classified 42 articles written over the last 7 years around the use of SDN techniques in a residential environment and have observed different patterns and trends.

The following figure shows a taxonomy of the different use cases they encountered in these research papers.

Roughly two major themes have been identified, where one category is more generalized and centers more around overall improvements for home network management, while the others are more specialized around a heterogeneous set of different use cases.



**Figure 5 - Taxonomy of SDN use cases in the home**

## **2.2. Generic theme around Home Networking Management**

A few years after the initial publication of the Openflow specification in 2008, researches started to wonder on how these technologies could be applied in the home context. The first articles on this topic were developed as part of the homework project conducted in the University of Nottingham (Mortier & Rodden, 2012). This project was aimed as a fundamental redesign of the home-network infrastructure centered around the home gateway and based on applying SDN concepts to provide the home users with a much better view and understanding of their home network, while enable them to control their gateway with novel user friendly user interfaces.

Most of the other works in this category since then have followed the same line of thinking, whilst proposing other technical architectures and solutions to achieve the same improvements around usability of the home gateway. They all agree on virtualizing the home network and delegating the management and control of the home network to an entity in the cloud or datacenter, most probably the Internet Service Provider (ISP) or dedicated third party. Removing the management burden from the end user, and providing them with easy to understand user interfaces to control their home environment was typically implemented using the power and flexibility that Openflow provided to intercept different packet's and acting upon them, while leveraging dedicated configurations in standard network daemons like DNS and DHCP to stay compatible with the legacy network protocols.

The necessity of having to stay compatible with existing devices and network protocols in the home is a specific constraint in the home context which differs from the evolution of SDN in the datacenter and even more particular in the application of Openflow in Campus networks, where SDN provides more flexibility to allow a total clean slate approach to experiment and deploy new networking technologies or implementations.

## **2.3. Specialized SDN in the Home themes**

The rest of the surveyed research papers around SDN in the home, put the emphasis of their work on particular aspects of home networking and are summarized in 10 different subcategories.

The most popular subject in this subcategory is around improving Quality of Service (QoS) or more general improving the Quality of user Experience (QoE), typically by optimizing the bandwidth allocation for different applications in the context of video streaming or multimedia applications. These types of optimizations are either done in a more static way by allowing the user to express their preference related to these applications or can be done dynamically or adaptive by a combination of local traffic shaping based on collected traffic statistics. Most of the articles describe implementations that are running these algorithms from the cloud, while some of the articles run these algorithms locally using an in-home SDN controller to avoid latency issues.

Another theme is to address the issues related to the proliferation of IoT devices in the context of smart home. Most of these articles focus either on troubleshooting the smart home, or address the difficulties related to on boarding and integrating this heterogeneous set of devices in the smart home. Some of these papers are focusing on drastically improving the security aspects of these IoT devices, again relying on machine learning capabilities in the cloud to categorize the vast amount of different device types based for example on fingerprinting their traffic patterns observed in the home (Miettinen, Marchal, & Hafeez, 2016). Although most of the existing Low power Wireless protocols are not IP based and hence are difficult to integrate with an OpenFlow approach, the evolution of new IoT radio protocols to all IP based like Zigbee and Thread will likely make these types of approaches more viable in the near future.

Independent from IoT, network security in general is a common theme in some of the research papers. In most of these security solutions users out-source the management task related to security to a third-party controller who has the expertise and capacity to monitor and coordinate these activities over the Internet. It is mostly in this context that the limitations of the current Openflow specification are becoming very clear, due to the lack of very fine grain control and flow definitions in the current spec, which are due to the pragmatic choices that have been made originally.

Because the capacity of internet usage is a particular concern for home users and their family, several white papers address the management of internet usage through an SDN architecture, where specific configured policies can dictate how different users, devices and application can consume a fair share of the internet bandwidth. Most notable example here is of course providing parents a view and control mechanism on how their children can have internet access at dedicated time slots of the day.

Another group of papers address the specific issues arising from managing home WiFi access points, or in general all multi-technology wireless network devices. As wireless has become the dominated in-home networking technology, it is not a surprise that Openflow get's extended more and more to cover as well the Wireless space. Besides management of Wireless access points and clients, SDN concepts are also used to address improvement in finding algorithms that obtain a better usage and allocation of WiFi channels to improve the WiFi coverage in the home. Although there is new standardization, Ex. EasyMesh to interact with WiFi extenders, using extenders with an Openflow controller can be seen as an alternative to provide more flexibility between WiFi routers and extenders.


Few papers adopt the concept of network slicing (Yiakoumis, Kok-Kiong, Katti, & McKeown) (Boussard, 2018). Network slicing is a promising technique that creates different slices over the same physical home network, so that each slice is independently controllable and can be isolated for different services. In this context the current limitation in the OpenFlow specification around multi tenancy support for different stakeholders become apparent. Several suggestions are made as for example using a hypervisor approach like FlowVisor on top of Openflow to avoid conflicts with flow configuration of different tenants while still providing compatibility with the current OpenFlow API. Recent papers on this topic (ex. Nokia) suggest an even more novel approach by introducing the concept of a Software Defined

Lan, where Openflow access is more partitioned per different service (Spaces) and controlled by a separate cloud service.

Finally, the last 4 papers in the survey are focused on specific applications. The first one is the most generic one and uses the Openflow APIs for instrumentation to enable troubleshooting. The second one uses multihoming for different types of multimedia applications, while the third one is focused on using an SDN controller for home device recognition and registration. The last one is interested in the advent of upcoming 5G networks and subsequent evolution towards more heterogeneous wired and wireless technologies and uses Openflow for utilizing redundant links for flow rerouting and performing link switching between wired and wireless technologies both under normal conditions and in case of link failures


## 2.4. Some basic observations

Despite the slight regression in 2017, the last 3 years witnessed an increased interest of researchers in the topic of software-defined home networking, see figure.




**Figure 6 - Statistics on SDN in the Home related research papers**

Following Figure compares the number of proposals that are based on a third party or the cloud (e.g. an ISP) to the number of proposals that manage the home network locally within the home itself without outsourcing any task to the cloud.



**Figure 7 - Statistics on Cloud solutions versus CPE local solutions**

Following figure depicts the number of works based on the OpenFlow protocol. As expected, excluding a few works where there are no exact implementation details, OpenFlow is used in most of the works. This comes as no surprise for the most popular open protocol in SDN implementations.



**Figure 8 - Openflow usage for in home SDN scenarios**

## 2.5. Conclusions

Looking at the challenges and the complexity of modern home networking environments and their slow evolution over the last two decades, this chapter gives a good overview on how the flexibility and open standardization of a local in home SDN approach can bring the power to address a myriad of compelling use cases. The separation of control from the data plane enables the isolation from the traditional network stack from the legacy infrastructure to provide the home user, or any other third party, with a clean interface to control network operation and exploiting this new programmability of networks. It helps improving and simplifying the network management in a home environment for the average home user in ways which were not feasible without the move towards an SDN approach.

Most of the research work reviewed have focused on this task but produced different architectures and prototypes to prove these concepts and demonstrate the design. Although the majority of these white papers have an SDN approach at their core, they are independent of each other in their approach and implementation and as such most likely incompatible from each other. This mandates the need for further innovation and evolution in this space to come to a unified framework that can combine these different use cases, but for sure the potential to integrate the proposed ideas is great.

Most of these proposed solutions put the user in control to manage their home environment, in most cases with the help of a cloud backend, which amplifies the need for a novel user centric experience. The role of this interface is crucial to make the various functionalities and the information of the underlying home network visible and should allow the user to make changes in an effortless and intuitive way. With the advances in speech recognition and personal assistance concepts these user interactions will be expanded to make further improvements to the user experience.

A last important thing to note is that most of the examples suggest the involvement of a cloud based third party (ISP, or independent 3 party) to assist in the management of the virtualized home network. The biggest challenge here is to find the balance between the home users privacy and security versus the visibility offered to third parties and an economical viable split between local edge functionality versus moving functionality deeper in the network.

### **3. Network function virtualization**

As a recap from previous chapters, SDN technology is a novel approach that facilitates network management and enables programmatically efficient network configuration in order to improve network performance and monitoring. Thanks' to its flexible architecture and open API's the technology can be used in the home to drastically improve the user experience and interaction and address many novel use cases as discussed in previous chapter.

Network Function Virtualization (NFV) on the other hand is an architecture philosophy that utilizes Virtualization technologies to replace hardware proprietary network elements with Virtual Network Functions (VNFs) that can run as software on virtual machines and standard computing hardware. With the appropriate infrastructure, VNFs can be deployed anywhere in the network when needed, creating a network-as-a-service model for network utilization.

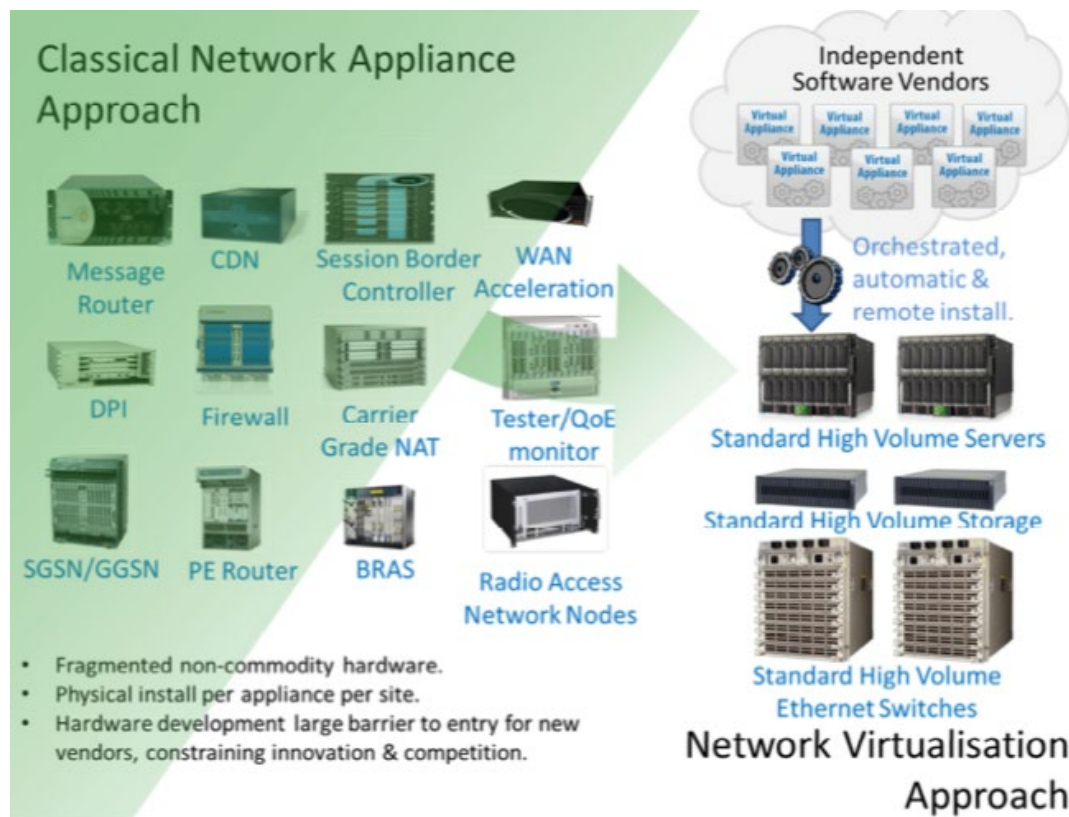
SDN and NFV are mutually beneficial but are not dependent on one another. You do not need one to have the other. However, the reality is SDN makes NFV more compelling and visa-versa. SDN contributes network automation that enables policy-based decisions to orchestrate which network traffic goes where, while NFV focuses on the services.

While SDN and NFV don't necessarily require each other to add value to an enterprise, collectively they are joined at the hip. Together SDN and NFV technologies allow better control for the entire network. These technologies can allow companies to increase the efficiency of their infrastructure and make it more functional. SDN and NFV offer the ability to add new services or modify existing services without making extensive changes to the physical network itself.

#### **3.1. Evolution of NFV**

NFV replaces network services provided by dedicated hardware with virtualized software. This means that network services, such as routers, firewalls, load balancers, core networking functions like DNS,

DHCP and WAN optimization devices, can be replaced with software running on virtual machines. Virtualized network functions are under the control of a hypervisor, which is the role that SDN fulfils in such a scenario. (See figure)



**Figure 9 - Network Function Virtualization Approach**

NFV services are deployed on commercial off-the-shelf (COTS) hardware platforms, typically running on Intel x86-based hardware and standard switching hardware (White Box). The combination of a base level hardware with this novel type of software creates a virtualized network that is not dependent on specific hardware.

Service providers have been leading the evolution towards NFV as a means to improving time to market and decrease their CapEx and OpEx related with their physical infrastructure. Today, [NFV](#) falls under the aegis of [ETSI](#), the European Telecommunications Standards Institute with a close cooperation with the Open Network Foundation (ONF) that brings together opensource initiatives and hardware reference designs.

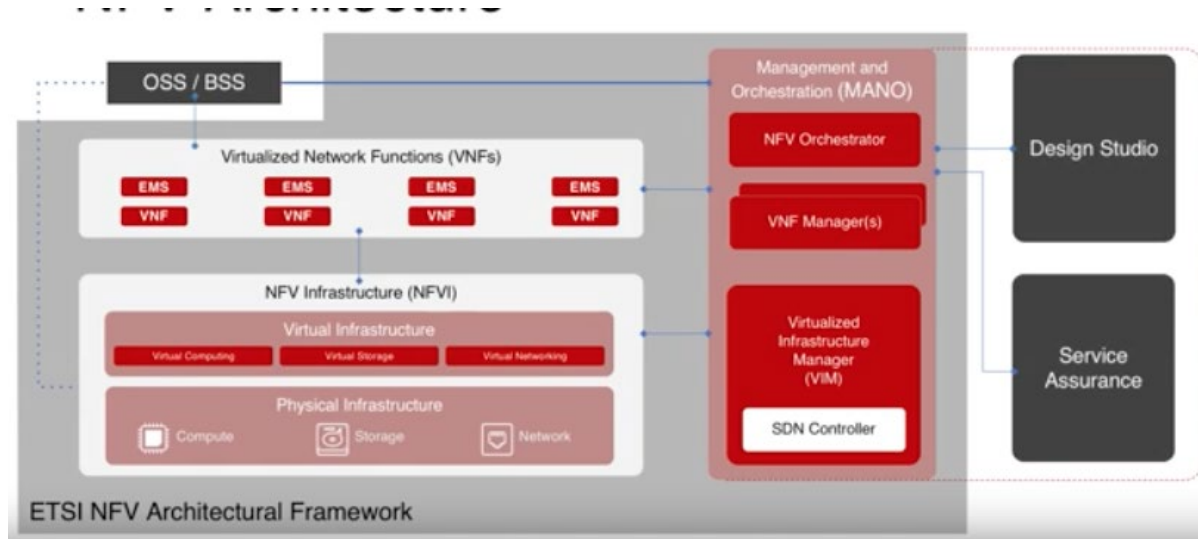
The goal of ETSI was to take various use cases and put it onto a common platform using virtualisation technologies, initially very hypervisor, virtual machine (VM) centric leveraging both opensource as vendor specific extensions. ETSI has played a pivotal role in the definition of the NFV architecture and numerous specifications.

NFV provides for an open architecture with many flexible options for deploying an NFV solution. The typical architecture of NFV consists of three distinct layers:



- Network functions virtualization infrastructure (NFVi) – the hardware and infrastructure software platform required to run network applications. (Openstack is by far the most popular opensource VIM, aka NFV cloud stack)
- Virtual network functions (VNFs) – software applications that deliver specific network functions, such as routing, security, mobile core, IP multi-media subsystems, video, etc.
- Management, automation and network orchestration (MANO) – the framework for management and orchestration of NFVi and various [VNFs](#).

A simplified architectural diagram is provided hereunder:

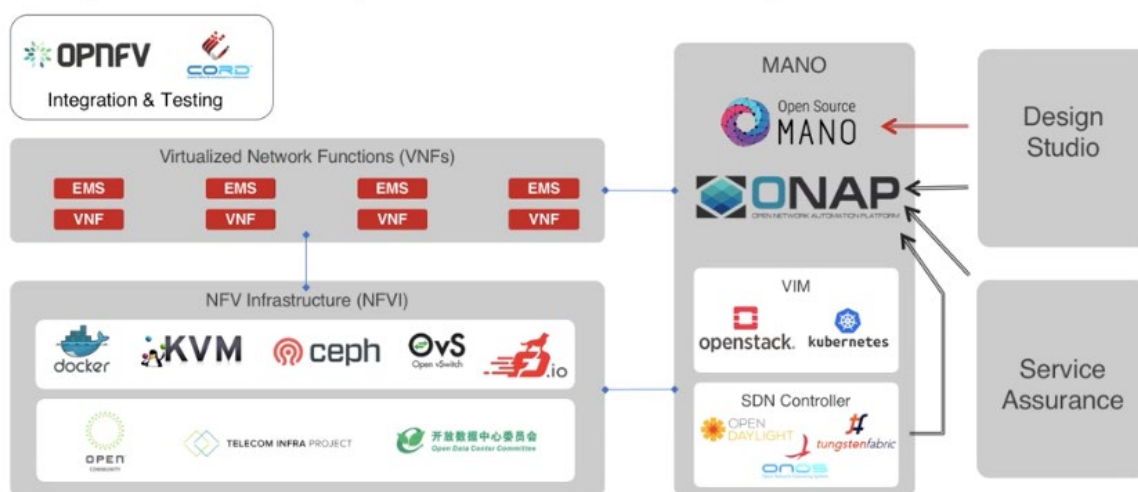


**Figure 10 - ETSI NFV Architectural Framework**

Open source has played a critical role in this evolution to accelerate the development of the NFV architecture in a way that it tries to complement standards and make them more open and accepted by the community. They have provided reference implementation to accelerate the adoption, validating the specifications while reducing Vendor Locking and creating transparency. Although there exists today an abundance of different open source components, most of the VNF implementations themselves so far have remained highly proprietary. The figure hereunder indicates some of the most relevant open source components used in the architecture.

# Major NFV Open Source Projects

CableLabs



**Figure 11 - Major NFV Open Source Projects**

Without going into the details of all these components, some of them deserve a closer look to understand their relevance:

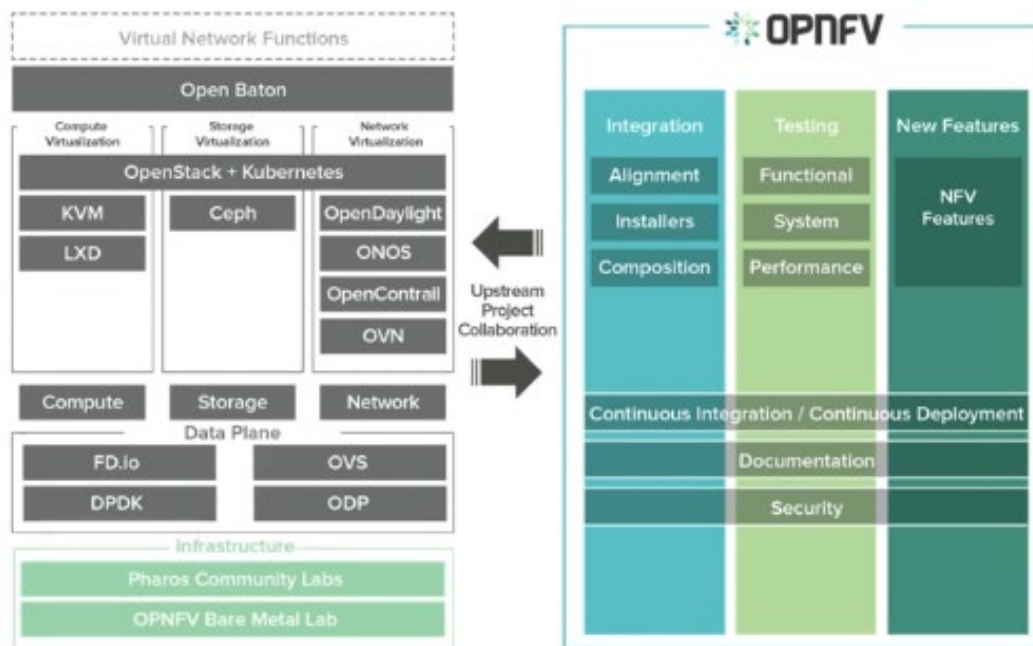
NFV orchestration plays a critical role within the broader set of MANO (management and orchestration) capabilities: NFVO tools are responsible for resource orchestration (the orchestration of NFV infrastructure resources across multiple virtualized infrastructure managers) and network service orchestration (the lifecycle management of network services).

Today the two major open source implementations are Open source MANO (OSM) supported by Telefonica and BT and Open Network Automation Platform (ONAP), which grew out of the merger of open source ECOMP(AT&T) and Open Orchestrator Project (Open-O) (China Mobile). Although they both cover NFV orchestration, ONAP also includes a unified design framework (supporting [TOSCA](#) and YANG inspired by Gigaspace commercial Cloudify solution) helping with end-to-end real time service orchestration and automation and provides the components around service assurance, closed loop monitoring and analytics. No real convergence has happened between these two stacks' so far.

Important to note is that in the movement from specialized hardware to more standardized off the shelf equipment like White boxes and enterprise datacenter/switching platforms, effort has been spent in the industry to also create open source hardware blueprints or reference designs.


The two major communities in this space are open compute platform (OCP) and Telecom Infrastructure Platform (TIP).

Open platform for NFV (OPNFV) aims at integrating various NFV related open source projects very close related to the ETSI NFV standards, providing a comprehensive testing framework while providing carrier grade functions. Several competing NFV technologies are provided offering a kind of ala carte menu to select from. (See figure)



**Figure 12 - OPNFV NFV Stack**

AT&T has been the driving force behind the Central Office Re-architected as a Data Centre (CORD) program to help the industry on converging the requirements for edge datacenters. Although initially focused on TELCO infrastructure, CORD today offers profiles and specific NFV functions for different industries like R-CORD for residential, M-CORD for mobile edge compute (MEC) and E-CORD for enterprise use cases. The open source activities are guided by the Open Networking Foundation (ONF), which is an operator led consortium. CORD core stack is depicted in this figure.

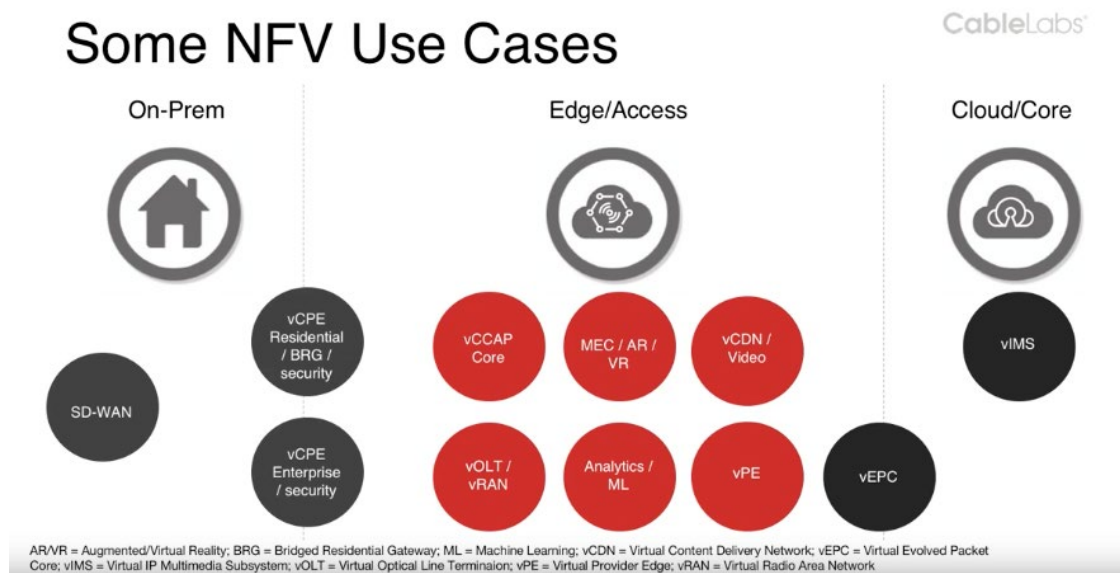


**Figure 13 - CORD Technology Stack**

### 3.2. NFV use cases

The amount of use cases that can be addressed by networking virtualization techniques is sheer endless (ETSI, 2017). There does not exist a proper consolidated taxonomy of all these use cases as almost all industries have a set of different problems they want to solve.

Hence for simplicity reasons I will give a limited overview of potential use cases in the Service provider space, covering both the Mobile and fixed access telecom and cable industry. (See figure)



**Figure 14 - Some NFV use cases**

The majority of use cases for communication Service Providers (CSP's) is centered around the transformation of their current infrastructure towards NFV by replacing existing network components by virtualized ones in order to reduce CAPEX and OPEX and have a more agile infrastructure. This typically includes, standard networking features like DHCP, DNS, Carrier grade NAT, DDoS protection, Firewall and Deep Packet Inspection (DPI).

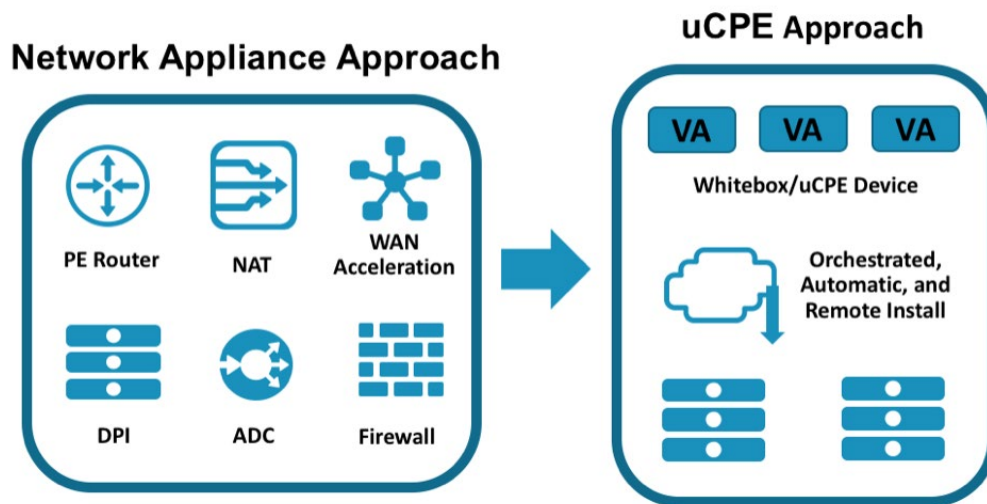
Also, the traditional components in the CSP Access/Edge or core network are being virtualized.

In the mobile industry this means virtualization of the Mobile core network, with NFV's like the virtual Evolved packet core (vEPC) and virtual IP Multimedia subsystem (vIMS). On the Edge this means virtualization of the Mobile base station or virtualized Radio Access Node (vRAN).

For Cable operators this centers around the virtualization of their Converged Cable Access Platform (vCCAP), while for Telco this evolution has been focused on their access infrastructure with components like the Optical Line Terminating (OLT) being virtualized (vOLT) and virtual provider edge router (vPE). Also virtualization of customer premise equipment (CPE) has received a lot of attention, both for residential deployments as for enterprise and small office/home office (SOHO) business related use cases. Albeit the first attempts were focused on keeping these vCPE devices as simple as possible, which resulted in a failure for residential deployments so far, the industry doesn't believe anymore that L2-based lightweight CPE will see significant traction, particularly in the enterprise space where L3-based full-function CPEs will be favored.

Current business vCPE have evolved towards a universal CPE (uCPE) device, which provides a better balance and flexibility between functions running at the CPE side versus functions running at the cloud side (ECI). The term uCPE has been coined by AT&T, which recently released their disaggregated network operating system (DANOS) to the Linux Foundation, together with releasing their hardware specifications of an uCPE box towards the open compute platform (OCP)

uCPE allows customer service providers to offer their enterprise and SMB customers enterprise functions as VNFs on a white box server or more commonly on a purpose-built device running at the customer premises. (See figure)




**Figure 15 - Universal CPE (uCPE)**

The most compelling use cases compared to just virtualizing the basic network functionality for potential cost reasons are those use cases which are innovative and typically are transversal or end to end and provide tangible and clear business advantage.

Key use cases in this area are:

Software defined WAN (SD WAN) (See figure) to reduce cost by using broadband connectivity, simplifies connectivity between remote customer sites and cloud providers and increases security by site to site encryption and micro segmentation of enterprise services.

## SD-WAN Topology (Software Defined Wide Area Networking)

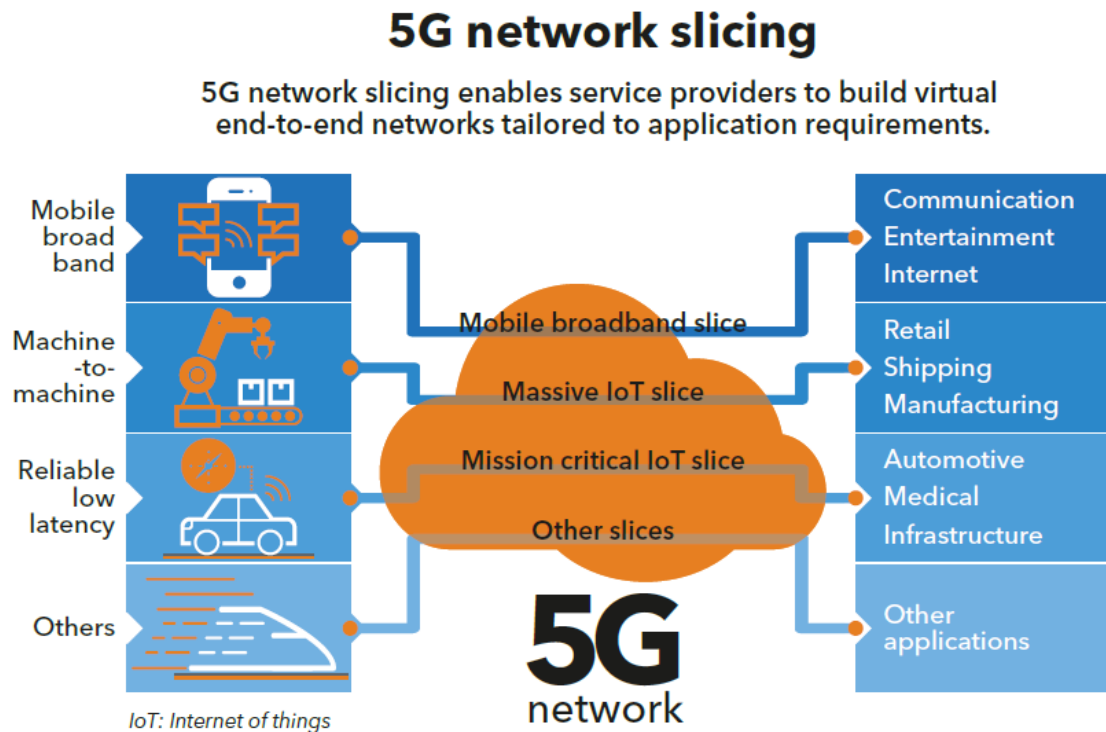


CableLabs

- Provides SD-WAN benefits: cost, control, visibility, optimized performance, efficiency, flexibility
- Add business services e.g. WAN acceleration, security, caching, and others

**Figure 16 - Software Defined Wide Area Networking**

Network Slicing, which allows running multiple logical networks on a common physical infrastructure, hence making the operator infrastructure multi-tenant for different applications. The key benefit of the **network-slicing** concept is that it provides an end-to-end virtual **network** encompassing not just **networking** but compute and storage functions too. This is one of the main use cases in the 5G-network evolution but can be applied for fixed line use cases as well. (See Figure)



**Figure 17 - 5G Network Slicing**

### 3.3. Challenges

This first phase of NFV, starting from 2014 to 2018 has been proving complex and difficult for many operators to deploy at scale. On one hand of the spectrum this process has been hindered due to the fact that it is a transformational process that forces operators to rethink their operational processes and have to go through a deep learning curve. On the other hand, the technical evolution around SDN and NFV has gone so fast that it was hard to keep up with all evolutions and challenges. Some of the problems that operators facing with the technology are the following:

- Lack of insight and information results in a service provider becoming locked into full stack virtualized solutions from a limited set of vendors.
- Lack of interoperability to access a full range of best of breed, trusted VNF's that can be easily and cost-effectively deployed.




- CPU, server cost, rack space and power required to meet the same performance cost of a dedicated custom designed solution, coupled with the operational cost of ownership is for most use cases more expensive.

Since most of the open source NFV projects are quite new, maturity and stability pose a significant barrier for carrier grade deployments. Operationally, VNF's are still multi-sourced virtual appliance, where each of them has to go through a complete lifecycle of sizing, deployment, configuration and upgrades which makes the on boarding process long and complex. Typical length is around 9 months', which negates the original objective around agility. The size of VNF images is quite huge which has to do the so called "lift and Shift" type of development and deployment. Meaning the first generation of VNF's where just repackages traditional network software. The software wasn't architected for virtualized environments.

Due to this lift and shift approach, these implementations where not designed of making use of the power of the cloud with respect to resilience and scaling and typically also lacked good open API's to automate them.

### 3.4. Future Evolution around NFV

Although the first generation of network virtualization has created a rich ecosystem of system integrators, equipment vendors (white boxes, enterprise infrastructure), software providers, open source frameworks and standards the amount of successful deployments and use cases have been rather limited. The resulting VNF's where rather large monolithic and siloed solutions, with a lack of interoperability running mainly on an Openstack-type of infrastructure. (see figure)



**Figure 18 - Current NFV Industry Reality**

Comparing to the Web scale company's aka cloud giants like Facebook and Google, one can notice that the challenges faced by CSPs are like those faced by the Web scale companies ten years ago. Web scale companies started by building their businesses on single monolithic applications. Over the time, they developed so-called micro service architectures. By deconstructing an application into smaller components which can be reused for other applications – and restructuring the IT organization into fully accountable micro service teams – companies can create more flexible, scalable and dynamic software development capabilities.



This service-based architecture paradigm in cloud-native distributed applications advocates for “smaller” services, i.e. micro services, to compose complex applications. These micro services are independent, small-scale processes that communicate through pre-defined APIs. Container technologies, including Docker, Kubernetes, DCOS etc., are in fact ideal for cloud-native applications as containers offer a lightweight atomic unit of computing compared to their very server centric VM predecessors.

In the decentralization of the operator’s network towards more and more edge infrastructure, one of the main advantages of container based micro-services is their portability. Cloud-native applications are designed to be portable to different deployment environments: for example, in a public, private, or hybrid cloud. They are well suited for edge and customer-premises solutions where compute resources are limited by space and power.

In this respect, the industry has started to move to take advantage of the low overhead of containers to deliver higher performance using cloud-native technologies to build the network functions, so they run in the same network and user space as the applications. Network functions are becoming part of the service topology. Network functions are truly just another service and can be developed and deployed using the same tools as the applications with the same velocity. CISCO has coined the term Cloud-native Networking Functions (CNF) for this new generation of virtualized network functions (CISCO, 2018) (Ericsson, 2018).

Today almost all of the traditional vendors have extended their solutions by integrating Kubernetes has a first-class citizen in their offerings providing a hybrid solution to support legacy VNF on VM's and new cloud native container network functions. They adapted their management plane to support a number of development and deployment pipelines for NFV management and orchestration (MANO). Major open source frameworks have been adapted as well as OPNFV and carrier grade solutions like SNAP from CableLabs.

The new 5G architecture heavily promotes a cloud native approach at their core and for multi-access edge computing (MEC) (5G\_PPP, 2018). The ONAP orchestration layer currently run’s on Openstack as on Kubernetes.

Most likely new edge deployment will operate with a much cleaner and smaller stack, removing the legacy Openstack infrastructure to keep deployments cost effective. Over time with improved security around containers like kata containers or gVisor from Google the difference between containers and VM’s will likely disappear as abstraction.

In the long run, autonomous networks are the desired future in which every element of the network is automated. High-resolution data would be evaluated to continually optimize the network for current and projected conditions. The result will be programmable mobile networks that leverage telemetry and analytics, automatically scale and self-heal, as well as deliver the highest service assurance.

The following figure presents the different evolutions has described in this chapter.

# The Three Waves of NFV

CableLabs®

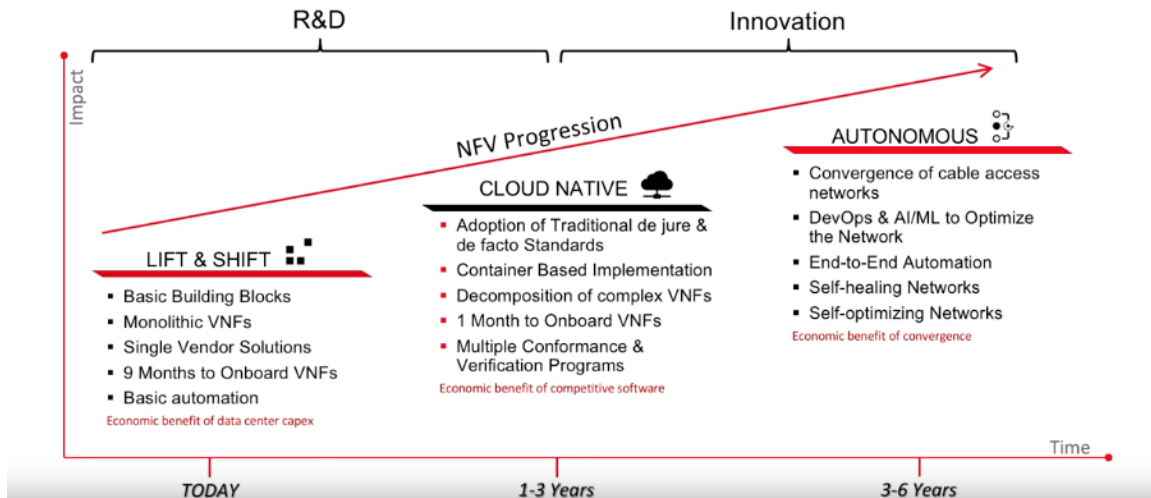


Figure 19 - The three waves of NFV

## 3.5. Observations and Conclusion

Communication Service Providers are under pressure mainly from the rapid advances that OTT players, web scale companies and new born in the cloud players, are making, thanks to the self-service model and extensive ecosystem's that public cloud environments provide. With increased user demands, the growth of internet of things and the emerging roll-out of 5G cellular, networks are becoming increasingly complex, while the pace of technological advances and innovations are creating a constant evolving architecture.

With an industry that evolves as quickly as the telecom industry, it is easy for hype to overcome progress. Network virtualization has been at the foreground of operator's strategy to improve their margins, agility to innovate and propel them in their journey towards digital transformation.

Nevertheless, the first wave around network virtualization has not enabled them to live up to the promises that the technology potentially could provide.

They have been initially centered around a monolithic approach, aligned with the first advances of server or machine virtualization technology platforms like Openstack and a type of lift and shift approach of traditional basic network technology stacks into network virtualized functions. Due to the fragmentation and immaturity of open-source initiatives and the reliance on traditional vendors who have a conflict of interest with the CSP transformation agenda to provide them with truly open systems. NFV 1.0 has not allowed telco's to scale or bring them the cost savings they need. While this transformation eventually results in a more efficient network with significantly lower TCO, typically the intermediate evolution stages will experience TCO increases resulting from transformation costs and manual processes.

Interoperability issues between different VNF solution providers have also played a key role in the slow progress of the industry.

With the advances of container technology, and the understanding that the new generation of network function's will have to adhere to more "cloud native" principles in order to horizontally scale out and provide the necessary resilience and high-availability in case of failure, the next wave of cloud-native network functions are embracing container orchestration frameworks like Kubernetes to be able to run more lightweight network functions in any public, private or hybrid cloud environment. This will become even more critical in edge computing, where resource utilization is key.

The initial slow progress around network virtualization and the technical and organizational difficulties have given operators a more realistic view on their timelines to evolve their infrastructure. Advance in 5G and automation technology will be key enablers for this next phase.

Despite all these difficulties the industry has learned valuable lessons to improve, while on the other hand a rich ecosystem of partners, equipment vendor's, standards and open source initiatives has matured to a point that gives more confidence towards the future. The industry is converging around a number of key tenants to achieve this progress.

A more pragmatic approach around standardization is necessary to achieve interoperability and the need for consolidation of the many different standardization bodies.

In this respect a better balance between ad-hoc standardization as witnessed in the public cloud (Ex Kubernetes and Docker) and standardization driven by standardization bodies need to be found, while standardization bodies should evolve towards more open standardization where the existence of an open source reference implementation will be key to drive progress and validate the design. The adagio of just enough standards should prevail instead of the very fragmented approach, which is typical in the early stage of technology evolution and disruption. The bottom line is that we need to accept that "the only constant is change." Innovation in software can bring many good things, but we need to learn how we can eliminate the silos, guard against new ones forming, create better interoperability, and simplify operational complexity.

Open source software—and now even open source hardware—will be critical components. Open source is all about public debate, and unlike previous transitions that were driven more by standards bodies driving consensus, the next generation of communications technology will be driven more by open source organizations like [OPNFV](#), [OpenStack](#), [OpenDaylight](#), [DPDK](#), [FD.io](#), [ONOS](#) and [ONAP](#) as the underpinnings of 5G will be virtual and very cloud-centric. Open source should drive standards and not the other way around.

Collaboration will become the preferred mechanism for driving change. No single vendor can deliver the full stack, and proprietary technologies will not keep pace with these future needs. This transformation will be delivered in virtualized (not physical) technologies, open source and multivendor, relying on significant integration work across many in the industry to be successful. In this respect we see already a shift in the value chain with new type of ecosystem player's, taking up the role of open source NFV/SDN-trained large system integrators. Examples here are companies like AMDOCS providing a fully integrated ONAP platform and even an AWS based development platform to allow quick development and players like RADISYS providing a first turnkey reference solution for M-CORD. Provided with additional services like a build, operate and transfer (BOT) business model it helps operators to focus on their use cases instead of spending their cycles on integrating, maintaining and evolving their infrastructure.

Improving automation and using a DevOps mentality to orchestrate their infrastructure in an end-to-end fashion will be key. The distinction between virtualized network functions and normal application

components will gradually disappear, and hence leveraging the evolutions in the cloud native automation DevOps domain will be key.


## 4. Virtualization on residential CPE

### 4.1. The first wave residential vCPE

While SDN and network virtualization have a rather strong business rational for enterprise and SOHO use cases and see evolution to uCPE, the residential application of these technologies has not been very successful. Most of the focus initially has been on cloud-based vCPE scenarios rather than running VNF's on premise or in a hybrid scenario, with the main objective of trying to achieve cost reduction by simplifying these CPE devices to mainly layer-2 based network devices and moving most of the layer-3 networking like NAT, DHCP, firewall and routing inside the access network.


Depending on the amount of flexibility required, two scenarios have been envisaged (Hermesh, Shulman, & Ray, 2016).

Using software defined forwarding at the CPE side, allowing an SDN controller to control and manage the network topology at the CPE side, see figure.



**Figure 20 - vCPE with SDN capabilities on CPE side**

An alternative to a fully virtualized network and using SDN capabilities at the CPE side is using tunnels as a bridge between SDN domains. Tunnels are logical point to point connections on top of an underlying network. Popular tunneling techniques include Layer 2 General Routing Encapsulation (L2GRE) over IP and Virtual Extensible LAN (VXLAN) over UDP. (See figure)



**Figure 21 - vCPE with tunneling on CPE side**

Since in the later scenario some GW functions remained locally and the complexity and operational expenses of the CPE remained, most of the initial marketing around residential vCPE focused on the first scenario or a hybrid where tunnels were configured by an SDN controller at the CPE side to allow a gradual transition from the existing architecture to a fully compliant SDN architecture.

Despite all hype and effort, the excessive attention towards a thin physical CPE device has never lived up to his promise. This is mainly due to a couple of reasons.

As described in the previous chapter on the NFV evolution, the rather slow progress and the transformational complexity has hindered the progress of vCPE deployment for residential use cases.

For most of the basic network functions, despite the flexibility a cloud model provides, the cost model cannot compete with similar function's executed locally on the CPE, since most of these functions require modest CPU performance and where the energy consumption is subsidized by the consumer.

The lack of compelling business use cases where consumers are willing to pay a premium for the service offered.

The closed nature and fragmentation of current CPE vendor landscape, which does not allow third parties to easily develop functionalities on a CPE device and deploy and evolve them.

The amount and cost of bandwidth required of sending a lot of the traffic towards the backend infrastructure and the rather asymmetrical nature of current access network, leading to trombone effects that becomes a bottleneck.

Latency issues with the approach which provides a sub optimal user experience due to round trip delays in the network.

The limited support of hardware offloading and acceleration techniques of typical low-cost System on Chip (SoC) designs to support integration of software based switches like OVS.

The lack of standards and API's inside the current CPE devices to become truly open and the non-existence of an execution environment that allows easy portability of different software agents on those devices which have a rather monolithic firmware approach.

The very VM centric approach of the first generation VNF's designs, making them unsuitable for deployment on these very constrained devices.

## 4.2. Current residential CPE landscape

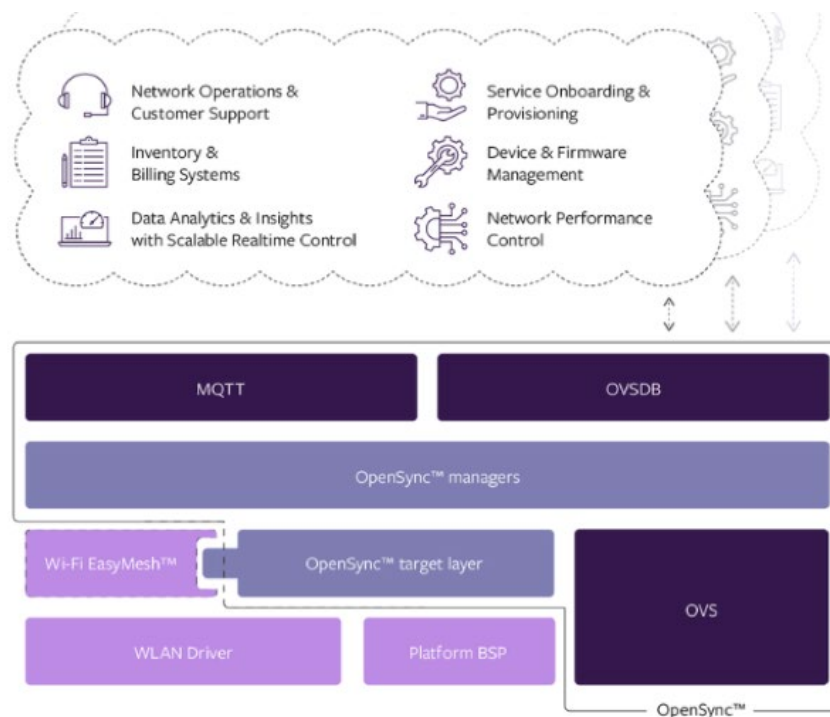
While the first wave of residential vCPE was heavily focused on running NFV's in the datacenter, new innovative players have shifted gear and have started to focus on how to leverage existing residential gateways to integrate their applications. These applications typically address the shortcomings in different domains of the home networking stack to provide compelling services around security, in home WiFi coverage and configuration, IoT use cases and enhanced monitoring capabilities.

Today these applications are built in a siloed fashion, leveraging their own backend/cloud solutions for control and management and relying heavily on analytics and machine learning.

Most of these solutions today are using a mixture of public cloud infrastructure, mainly AWS or Google and open source components to allow them some cloud agnosticism for their core infrastructure.

Some of the major players in this domain are Cujo (Security) and Plume (WiFi and security). They both recently open sourced their CPE agent and are both leveraging SDN network programmability, more specifically Plume is using OVS (Openflow) in their agent which is called OpenSync, while Cujo is using their own internal developed linux kernel module called NFLua which uses similar concepts as in the early days of active networking.

Plume's Opensync architecture is depicted in following figure.

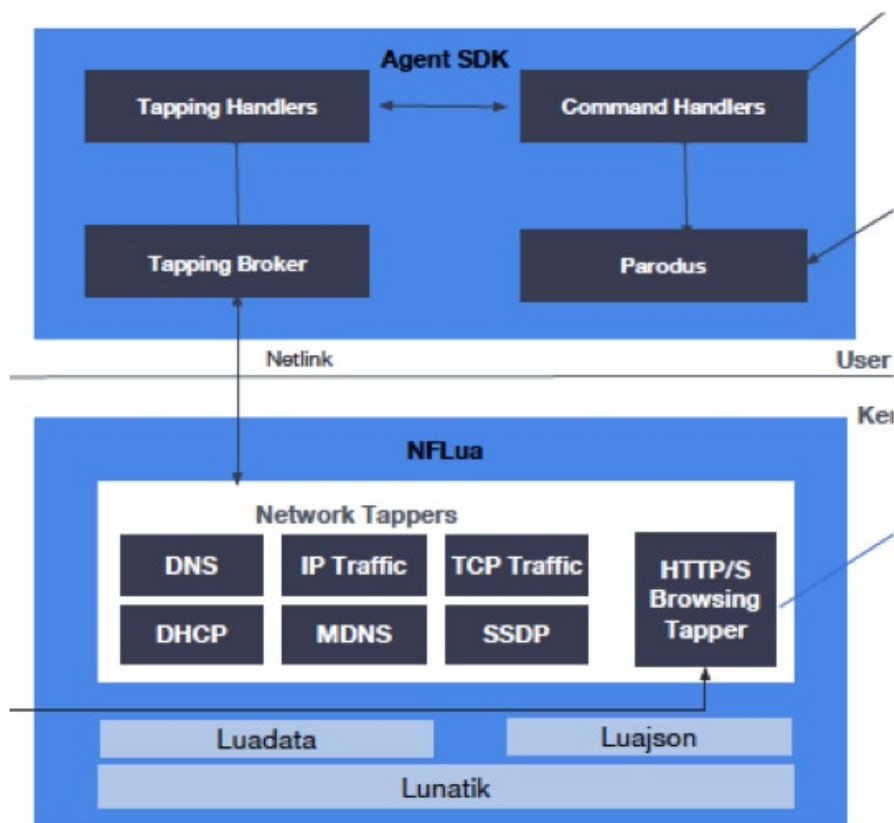


**Figure 22 - Plume's OpenSync Architecture**

For cloud connectivity, they are using OVSDB, which is a part of OVS, for control and management, while using MQTT for ingesting monitoring events in their cloud infrastructure.

The OpenSync framework on the gateway is centered around OVS, while they have created an abstraction layer, called the target layer to allow interaction and integration with the legacy networking daemons like Wireless, DHCP, NAT etc.

Cujo's agent architecture is depicted in following figure.



**Figure 23 - Cujo CPE Agent**

For cloud connectivity they are using XMIDT, that has been recently open sourced by Comcast, which has been built from the ground up in a cloud native way for scalability reasons.

Since their core use case focusses on security and requires quite some programable flexibility inside the data networking plane they opted to write an in-house performant and low footprint engine that could cope with the very constrained nature of these type of embedded devices with respect to memory and CPU budget. This engine which resides in the Linux kernel uses the LUA programming language, which is a data centric scripting language which has been very popular for low footprint embedded devices and gives them the flexibility they require to evolve and adapt their system towards new requirements.

Although these embedded agents are not following the exact definition as used for their cloud oriented NFV counterparts, the fact that they are using software defined networking techniques are putting these



types of agent's in the same category. Thanks to this new evolution, the industry has started to finally embrace technologies like OVS for residential gateways, and SoC manufacturers and traditional gateway OEM's are forced to start optimizing the integration of OVS to make use of hardware offloading and acceleration capabilities in their existing designs.

Most likely some convergence will start to happen around these two approached, centered around extending OVS and Openflow around residential use cases.

Today these agents are integrated as third party application's inside the legacy monolithically firmware of traditional CPE's, and hence are following the same slow feature evolution and deployment cycle as the rest of CPE functionality.

### **4.3. Containerization and Edge compute**

Next generation networks, driven by mobile, enterprise and internet of things evolutions are pushing the boundaries of computing more and more to the edge of the network in close proximity of the user, to cope with increased network utilization and to be able to guarantee performance and quality of service (QoS).

Due to the heavy footprint of current NFV platforms, the industry is shifting towards more lightweight solutions, that are easier to manage and operate at scale.

As an example, 5G mobile networks will utilize Mobile or Multi-Access Edge Computing (MEC) platforms at the edge of the network, in mobile access devices (E.g. base stations) with a key objective of having an IT service environment with cloud computing capabilities in close proximity to the mobile subscribers.

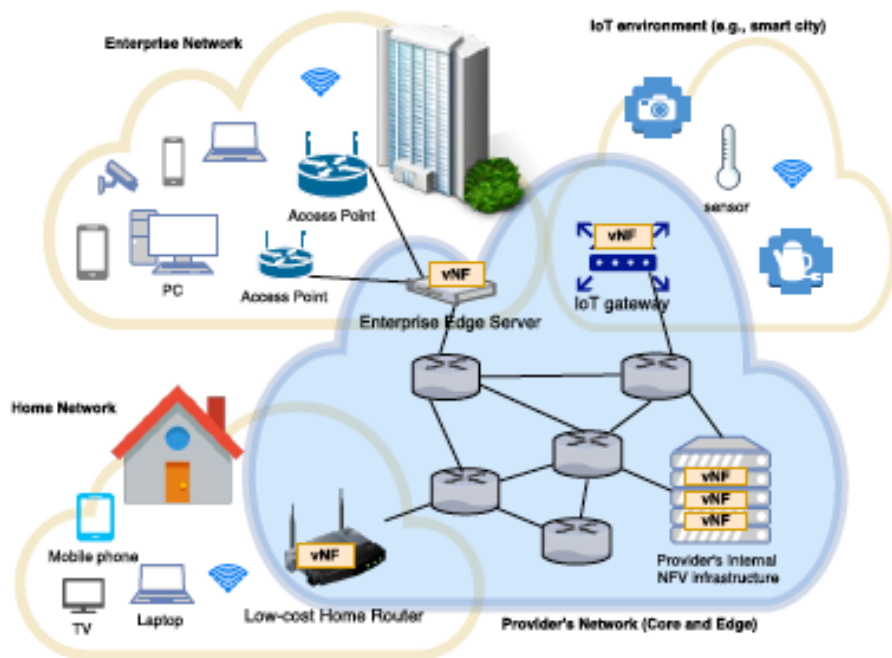
Other edge devices including enterprise edge servers, residential wireless routers and even IoT gateways that connect IoT devices to the internet.

As edge devices are located close to the users, services running at the edge provide higher forwarding performance (high throughput, lower latency) than running services remotely, and therefore save the utilization of the WAN infrastructure, which can correspond to savings of million dollars per year.

Due to this evolution, NFV's are starting to move to lightweight virtualization technologies, like containers to avoid the additional hardware requirements and overhead which is typically associated with using hypervisors and VM's.

One of the more recent research proposals in this field is published by the university of Glasgow, presenting Glasgow Network functions (GNF), that brings together NFV and edge computing by using generic lightweight Linux containers to host VNF's in a distributed, heterogeneous edge infrastructure (Cziva & Pezaros, 2017). The overall architecture is depicted in following figure.





**Figure 24 - Glasgow Network Function High Level Architecture**

In the past years different container technologies have emerged as a key technology for application isolation, packaging and delivery technology which are using core Linux kernel evolutions such as Control groups (Cgroups) for resource management and Namespaces for process isolation. Linux Containers (LXC) is a prominent example of such a technology, while today Docker has become the defacto standard for containerization, driven by its success of cloud deployments. Compared to LXC, the success of Docker is driven by its focus to deliver the developer with an easy user experience to package software inside a container, delivering a workflow that can easily be automated, and having a vast eco system of tools to be able to share images for different environments.

Recently, a more performant and lightweight version of Docker has become available in open-source by a company called Balena, which provides a docker runtime that can easily be integrated in constrained devices like home gateways and STB, with a footprint of under 30MBytes. Most CPE vendors start to offer support for both LXC as Docker as first-class citizens on their platforms, as a way to improve the integration effort for third party software and services.

In line with the cost/performance curve, the next generation of the broadband GWs will integrate a Modem, a router, best in class Wi-Fi, IOT radios (802.15), Gigabit Ethernet interfaces, multi core CPU (e.g. Quad core ARM), 512+ MB of RAM, 512 MB+ of Flash etc.; all for a familiar price point that is amenable to MSOs. Such integrated Gateways are capable of doing computing at the edge / customer premise

From an operator perspective, the Edge means the base station and their core network infrastructure, but extending their reach towards the physical edge, and applying these lightweight virtualization techniques on their customer premises equipment, which could give them end to end control and flexibility around their efforts on network virtualization, and allow third parties to innovate with new IoT use cases, directly on CPE devices.

#### 4.4. Software Life Cycle Management and Orchestration

While container technology provides a technology agnostic execution environment that is capable of running all types software and even network functions, written in a multitude of different programming languages, the scale of moving these type of computation towards the edge poses additional challenges on how to automate and operate these environments, as the amount and number of edge platforms start to increase. Certainly, when CPE devices are coming in the mix, bumping this challenge towards millions of these platforms.

CPE have been traditionally managed with technologies dating from original concepts as SNMP or TR69 for their configuration and Firmware life cycle management, which are costly and difficult to scale.

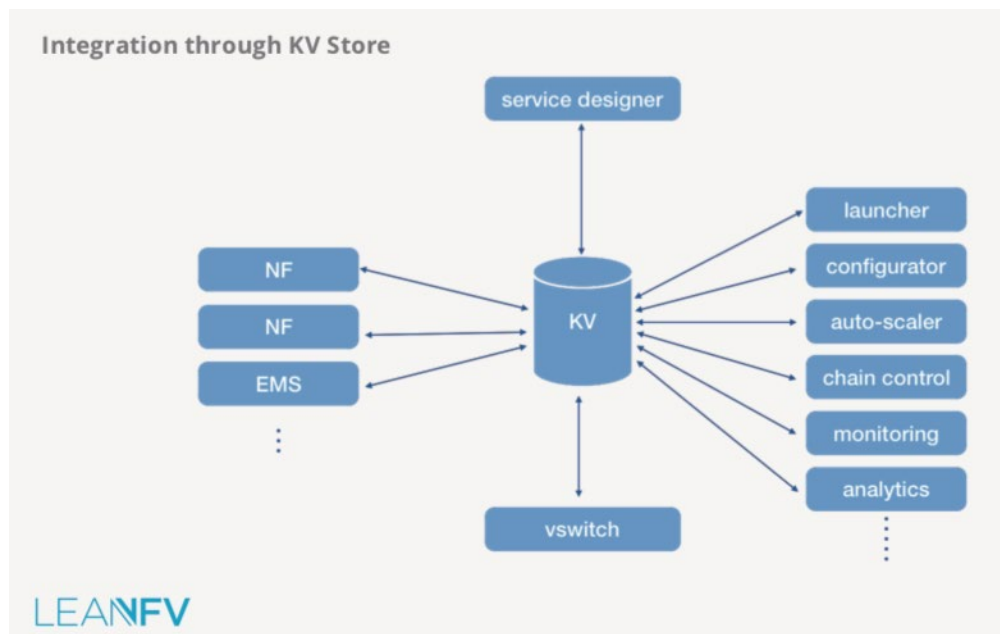
Current NFV platforms have the same type of challenge with respect to this type of software life cycle management or orchestration, even if they have moved to container orchestration frameworks like Kubernetes, and have shifted towards standardized frameworks like TOSCA to provide them with infrastructure as code blueprints to deploy applications and network functions. However, current NFV standards are falling short to provide orchestration that can cope with hundred thousand or even millions of disparate platforms.

Driven by the slow progress and frustration inside the NFV community, recently a group of NFV specialist from different leading vendors and operators have coined the term LeanNFV and released a white paper in order to pinpoint current design constrained that are hindering innovation and progress at the Open Network Summit 2019 in San Jose.

Their biggest complaint is the tight coupling of NFVs with the traditional software infrastructural components like Openstack or Kubernetes or SDN controllers, which increase complexity and complicates integration of VNF's from different vendors.

They propose the addition of a type of Key Value store with publish/subscribe semantics which serves as a universal point of integration and that provides a sort of central repository to enable VNFs to exchange state and discover each other.

This type of plugin integration enables NFVs to be plugged into any computational infrastructure, being it Openstack, Kubernetes or maybe in the future just a Docker enabled CPE devices. (see figure)



**Figure 25 - LeanNFV Key Value store integration**

The LeanNFV movement is off course still very immature and has certainly not gained wide acceptance yet in the industry. Nevertheless it addresses one of the core problems of the current NFV architecture, which will only become more relevant when NFV's are starting to migrate more and more to the extreme edge, E.g. CPE devices which will only run a limited environment for containers and where there is a necessity to decouple the life cycle management from the NFV implementation itself. As such, further API definition will serve as a guidance for NFV implementations to become truly cloud native while improving integration and addressing scalability concerns.

## 5. Conclusion, The future of residential CPE

“**The future** is already **here**. It's just not evenly distributed yet” is a famous quote from William Gibson, which perfectly captures the current state of technology around the further decentralization of cloud computing towards the edge and the softwarization that is happening due to network virtualization and software defined networking.

Despite the tremendous focus of the industry on the redesign of their core and access network, the acceleration that 5G will bring around MEC initiatives and the maturity around the white box ecosystem for enterprise and SOHO use cases with the introduction of universal CPE platforms, this decentralization is not a linear process.

In the shadow of all these evolutions new and compelling services and applications are getting directly integrated onto existing or next generation residential CPE equipment to enhance the home experience, being it network centric features or smart home, IoT centric evolutions. Albeit still vertically focused, resulting in siloed solutions, most of these providers have not waited for the operator's edge to materialize, and are working in a more top down approach using the public cloud as their backbone to deliver their end to end solutions, while heavily relaying on big data and analytics to augment their products while creating deep insight in user behavior and the home context.

In the current classical residential CPE architecture like home gateways network functions are deeply embedded inside the firmware in a monolithic fashion, while Soc's are deeply integrated and fully optimized in function of cost and forwarding capabilities, resulting in long integration and deployment cycles for new features.

With respect to network programmability there is no solutions fits all, hence today's evolving CPE market will have space for different shades of CPE.

The discriminating factor here is avoiding fundamental cost impact and the technology choice will be most likely made depending on the balance between Store and Forwarding needs of different use cases. This means that for devices like setop boxes where the cost impact is relatively small, DPDK will naturally have a bigger benefit for host originating traffic and storage driven applications, while for gateways which are heavily focused on forwarding use cases a better integration of OpenVswitch's data plane with SoC assisted hardware offloading will strike a better balance.

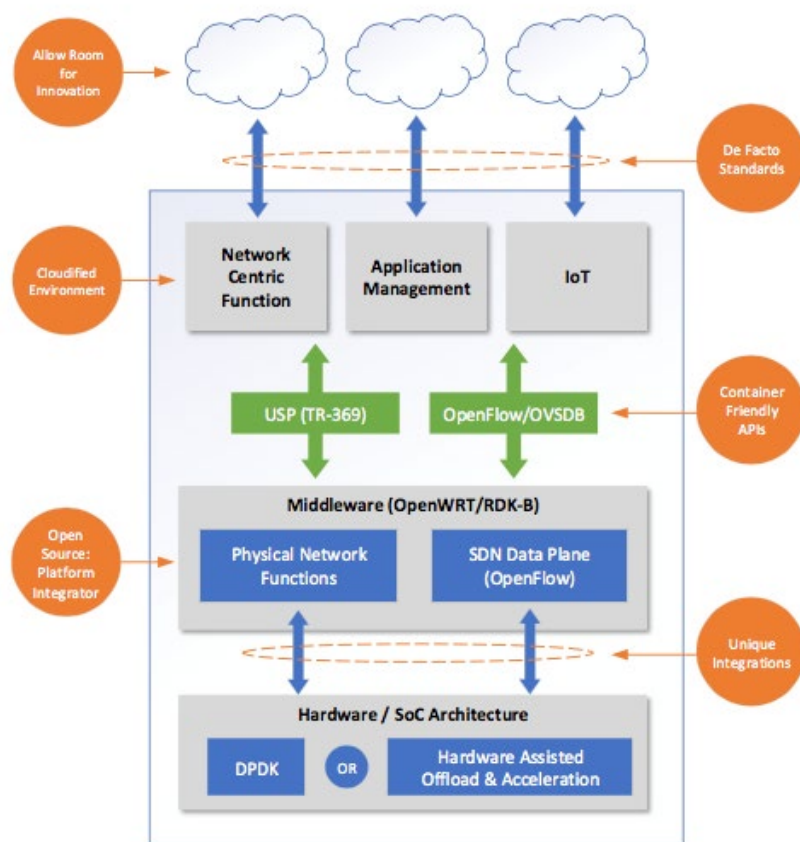
In order to be successful this process will be a gradual evolution from the existing legacy CPE architecture towards a more flexible and future safe platform rather than a revolution or big bang, which will have to strike a balance between offering compelling features and in home experiences and bringing them to market as fast as possible while leaving room for innovation and tackle the obstacles and technology gaps that still exists.

The main areas for improvements are better security and isolation, ultra-scalable orchestration and life cycle management for these new containerized components (pure application or network centric), improved network programmability of the SDN data patch in line with the hardware constraints of these type of devices (Bianchi & Bonola), a decentralizations and federation of the SDN controller path (Bhowmik & Tariq, 2015) and multi tenancy to allow different solutions to cooperate with each other and evolve to features like Software defined LAN and network slicing to offer true end to end QoS guarantees.

Due to the early stage of this evolution, this initial stage will be characterized with fragmentation in a lot of different initiatives, which is a typical cycle before the market will consolidate and mature standardization will set in. Today we already see different open-source initiatives from different vendors like Plume's Opensync and Broadband Forum USP providing different frameworks and others will follow.

To avoid the same pitfalls as witnessed in the SDN and network virtualization area which resulted in years of slow progress and severe fragmentation of the technology landscape the next figure will showcase a home gateway architecture that could drive an incremental approach and accelerate this evolution in drastic way by leveraging a natural evolution from the current legacy black box approach towards more of a "grey box". The figure also includes the basic building blocks to innovate on the network virtualization and SDN front.

The architecture is basically divided in three different areas that have distinct different life cycles in terms of design, development and go to market. (See figure)



**Figure 26 - Residential CPE architecture Layers**

At the lowest layer this cycle is closely linked to the hardware evolution, which is bound to a very slow evolution of ASIC or SoC development which usually takes a couple of years to be developed, while the resulting products then have a life time in the field of approximately 4 to 7 years. Since this layer includes the “data path” it has a clear need to contain “data path” programmability capabilities as expressed by SDN or to be exploited by network functions.

While not the perfect solution, OVS seems to be to most likely and promising candidate here thanks to its mature and vibrant ecosystem and since it offers both a standardized northbound API (Openflow or OVSDB), while it also supports the low level interfaces that allow SoC vendors to offload the data path further down in their ASIC’s. Please note that OVS is mainly limited towards level 2 forwarding, hence the industry will have to make an effort to extend OVS with more flexibility in this area while preserving the overall cost and performance of the total SoC cost. Ideally this should evolve to an industry specific domain programming language for the networking data path in line with the hardware constraint of these devices, being it something like P4 or a similar but optimized embedded version of it.

To allow faster experimentation of novel new approaches, the industry and research community should leverage ecosystems like off the shelf OpenWRT boxes or even trying to leverage the more Do it Yourself (DIY) ecosystems like the once we see maturing around the Raspberry Pi which deliver hardware that can easily be extended.

The middle layer contains the current more traditional networking stacks that we find in the current generation of residential gateways. Although there is quite some fragmentation with different existing solutions and vendors, today we see a consolidation of two major ecosystems emerging. On one hand OpenWRT has played a dominant role in this space to deliver a community based opensource ecosystems of classical Linux based network components, while Comcast has a more syndicated and controlled ecosystem around RDK both for routers as for set-top boxes.

Typically, the life cycle at this level is characterized by lengthy feature development and integration cycles of approximately 6 to 9 months, which require rigorous testing and qualification cycles for radios like WiFi. Operators are typically very conservative in deploying these new releases in the field, and this problem is aggravated due to the rather old school remote management solutions like TR69 or SNMP, which are costly to scale and don't provide a lot of feedback and monitoring to seamlessly upgrade the installed base without introducing considerable risk.

Due to the rather monolithically design and implementation of this layer, this typical lengthy cycle cannot be fundamentally improved. Nevertheless, this layer today contains a lot of in home and access networking features which are mature and have been added over a decade of deployments.

Rather than trying to redesign these basic functions or even try to move them away towards the cloud or datacenter like in the original vCPE effort, this functionality should be leveraged in combination with the new upcoming SDN functionality since, as indicated in most of the SDN in the home research, most use case will require a combination of the two approaches. The control plane management of WiFi is a good example of that but basic NAT, DHCP and DNS are similar examples.

Today this middleware layer also includes support for life cycle management of different containerized execution environment, which will serve as the foundation of the next layer.

In terms of northbound API's to manage these build-in network functions or allow eventing for monitoring purposes they should be chosen in such a way that minimal intrusion or rework is necessary in this layer. One of the most promising API standards as candidate is the newly released broadband Forum specification called user service platform (USP), which addresses the shortcomings of the current TR69 standard with respect to multitenancy, provides more modern container friendly API's and leverage the existing TR69 data model.

At the highest layer, the experience to develop or integrate new features should be as frictionless as possible, mimicking the self-service experience that made development in cloud environments so popular. Hence the use of containerization at this layer to isolate these components from each other is a necessary step to take. Besides isolation which provides a better security model, this containerization also allows developers to use whatever technology or programming language they are most familiar while providing excellent portability decoupled from the Linux environment they are running on. A key challenge here is to provide the proper software life cycle management or orchestration of these components, which requires an ultra-scalable implementation of this management layer.

Today, most of the third parties are using their own design or implementation for this cloud communication layer. Examples are either in house developed implementation like Comcast Xmidt, existing more enterprise or datacenter originated solutions like OVSDB itself or some are using public cloud provider offerings like AWS IoT. The later could be a good inspiration as an ultra-scalable cost-effective solution.

Also, Broadband Forum has realized that there is not a solution that fits all, and hence has decoupled their USP implementation from the underlying communication protocol and design.

In order to not jeopardize the needed agility and innovation in this layer, vendors should have the flexibility to select their own implementation at this moment to avoid premature standardization. Nevertheless, the industry should start to follow closely the efforts from the LeanNFV movement in order to come to consensus on a set of APIs for this cloud communication that could both entail management and control and synchronizing of state for these components between the device and the cloud.

The key aspect of this layer is that it allows secure end to end ownership for third party solutions decoupled from the traditionally slower life cycle of the monolithic firmware in such a way that every of the 3 layers can evolve at their own specific pace.

# Abbreviations

AA	Active Applications
API	Application Programming Interface
CapEx	Capital Expenditure
Cgroups	Control Groups
CNF	Cloud-native Network Functions
COAP	Constrained Application Protocol
CORD	Central Office Re-architected as a Datacenter
COTS	Commercial off-the-shelf
CPE	Customer Premise Equipment
CPU	Central Processing Unit
CSP	Communication Service Provider
DANOS	Disaggregated Network Operating System
DPDK	Data Plane Development Kit
EE	Execution Environment
ETSI	European Telecommunication Standards Institute
FPGA	Field-Programmable Gate Array
GNF	Glasgow Network Function
GPU	Graphics Processing Unit
ISP	Internet Service Provider
LXC	Linux Containers
L2GRE	Layer 2 General Routing Encapsulation
MANO	Management, Automation and Network Orchestration
MEC	Multi-access Edge Compute
NAT	Network Address Translation
NIC	Network Interface Card
NFV	Network Function Virtualization
OCP	Open Compute Platform
ONAP	Open Networking Automation Platform
ONF	Open Network Foundation
OLT	Optical Line Termination
OpEx	Operational Expenditure
OTT	Over The Top
OPNFV	Open Platform for NFV
OSM	Open Source MANO
QoS	Quality of Service
QoE	Quality of Experience
SDN	Software Defined Networking
SDWAN	Software Defined Wide Area Network
SOC	System On Chip
TCO	Total Cost of Ownership
TIP	Telecom Infrastructure Platform
USP	User Services platform
vCCAP	Virtual Converged Cable Access platform
vCPE	Virtual CPE



vEPC	Virtual Evolved Packet Core
vIMS	Virtual IP Multimedia Subsystem
VM	Virtual Machine
VPP	Vector Packeting Processing
vRAN	Virtual Radio Access Node
VXLAN	Virtual Extensible Local Area Network
SCTE	Society of Cable Telecommunications Engineers

# Bibliography & References

- 5G\_PPP. (2018). From Webscale to Telco, the Cloud Native Journey.
- Alshnta, M. A., Mohd, F. A., & Al-Haiqi, A. (2018). SDN in the home: A survey of home network solutions using Software Defined Networking. Malaysia.
- Bhowmik, S., & Tariq, M. A. (2015). Distributed Control Plane for Software-defined Networks: A Case Study Using Event-based Middleware. Stuttgart.
- Bianchi, G., & Bonola, M. (n.d.). Data Plane Programmability the next step in SDN. Roma.
- Boussard, M. (2018). Future Spaces: Reinventing the Home Network for Better Security and Automation in the IoT Era. Nozay.
- Calvert, K. (n.d.). Reflections on Network Architecture: an Active Networking Perspective. Kentucky.
- CISCO. (2018). Cloud-Native Network Functions (CNFs) White Paper.
- Cziva, R., & Pezaros, D. (2017). Container network functions: bringing NFV to the network edge. Glasgow: IEEE Communications Magazine.
- ECI. (n.d.). The definitive guide to vCPE, Tempering expectations and making NFV a reality.
- Ericsson. (2018). NFV Transformation Journey eBrief.
- ETSI. (2017). NFV Use Cases.
- Feamster, N., Rexford, J., & Zegura, E. (n.d.). The Road to SDN: An Intellectual History of Programmable Networks. Princeton.
- Hermesh, B., Shulman, S., & Ray, G. (2016). HEAD IN THE CLOUD, FEET ON THE GROUND: VIRTUALIZATION OF THE RESIDENTIAL GATEWAY.
- Miettinen, M., Marchal, S., & Hafeez, I. (2016). IOT SENTINEL: Automated Device-Type Identification for Security Enforcement in IoT. Darmstadt.
- Mortier, R., & Rodden, T. (2012). Control and Understanding: Owning Your Home Network. Nottingham.
- Santamaria, C. J. (2017). Management of a heterogeneous distributed architecture with the SDN. Reims.
- Sapien, M. (2016). Evaluating the trade-offs in applying NFV to enterprise services delivery. Ovum.
- Yiakoumis, Y., Kok-Kiong, Y., Katti, S., & McKeown, N. (n.d.). Slicing Home Networks. Stanford.

# Security Analysis Of 5G Mobile Networks

A Technical Paper prepared for SCTE•ISBE by

**Tao Wan**

Principal Architect  
CableLabs  
858 Coal Creek Circle, Louisville, CO 80027  
303.661.3326  
t.wan@cablelabs.com

**Mansour Ganji**

Lead Security Architect  
Rogers Communications  
8200 Dixie Rd, Brampton, ON, CA, L6T 4B8  
647.289.4679  
mansour.ganji@rci.rogers.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Background .....	3
Threats against Broadcasting .....	4
Threats against Paging .....	5
Threats against Unicasting .....	6
1. Pre-Authentication .....	6
2. Authentication .....	6
3. Post-Authentication .....	7
Conclusion .....	8
Abbreviations .....	9
Bibliography & References .....	9

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Simplified 5G System Architecture .....	3

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Messages to be sent and received out of authentication context.....	7

# Introduction

Cellular mobile networks have evolved from 2G to 5G over the past three decades. Mobile services offered by 2G, 3G, and 4G networks have always been voice calls and data network access. The introduction of 5G has changed this protocol by providing the communication technologies for many more use-cases tailored for each specific requirement. More specifically, 5G will provide network connectivity not only for human-to-human communications but also for human-to-machine, and machine-to-machine communications. 5G user equipment will fall into a broad range of devices where at one end they are fully-fledged computers, and at the other end they are single-purpose and resource-constrained IoT devices.

Because of the potentially significant impact on our society by 5G, its security is of critical importance and must be treated systematically. Researchers from both industry and academia have been working on improving security in 5G for a while. For example, the 3GPP SA3 working group has been studying and defining security specifications for 5G systems since 2017. Academic researchers are also helping to identify flaws in 5G specifications and are proposing enhancements. In this paper, we conduct a summary of security threats to 5G and prior generations of mobile networks and discuss how some of these threats are being addressed by the 3GPP 5G security standard.


Threats against cellular mobile networks can be generally classified into three categories: threats against user equipment or subscribers, threats against radio access networks, and threats against mobile core networks. In this paper, we focus on threats against subscribers. More specifically, we consider how subscriber security can be attacked by exploiting design constraints or flaws in control channels including broadcasting, paging and dedicated unicasting channels.

Due to the fact that neither broadcasting nor paging messages are authenticated in 5G (release 15) and prior generations, they are subject to spoofing, enabling many of the attacks against subscribers. Unicasting messages may or may not be security protected. Unprotected unicasting messages are also subject to spoofing and can be exploited to attack subscribers.

Through a summary of security threats against and defenses by 5G networks, we hope that a realistic understanding of expected 5G security can be established across the networking community, and hopefully among the general public as well.

## Background

A cellular mobile network including 5G consists of user equipment (UE), access networks and core networks (see Figure 1).



**Figure 1 - Simplified 5G System Architecture**

A UE is a device connecting to the cellular network to consume the services offered by the network, e.g., voice calls and data network access. A UE usually consists of an application processor running a general-purpose OS such as Android, and a baseband processor running mobile network protocol stacks (e.g., LTE

or 5G). A UE often contains a Universal Integrated Circuit Card (UICC) hosting at least a Universal Subscriber Identity Module (USIM) application, where a cryptographic key is stored and shared with the subscriber's home network and is the basis for mutual authentication of the UE and the network.

An access network is usually based on radio technologies, although other types of access networks including wireline access technologies are supported, e.g. in 5G (release 16). Radio access networks (RAN) manage radio resources between the UE and the next generation NodeB (gNB) to provide connectivity between the UE and the rest of the networks. 5G radio resources can be organized into local channels, including a Broadcasting Control Channel (BCCH), Paging Control Channel (PCCH), Common Control Channel (CCCH), Dedicated Control Channel (DCCH) and Dedicated Traffic Channel (DTCH).

5G core networks consist of virtualized network functions communicating with each other using web-based service requests and responses. The adoption of service-based architecture and virtualization technologies by 5G also result in the introduction of new network entities in 5G core networks, including Security Anchor Function (SEAF), Authentication Server Function (AUSF), Unified Data Management (UDM), Authentication credential Repository and Processing Function (ARPF) and Subscription Identifier De-concealing Function (SIDF).

## Threats against Broadcasting

A broadcast channel is used by the network (e.g., gNB) to broadcast system information for the UE to select and connect to the network. In 5G, gNB broadcasts a Master Information Block (MIB) and a number of System Information Blocks (SIB), some of which are always transmitted periodically, and others are only transmitted on-demand by UE. MIB contains physical layer information required by UE to establish radio links with gNB to receive the first SIB for cell selection. SIB1 contains information for UE to connect to a network including PLMN identifiers, track area code, cell identifier, etc. SIB2 to SIB5 contain information about cell re-selection, SIB6 to SIB8 contain public warning information (e.g., earthquake and tsunami warnings) and SIB9 contains information about time (e.g., UTC time and local time).

All user equipment needs to engage with all broadcast messages from all available eNB/gNB radio towers. This is to choose the network that it wants to connect to and then choose the frequency and channel that the base station is mandating it to use.

Since broadcasting messages are intended for all devices in an area, they are transmitted in clear text. Further, they are not authenticated for origin, nor protected for integrity in 5G (release 15) and prior generations. Therefore, all broadcasting messages are subject to spoofing and tampering. We consider four types of possible attacks:

First, the cell selection information can be forged to lure UE away from a legitimate cell, e.g., to a fake base station. More specifically, a fake base station can intercept all broadcast information from a legitimate gNB and rebroadcast the same information with higher power and with some modified elements (e.g., tracking area code) to fool the UE that it has entered a new tracking area and then reselect the fake cell. This is a known issue and has been actively exploited, e.g., to send fake short messages for fraud purposes [10, 11, 15].

Second, SIB3 to SIB5 contain a black list of cells which UE should not select. If this list is forged and cached, the UE may be subject to denial of service attacks if all available cells in an area are included in a faked blacklist. This attack has not been reported before and it is not clear how practical it is.

Third, SIB6 to SIB8 contain public warning information, which if spoofed may cause a public disturbance and instability. This is a known attack and has been demonstrated in LTE by broadcasting fake presidential alerts to a crowd [12] and it is applicable to 5G (release 15).

Fourth, SIB9 contains timing information, which can be spoofed to influence the time setting in UE. Since time is also critical to security, particularly in public key certificate validation (e.g., validating if a certificate has expired), spoofed timing information may lead to other attacks. We have not seen such an attack yet, but it is certainly possible.

## Threats against Paging

One of the requirements for the handset is to stay in a dormant mode while not actively using the network. This is both to reduce the battery consumption and also to minimize the network resource usage. While in this state, if there's an incoming call or a message to be delivered to the UE, the mobile network first pages the subscriber over its last known tracking area. Paging messages are sent over the paging channel in clear text without any authenticity or integrity protection. To protect user privacy, the subscriber's permanent identifier (SUPI) is not included in any paging messages in 5G. Instead, a Global Unique Temporary Identifier (GUTI), namely 5G-GUTI, is used.

5G-GUTI is assigned to the UE by the network (i.e., AMF) in the following situations [3]: 1) upon receiving a Registration Request message of types: a) initial registration; b) mobility registration update; and c) periodic registration update; 2) upon receiving a Service Request message in responding to a paging message. In this case, a new 5G-GUTI is sent to the UE by a UE Configuration Update Procedure. Note that in all cases, a new 5G-GUTI are also sent out to UE after NAS security context has been activated. An operator may implement a more frequent change of 5G-GUTI.

Attacks exploiting the paging messages can be classified into three categories: 1) location tracking; 2) denial of services; and 3) SUPI disclosure.

First, paging messages can be captured and used to determine the coarse-grained location of a UE upon the observation of the presence of an UE identifier of interest. Depending on the size of the area to which paging messages are sent, tracked location can be large or small. For example, in 4G/LTE, a UE can be tracked to an area of 2 km<sup>2</sup> (the size of an LTE cell) when the smart page is implemented. Since a paging area may become even smaller in 5G, location tracking can be more precise. Although the use of a temporary identifier (e.g., 5G-GUTI) with frequent changes can mitigate a location tracking attack, other flaws can still make it possible. For example, Torpedo (tracking via Paging message distribution) [4] exploits side channel information to track the user's location.

Second, paging can be exploited to deny the service of a UE. For example, an adversary can listen to the paging channel and respond to a paging request quickly so that the response from a victim is ignored by the network. In this case, the victim will not be able to receive its service (e.g., an incoming call or a text message). The attacker can also forge false paging signals and send them to the victim's handset device. Depending on the LTE or a 5G baseband modem on the target's phone, the forged messages may push the handset into a detached state. And if the attackers can continuously send the forged messages, they can cause a DoS attack on the victim.

Third, paging messages may disclose some information about a UE's SUPI even though 5G-GUTI is used and changed frequently. For example, it is discovered in [4] that an IMSI is used to calculate a paging occasion [16] which can leak the last 7 bits of the IMSI. This flaw is fixed in the next version of TS 38.304 [17].

# Threats against Unicasting

Threats against unicasting messages can be further classified based on the states of the UE in the process of authentication. More specifically, we classify such threats into three sub-categories, namely, threats against unicasting prior to authentication; threats against the authentication protocol itself, and threats against unicasting messages after authentication.

## 1. Pre-Authentication

The 3GPP standards contain specifications for securing the communication channel between user equipment and the network. But in all of the releases before 5G (Rel.15), they come into play only after the device has been authenticated and the security context has been built. Encrypting traffic between user equipment and the network needs a key, and that key would only be built during the security context creation. Therefore, every communication before that stage had to be made in clear-text. This issue has always been a challenge as this leaves space for eavesdropping on the communication channel and obtaining information that is sensitive in nature. The most important piece of data that can be revealed during this stage is the IMSI (International Mobile Subscriber Identification) which is unique to every subscriber. Obtaining IMSI is a big privacy concern as it allows tracking the subscribers' location. Besides, many of the attacks using DIAMETER signaling protocol require the attacker to know the victim's IMSI beforehand. So disclosing the IMSI opens the door to many more attacks. Although 3GPP had included using of temporary subscriber identifications like TMSI (Temporary Mobile Subscriber Identification) and GUTI (Globally Unique Temporary Identifier), there are some attacks reported to be run successfully that could force the user equipment to disclose the IMSI during a clear-text communication. [3]

This has been improved in the latest 3GPP technical specifications for 5G by adding an asynchronous encryption covering the whole authentication process. With this approach, the network operator will use a public-private key pair and the public portion of the key will be pre-provisioned on every subscriber's UICC (Universal Integrated-Circuit Card). The subscriber identifier has been renamed to SUPI (Subscriber Universal Public Identifier) in the recent 3GPP release document and it is encrypted before being transmitted over the air interface. The encrypted identity is called SUCI (Subscriber Universal Concealed Identifier) and is the main part of the information that is being transmitted for the authentication period [3].

Prior to the authentication and key agreement, certain RRC layer messages need to be exchanged between the UE and the network, which are subject to spoofing and tampering. Two examples of such RRC messages are RRC\_UECapabilityEnquiry and RRC\_UECapabilityInformation. Some NAS messages (e.g., NAS Service Reject) may also be sent out to UE prior the establishment of security context. Those unprotected messages can be exploited to attack both the UE and the network.

## 2. Authentication

The authentication process in 5G continues to use an AKA algorithm like the previous 3G and LTE generations; the algorithm is called 5G AKA. There are two new authentication methods that have been added to the list and they are EAP-AKA' and EAP-TLS (only in Non-Public network or isolated deployment). Choosing which authentication method to utilize will be a decision made by the home network.



As mentioned earlier, the authentication process has improved from the previous generations. One of the big improvements is adding the home network's public key to the process. With the subscriber (UE) having the home network's public key, it can start encrypting sensitive authentication data such as SUPI right from the beginning of the authentication process. The second improvement is reducing the trust in the serving network for authenticating the roaming subscriber. In the previous network generations, the serving network had the option to fake the presence of a subscriber, thus tricking the home network into updating a subscribers' location, dropping the legitimate security context and running a DoS as a result, or even redirecting SMS and connecting to a malicious or compromised serving network. The new advances in 5G roaming authentication prevent all these issues because the home network will not authenticate a roaming subscriber, unless it implicitly receives the data that it expects from the subscriber in the authentication process. It is only after this stage that the home network passes the encryption and integrity checking keys to the serving network along with the SUPI. [6]

However, these improvements require the channel between the home network and the serving network to be authenticated and encrypted as it now passes sensitive information like  $K_{SEAF}$ . Since this is a design consideration rather than a 5G specification, we will not dive deeper into this context.

There are also a number of papers [4,5,6] which find some security issues with the 5G authentication and key agreement protocols. Those issues include information disclosure from side channel, race condition exploitation, and improper protection of SQN. 3GPP is working on to mitigate some of those issues.

### 3. Post-Authentication

There are certain unicasting messages that cannot be protected due to design limitations, even after security contexts have been established and are in use between UE and the network after a successful authentication and key agreement. Those messages can be sent out in clear texts with neither confidentiality nor integrity protection. Thus, they are subject to spoofing and tampering attacks. These messages include:

**Table 1 - Messages to be sent and received out of authentication context**

Network to UE	UE to Network
IDENTITY REQUEST	IDENTITY RESPONSE
AUTHENTICATION REQUEST	AUTHENTICATION RESPONSE
AUTHENTICATION RESULT	AUTHENTICATION FAILURE
AUTHENTICATION REJECT	SECURITY MODE REJECT
REGISTRATION REJECT	REGISTRATION REQUEST
DEREGISTRATION ACCEPT	DEREGISTRATION REQUEST DEREGISTRATION ACCEPT
SERVICE REJECT	

These messages are not confidentiality or integrity protected as they may sometimes need to be communicated out of the security context. But, this leaves a possibility for an attacker to spoof either the subscriber or the network and cause a service disruption to a subscriber. If the attacker is in a position to continuously send a malformed message to the subscriber or to the AMF, they can run a targeted DoS against a specific subscriber. [13]

This also applies to the emergency service requests where both confidentiality and integrity need to be set to “null” due to the nature of the call.

## Conclusion

There has been significant improvement in 5G authentication in comparison to the previous cellular network generations. Nevertheless, there are still areas that are susceptible to potential attacks or misuses. Most of the cases where a vulnerability still exists are information exchanges that can or should be allowed without a security context in place. These are messages that are supposed to be available to all subscribers in an area or channels to be present for emergency communications. The current trust model for cellular mobile networks is based on isolated trust within each individual service provider. In other words, there are many trust trees in the communications industry and each tree has a root within a specific service provider. Outside that trust tree, the subscribers and service providers have no choice but to interact with “untrusted” entities that have the potential to be malicious. With this model, any improvements in the authentication process and trust establishment will remain a localized attempt and will not help with the global trust.

One possible solution for this issue is to follow the same model or models that have been proven to work in the public Internet access area. The trust model built for safe web browsing, for example, could be a good start for more developments in the cellular network. Having global trust anchors endorsed by the home network that all of its user equipments can trust could be a viable solution. In such a model, even the communications that have to be outside subscribers’ security context, can still be digitally signed so that the subscriber can verify that a message is coming from “a legitimate service provider” even if it’s not its own home network. We hope that 3GPP will embrace digital signature based solutions which can mitigate most of the threats discussed in this paper.

# Abbreviations

3GPP	3 <sup>rd</sup> Generation Partnership Program
AKA	Authentication and Key Agreement
AMF	Authentication Management Function
DoS	Denial of Service
EAP	Extensible Authentication Protocol
gNB	Next Generation NodeB
GUTI	Global Unique Temporary Identifier
HN	Home Network
IMSI	International Mobile Subscriber Identity
ISBE	International Society of Broadband Experts
LTE	Long Term Evolution
SN	Serving Network
SUPI	Subscriber Permanent Identifier
SUCI	Subscriber Concealed Identifier
TLS	Transport Layer Security
UE	User Equipment
UICC	Universal Integrated-Circuit Card
SCTE	Society of Cable Telecommunications Engineers

## Bibliography & References

[1] 3GPP TS 24.501. “Non-Access-Stratum (NAS) protocol for 5G System (5GS), Stage 3 (Release 15). Jan 2019.

[2] 3GPP TS 38.321. “NR; Medium Access Control (MAC) protocol specification (Release 15). V15.5.0, March 2019.

[3] 3GPP TS 33.501. “Security architecture and procedures for 5G system, (Release 15)”

[4] Hussain SR, Echeverria M, Chowdhury O, Li N, Bertino E. Privacy Attacks to the 4G and 5G Cellular Paging Protocols Using Side Channel Information. In 26th Annual Network and Distributed System Security Symposium, NDSS, San Diego, CA, USA, February 2019.

[5] Borgaonkar R, Hirschi L, Park S, Shaik A. New Privacy Threat on 3G, 4G, and Upcoming 5G AKA Protocols. Proceedings on Privacy Enhancing Technologies 2019.

[6] David Basin, Jannik Dreier, Lucca Hirschi, Saša Radomirovic, Ralf Sasse, and Vincent Stettler. "A formal analysis of 5G authentication." In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS'18), October 2018.

- [7] Cremers, Cas, and Martin Dehnel-Wild. "Component-Based Formal Analysis of 5G-AKA: Channel Assumptions and Session Confusion." In 26th Annual Network and Distributed System Security Symposium, NDSS, San Diego, CA, USA, February 2019.
- [8] Golde N, Redon K, Seifert JP. Let me answer that for you: Exploiting broadcast information in cellular networks. In Proceedings of the 22nd {USENIX} Security Symposium ({USENIX} Security 13) 2013 (pp. 33-48).
- [9] Shaik A, Borgaonkar R, Asokan N, Niemi V, Seifert JP. Practical Attacks Against Privacy and Availability in 4G/LTE Mobile Communication Systems. In 23th Annual Network and Distributed System Security Symposium, NDSS, San Diego, CA, USA, February 2016.
- [10] Marc Lichtman, Raghunandan Rao, Vuk Marojevic, Jeffrey Reed, Roger Piqueras Jover. "5G NR Jamming, Spoofing, and Sniffing:Threat Assessment and Mitigation"
- [11] Roger Piqueras Jover, Vuk Marojevic. "Security and Protocol Exploit Analysis of the 5G Specifications". IEEE Access Magazine, Volume 7, 2019
- [12] Gyuhong Lee et al. "This is Your President Speaking: Spoofing Alerts in 4G LTE Networks". MobiSys '19, June 17–21, 2019, Seoul, Korea
- [13] 3GPP TS 23.501 V16.0.2 (2019-04). System Architecture for the 5G System; Stage 2 (Release 16)
- [14] 3GPP TS 38.331. "NR; Radio Resource Control (RRC) protocol specification (Release 15). V15.5.0, March 2019.
- [15] Li, Zhenhua, Weiwei Wang, Christo Wilson, Jian Chen, Chen Qian, Taeho Jung, Lan Zhang, Kebin Liu, Xiangyang Li, and Yunhao Liu. "FBS-Radar: Uncovering Fake Base Stations at Scale in the Wild." In Proceedings of NDSS. February 2017.
- [16] 3GPP TS 38.304. "NR; User Equipment (UE) procedures in Idle mode and RRC Inactive state (Release 15). V15.0.0, June 2018.
- [17] 3GPP TS 38.304. "NR; User Equipment (UE) procedures in Idle mode and RRC Inactive state (Release 15). V15.1.0, September 2018.

# Authentication In 5G Wireline And Wireless Convergence

A Technical Paper prepared for SCTE•ISBE by

**Tao Wan**

Principal Architect  
CableLabs  
858 Coal Creek Circle, Louisville, CO 80027  
303.661.3326  
t.wan@cablelabs.com

**Yildirim Sahin**

Principal Wireless Engineer  
Charter Communications  
6360 Fiddlers Green Circle, Greenwood Village, CO 80111  
720.536.9394  
yildirim.sahin@charter.com

**Max Pala**

Principal Architect  
CableLabs  
858 Coal Creek Circle, Louisville, CO 80027  
303.661.3334  
m.pala@cablelabs.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Background on WWC .....	3
5G Authentication Framework.....	5
5G Authentication Functions .....	5
5G Authentication Framework .....	5
WWC Authentication.....	6
Trusted and Untrusted Access .....	6
Authentication in Wireline Access Networks .....	6
Authentication of W-AGF.....	7
Authentication of FN-RG .....	7
Authentication of 5G-UE.....	8
Authentication of Non-5G Capable Devices.....	8
Conclusion .....	9
Abbreviations.....	9
Bibliography & References .....	10

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Reference architecture for 5GC network for FN-RG connecting to W-5GAN (based on 3GPP TS 23.501).....	4
Figure 2 - Reference architecture for 5GC network for 5G-RG connecting to W-5GAN and 3GPP Access (based on 3GPP TS 23.501).....	4
Figure 3 - 5G Authentication Framework.....	6
Figure 4 - Architecture of WWC Authentication .....	6
Figure 5 - Registration of FN-RG to 5G Core .....	7
Figure 6 - Authentication Procedure of Non-5G Capable Devices in WWC.....	8

# Introduction

5G is one of the hottest technologies being trialed and deployed by network operators worldwide. Not only does 5G provide the superior services of fast speed, high bandwidth, and low latency, it also supports new use cases. One of these use cases is support for wireline and wireless convergence (WWC). In WWC, the 5G core manages both wireless access networks and wireline access networks (e.g., cable networks). This provides at least two benefits to residential network users. First, 5G user equipment with both cellular and Wi-Fi (WLAN) and/or wireline access can perform a seamless handover between cellular networks and residential networks. Second, residential user equipment without cellular access (e.g., a laptop or IoT devices at home) can also register to the 5G core to obtain services such as the Quality of Services (QoS) guarantee offered by 5G.

To enable WWC, authentication in cellular networks must evolve. More specifically, 5G authentication must allow the authentication of user equipment over wireline networks. This is in contrast to prior generations of cellular networks (e.g., 4G) which only allow authentication of subscribers over radio access networks. Further, 5G authentication must also allow user equipment without 3GPP credentials (e.g., a secret key stored in a UICC and shared with a network operator) to be authenticated by the 5G core. Prior generations of cellular networks authenticate only user equipment with 3GPP credentials.

In this paper, we provide a comprehensive analysis of 5G authentication that has been defined by 3GPP to support WWC. We include the 5G unified authentication framework which allows authentication to become agonistic to access networks and consistent between wireless and wireline networks. We also describe work-in-progress mechanisms that 3GPP is developing to authenticate non-3GPP-capable user equipment.

The rest of the paper is organized as follows. In Section 2, we provide background information on WWC. In Section 3, we introduce the 5G authentication framework which supports multiple authentication methods over multiple access types. In Section 4, we focus on the authentication of network elements in WWC. We conclude the paper in Section 5.

## Background on WWC

One of the objectives of the 3GPP 5G system architecture is to minimize dependencies between the access network and 5G Core (5GC) network in order to integrate different 3GPP access and non-3GPP access networks. Based on this objective, in 3GPP Release-15, the support for untrusted non-3GPP access in the 5GC network was introduced. In 3GPP Release 16, which is planned to be completed in March 2020, the support for trusted non-3GPP access and wireline access in the 5GC network will be introduced.


In 3GPP Release 16, the 5GC network supports connectivity to two types of Wireline 5G Access Networks (W-5GAN): Wireline 5G Broadband Access Network (W-5GBAN) and Wireline 5G Cable Access Network (W-5GCAN), which are specified by the Broadband Forum (BBF) and CableLabs® organizations, respectively. The 5GC network interfaces these wireline access networks via a gateway function called Wireline Access Gateway Function (W-AGF).

As depicted in Figure 1 [1] and Figure 2 [1], W-AGF provides N1, N2 and N3 interfaces towards the 5GC network. In the southbound direction W-AGF provides connectivity towards two types of Residential Gateways (RGs), called Fixed Network RG (FN-RG) and 5G-RG via Y4 and Y5 interfaces, respectively.


FN-RG is a legacy RG in existing wireline access networks that does not support N1 signalling and it is not 5GC capable. The FN-RG is defined by BBF and CableLabs organizations; and it can either be an FN-BRG or an FN-CRG depending it is part of W-5GBAN or W-5GCAN, respectively.

5G-RG is an RG capable of connecting to the 5GC network playing the role of user equipment (UE) with regard to the 5GC. It supports secure elements and exchanges N1 signalling with 5GC. The 5G-RG can either be a 5G-BRG or 5G-CRG depending it is part of a W-5GBAN or W-5GCAN, respectively.

As depicted in Figure 2, 5G-RG can be connected to 5GC via W-5GAN, NG-RAN or via both accesses. If 5G-RG connects to 5GC via NG-RAN, it is also called Fixed Wireless Access [1].



**Figure 1 - Reference architecture for 5GC network for FN-RG connecting to W-5GAN (based on 3GPP TS 23.501)**



**Figure 2 - Reference architecture for 5GC network for 5G-RG connecting to W-5GAN and 3GPP Access (based on 3GPP TS 23.501)**

In the reference architecture diagrams depicted in Figure 1 and Figure 2, FN-RG or 5G-RG can provide connectivity to devices behind them. Such devices could be 5G capable UEs or devices with no 5G capabilities, e.g. laptops, tablets, etc.

Detailed WWC architecture information can be found in 3GPP specifications [1] and [2].



# 5G Authentication Framework

In the following section, we first describe the network functions that are involved in the authentication framework and then discuss the authentication framework itself.

## 5G Authentication Functions

Service-based architecture (SBA) has been introduced for the 5G core network. Accordingly, new entities and new service requests have also been defined in 5G. Some of the new entities relevant to 5G authentication are listed below.

The Security Anchor Function (**SEAF**) is in a serving network and is a “middleman” during the authentication process between a UE and its home network. It can reject an authentication from the UE, but it relies on the UE’s home network to accept the authentication.

The Authentication Server Function (**AUSF**) is in a home network and performs authentication with a UE. It makes the decision on UE authentication, but it relies on backend service for computing the authentication data and keying materials when 5G-AKA or EAP-AKA’ is used.

Unified Data Management (**UDM**) stores subscription data associated with subscribers and selects an authentication method based on the subscriber identity and configured policy, but it relies upon Authentication Credential Repository and Processing Function (ARPF) where the shared keys are stored to compute the authentication data and keying materials for the AUSF if needed.

The Subscription Identifier De-concealing Function (SIDF) decrypts a Subscription Concealed Identifier (SUCI) to obtain its long-term identity, namely the Subscription Permanent Identifier (SUPI), e.g., the IMSI. In 5G, a subscriber’s long-term identity is always transmitted over radio interfaces in an encrypted form. More specifically, a public key-based encryption is used to protect the SUPI. Therefore, only the SIDF has access to the private key associated with a public key distributed to UEs for encrypting their SUPIs.

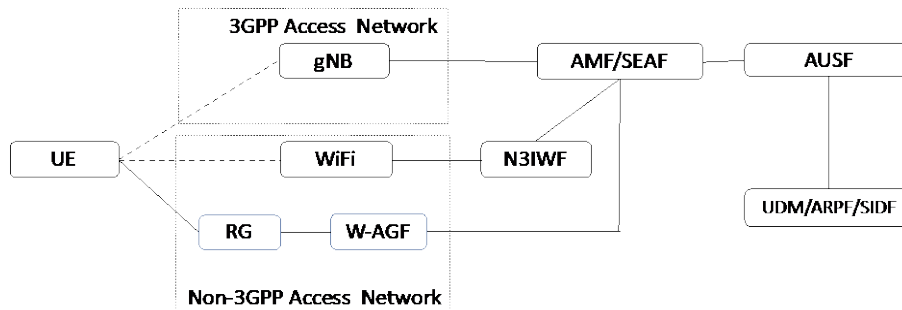
## 5G Authentication Framework

A unified authentication framework has been defined to make 5G authentication both open (e.g., with the support of EAP) and access-network agnostic (e.g., supporting both 3GPP access networks and non-3GPP access networks such as Wireless Local Area Network [WLAN] and wireline access networks) (see Figure 3).

When EAP (Extensible Authentication Protocol) is used (e.g., EAP-AKA’ or EAP-TLS), EAP authentication is between the UE (an EAP peer) and the AUSF (an EAP server) through the SEAF (functioning as an EAP pass-through authenticator).

When authentication is over untrusted non-3GPP access networks, a new entity, namely the Non-3GPP Interworking Function (N3IWF), is required to function as a VPN server to allow the UE to access the 5G core over untrusted, non-3GPP networks through IPsec (IP security) tunnels.

Several security contexts can be established with one authentication execution, allowing the UE to move from a 3GPP access network to a non-3GPP network without having to be reauthenticated.



**Figure 3 - 5G Authentication Framework**

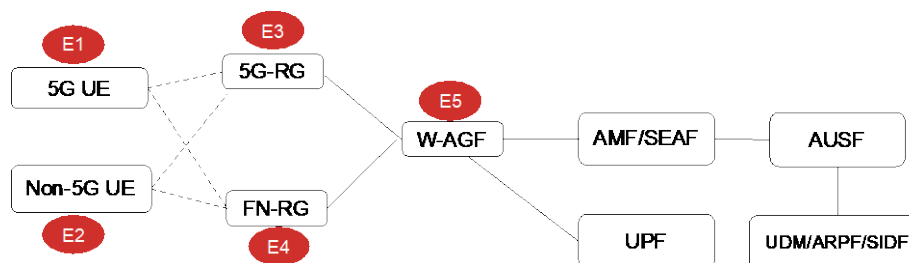
## WWC Authentication

### Trusted and Untrusted Access

Access networks in 5G systems can be either trusted or untrusted. For untrusted access networks, an IPsec tunnel is first established between a UE and a Non-3GPP Interworking Function (N3IWF). Afterwards, the UE can use one of the 5G primary authentication methods to authenticate to the 5G core. Non-5G UE will not be able to authenticate to the 5G core over the untrusted access network since it may not support the capabilities required to discover the N3IWF or to establish IPsec tunnel with N3IWF. For access networks considered as trusted, an IPsec tunnel may not be necessary, making it possible for non-5G UE to be authenticated to the 5G core, e.g., using an IETF EAP authentication method. This paper focuses on the authentication in wireline access networks with 5G core, which are considered as trusted access networks although not explicitly stated so in 3GPP specifications.

### Authentication in Wireline Access Networks

We consider the authentication of five network elements in 5G systems with wireline access networks. These five elements are 5G UE, non-5G UE, 5G-RG, FN-RG and W-AGF.



**Figure 4 - Architecture of WWC Authentication**

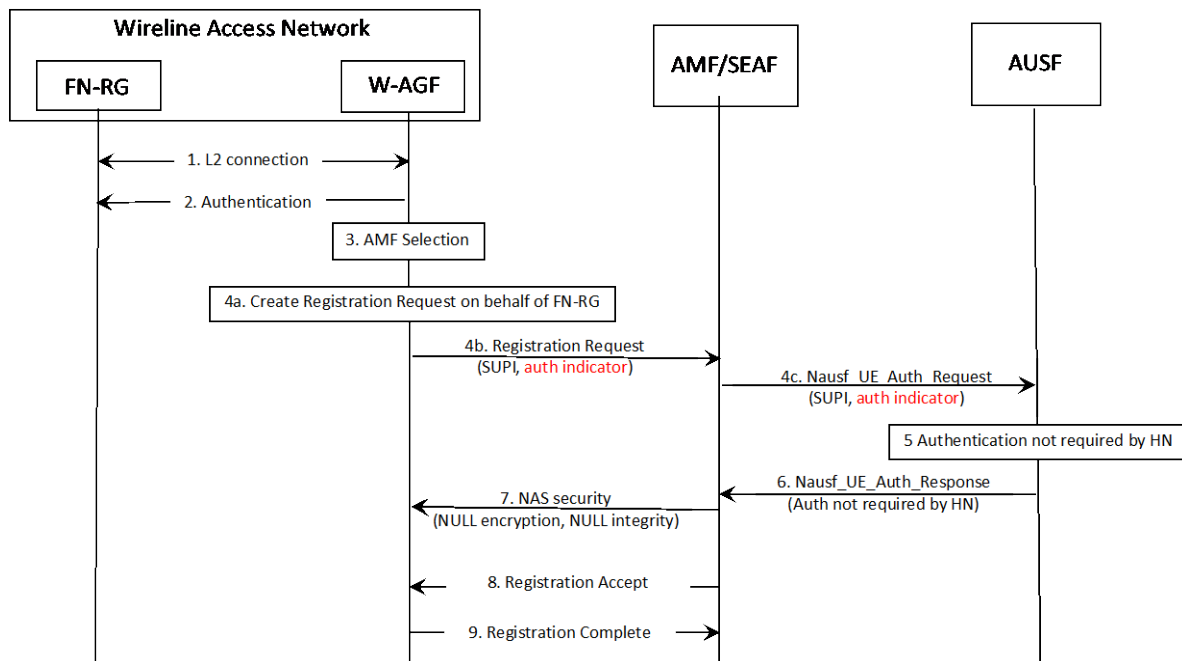
## Authentication of W-AGF

W-AGF (Wireline Access Gateway Function) is a gateway providing both signaling and user plane connectivity from the wireline access networks to 5G core. Since the network operator owns both the access networks and the 5G core networks, W-AGF is considered trusted and is authenticated by the 5G core network functions such as AMF and SEAF by establishing mutually authenticated TLS with the core. This requires W-AGF to be provisioned with server public key certificates for the authentication.

## Authentication of FN-RG

FN-RG is a legacy residential gateway which does not have any 5G capability nor does it interact directly with the 5G core. It is authenticated by the access network using the authentication method defined by either CableLabs or BBF.

After FN-RG is authenticated by the access network, it can be registered as 5G core via W-AGF using a transitive trust model. More specifically, W-AGF will indicate to the core that FN-RG has been authenticated in the registration message so that 5G core (i.e., AUSF) will skip the authentication of FN-RG during its registration (see Figure 5, based on TS 23.316 and TR 33.807).



**Figure 5 - Registration of FN-RG to 5G Core**

5G-RG is treated as a 5G UE and is authenticated by the 5G core using one of the primary authentication methods, e.g., 5G-AKA or EAP-AKA'. Since the wireline network does not support NAS capability, an EAP method (i.e., EAP-5G) is used to encapsulate NAS messages within EAP between 5G-RG and W-AGF.

Since 5G-RG should also support the authentication methods required by the access networks (e.g., BPI+), and will always be authenticated by the access network, we suggest that the transitive trust model assumed by the authentication of FN-RG may also be made applicable to 5G-RG. More specifically, 5G-RG can be authenticated by the wireline access network, and W-AGF indicates its authentication result to the 5G core without having to execute the authentication of 5G-RG by the 5G core directly. This would reduce the overhead of running the 5G authentication of 5G-RG, and make the authentication process consistent for both 5G-RG and FN-RG.

## Authentication of 5G-UE

5G-UE supports the primary authentication methods including 5G-AKA and EAP-AKA' and can use one of such methods to authenticate to the 5G core. If authenticated through wireline access networks, EAP-5G will be used to encapsulate NAS messages inside EAP frames.

## Authentication of Non-5G Capable Devices

There are large numbers of devices connecting to fixed access networks today which do not support 5G. It is desirable to allow those non-5G capable (N5GC) devices to also register to the 5G core so that they can be treated consistently as 5G UE.

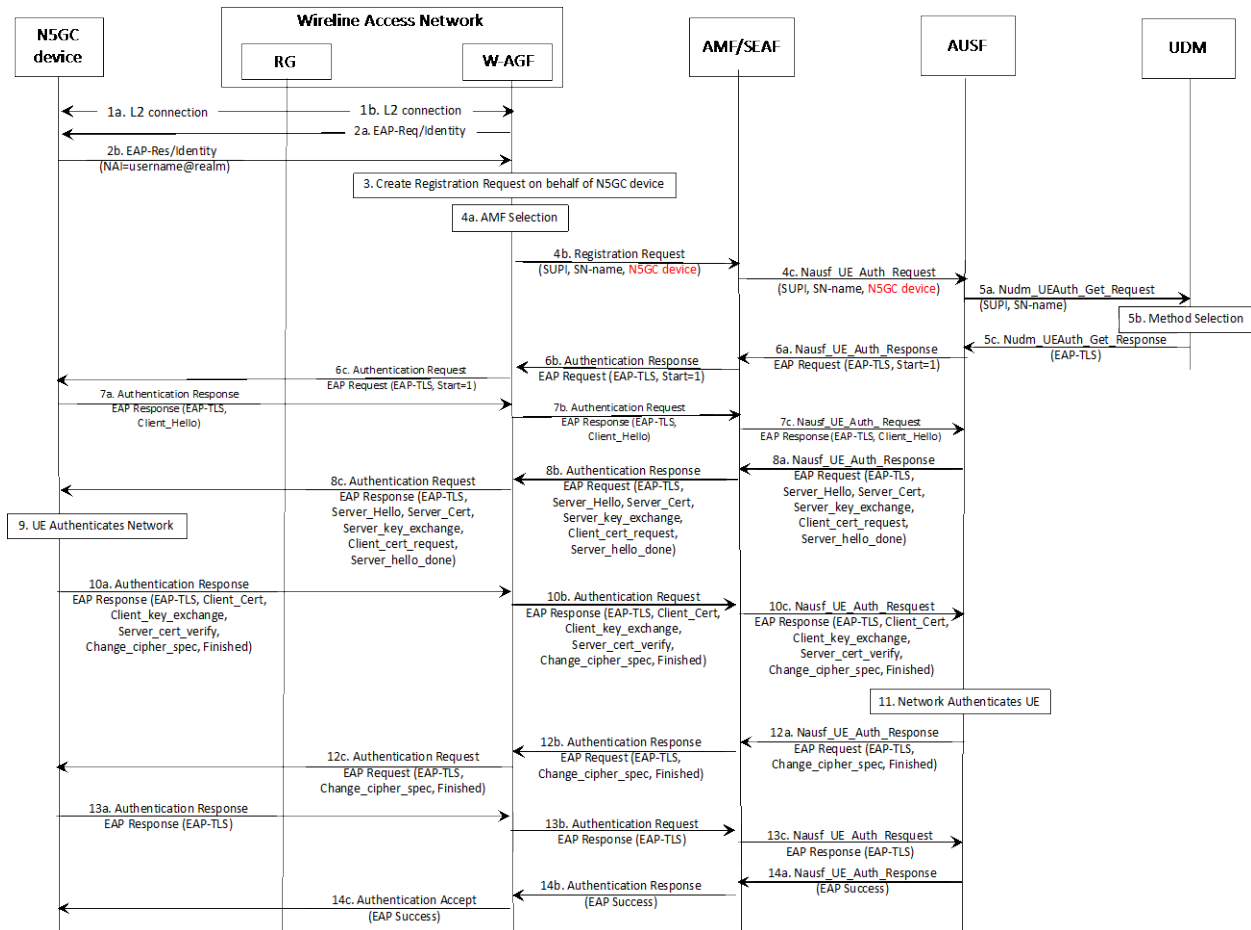


Figure 6 - Authentication Procedure of Non-5G Capable Devices in WWC

To do so, non-5G capable devices can use one of the EAP methods such as EAP-TLS to authenticate to the 5G core. EAP-TLS is defined in TS 33.501 Annex B, which, however, cannot be directly used to authenticate non-5G capable devices since it requires the EAP peer to be capable of receiving and processing 5G specific parameters such as ABBA and ngKSI, as well as performing 5G specific key derivation. We proposed a new procedure (see Figure 6) for authenticating non-5G capable devices, which is under discussion at 3GPP.

## Conclusion

5G standards are being actively developed, and one of the 5G use cases is to support convergence. In this paper, we discussed the authentication of each of the network elements in WWC. Due to the diversity of network elements and different trust models, WWC demands the support of different authentication methods and procedures, some of which have been agreed to in 3GPP and some have not. We hope this paper can serve the purpose of not only providing an overview of this subject, but facilitate the development of WWC authentication related specifications in 3GPP.

## Abbreviations

3GPP	3 <sup>rd</sup> Generation Partnership Project
5GC	5G Core
5G-RG	5GG Residential Gateway
5G-BRG	5G Broadband Residential Gateway
5G-CRG	5G Cable Residential Gateway
AKA protocol	Authentication and Key Agreement protocol
AMF	Access and Mobility Management Function
API	Application Program Interface
ARPF	Authentication Credential Repository and Processing Function
AUSF	Authentication Server Function
AUTH token	Authentication Token
AV	Authentication Vector
BBF	Broadband Forum
FN-BRG	Fixed Network Broadband RG
FB-CRG	Fixed Network Cable RG
FN-RG	Fixed Network RG
NG-RAN	Next Generation Radio Access Network
RG	Residential Gateway
SMF	Session Management Function
UPF	User Plane Function
W-5GAN	Wireline 5G Access Network
W-5GBAN	Wireline 5G Broadband Access Network
W-5GCAN	Wireline 5G Cable Access Network
W-AGF	Wireline Access Gateway Function
WLAN	Wireless Local Area Network
WWC	Wireless Wireline Convergence

# Bibliography & References

- [1] 3GPP TS 23.501, "System Architecture for 5G System", Release 16
- [2] 3GPP TS 23.316, "Wireless and wireline convergence access support for the 5G System (5GS)", Release 16
- [4] 3GPP, "Security Architecture and Procedures for 5G System" (Release 15), technical specification (TS) 33.501, v15.5.0 (September 2018).
- [5] Byeongdo Hong, Sangwook Bae, and Yongdae Kim, "GUTI Reallocation Demystified: Cellular Location Tracking with Changing Temporary Identifier," Proceedings of the Internet Society Symposium on Network and Distributed System Security (NDSS) (February 2018).
- [6] Internet Engineering Task Force, "Improved Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA)," Request for Comments (RFC) 5448 (May 2009).
- [7] Internet Engineering Task Force, "Extensible Authentication Protocol (EAP)," Request for Comments (RFC) 3748 (June 2004).
- [8] Internet Engineering Task Force, "The EAP-TLS Authentication Protocol," Request for Comments (RFC) 5216 (March 2008).

# **SD-WAN 2.0**

## **A Platform For Multi-Cloud, Security And Value Added Services**

A Technical Paper prepared for SCTE•ISBE by

**Charuhas Ghatge**

Senior Product and Solutions Marketing Manager  
Nuage Networks, a Nokia Company  
755 Ravendale Drive, Mountain View, CA94043  
(510) 299-2989  
Charuhas.ghatge@nokia.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
SD-WAN 2.0.....	3
1. The Challenges.....	3
1.1. Multi-Cloud Solution .....	3
1.2. End-to-End Security .....	3
1.3. Transforming the branch to a Value Added Services (VAS) Platform.....	4
2. The Solution – SD-WAN 2.0.....	4
2.1. Multi-Cloud Solution .....	4
2.2. End-to-End Security .....	5
2.2.1. Prevent .....	5
2.2.2. Detect .....	7
2.2.3. Respond .....	7
2.3. Transforming a branch to a Value Added Services platform .....	8
Conclusion .....	8
Abbreviations.....	8

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Solution For Public Cloud.....	4
Figure 2 - Adaptive Security Model .....	5
Figure 3 - Key New Requirements of SD-WAN: Branch-in-a-box.....	8



# Introduction

Enterprises worldwide have embraced the concept of SD-WAN, with leading telecommunications research firm IDC forecasting the worldwide market for SD-WAN to grow at a compound annual growth rate (CAGR) of 40% from \$830 million in 2017 to more than \$4.5 billion in 2022. Clearly, SD-WAN is here to stay.

SD-WAN 1.0 solutions have tried to solve the connectivity and automation challenges of branch offices, which have been underserved by the IP-VPN services. SD-WAN 1.0 also has been successful in reducing the bandwidth costs by offloading non-mission critical applications from the expensive MPLS to the cost-effective internet.

The new challenges for the enterprises are stemming from their quest and pursuit of Digital Transformation. The digital transformation almost mandates them to a multi-cloud strategy- from on-premises data centers to IaaS and SaaS public clouds and out directly to the branch offices and remote locations that constitute the intelligent edge.

With SD-WAN 2.0's reach from the branches to the DCs to the public clouds, security and governance become even more important as the attack surface increases, and hence the need for an end-to-end security model that is enterprise-wide: across hybrid-cloud, datacenter and branch network.

SD-WAN collapsed the functionality of a typical branch network, where many separate physical devices were needed to provide NAT, firewall, load balancers etc. into just one physical platform and many of these functionalities provided as VNFs. With the availability of this generic and powerful platform, the SD-WAN 2.0 must evolve this physical platform to deploy and manage Value Added Services including VoIP, Next Generation Firewall IoT and Wi-Fi access.

## SD-WAN 2.0

### 1. The Challenges

Enterprise IT needs are unmet in providing Multi-cloud solutions, end-to-end-security and Transforming a branch to value added Services platform. Let us look at these unmet needs and challenges.

#### 1.1. Multi-Cloud Solution

Hybrid Cloud has become the most popular approach to multi-cloud architectures. In Hybrid cloud you get the best of both worlds – the control and security of the private clouds and the flexibility and elasticity of the public clouds. Shortcomings in a wide area network (WAN) can exacerbate the complexity and management issue often associated with it. The network managers are often asking questions such as 'How can I easily move my workloads between public cloud providers and my branch or data center?' or 'How can I have fully redundant resilient connectivity using MPLS, Internet and LTE?'

#### 1.2. End-to-End Security

Increasingly, the threat landscape is getting more sophisticated with the rise of ransomware, web-based malware, botnets and phishing emails resulting in significant financial loss and data breaches. Malware like WannaCry ransomware that used lateral movement shows the importance of ensuring proper

segmentation both at the branch and datacenter to contain lateral spread and the need for a new analytics-based approach to detect and respond to these zero-day attacks.

Massive data breach at Equifax is a reminder to organizations on the importance of patching and quickly closing security vulnerabilities to secure key data before an attacker can use the security gaps to steal personal information. Organizations continue to get breached despite investments in security. Clearly there are gaps in organization's current security model for these attacks to happen.

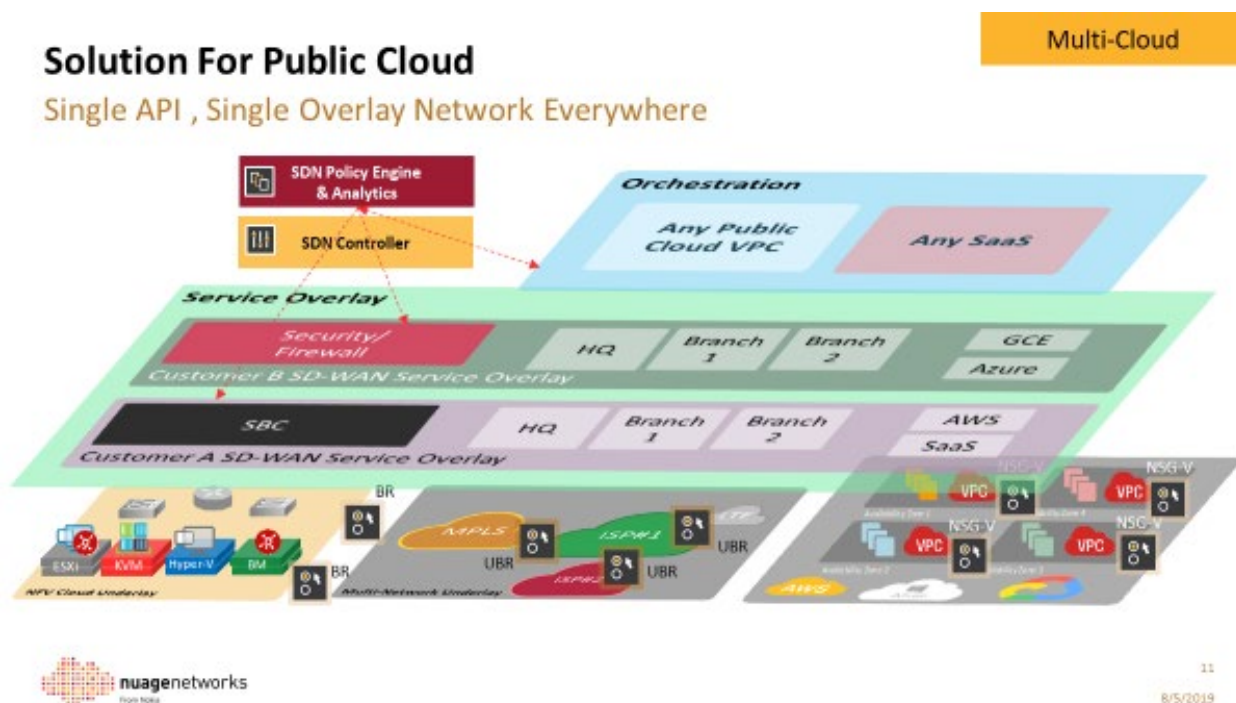
Current manual, perimeter-centric and reactive security model cannot effectively secure an organization data from emerging security threats in the cloud era.

### 1.3. Transforming the branch to a Value Added Services (VAS) Platform

The branch network of today is quite complex is comprises of many disparate physical devices offering spectrum of functionality and they include – NAT, DHCP, SBC/VoIP, security devices such as IPS/IDS, Firewall. They often are fragmented devices/appliances with rigid orchestration, they lack the flexibility to manage, configure and monitor and cost-prohibitive from the operational point of view. What is critically needed is a universal platform to ease the burden of manageability and flexibility to offer differentiated services at fraction of time and cost.

## 2. The Solution – SD-WAN 2.0

### 2.1. Multi-Cloud Solution



**Figure 1 - Solution For Public Cloud**

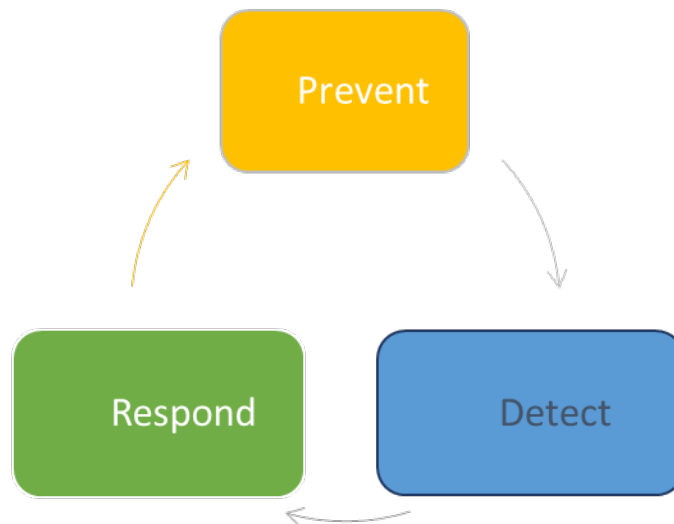
SD-WAN with its overlay feature, must provide the flexibility and choice of multiple underlay networks, whether those networks are MPLS or internet (from multiple ISPs). An ideal SD-WAN multi-cloud architecture should provide:

- Underlay
  - NFV DC underlay (private cloud underlay)
  - Multi-network underlay (MPLS, Internet- from multiple ISPs and LTE)
  - Public Cloud (AWS, Azure, GCP) Underlay
- A SD-WAN Services Overlay
  - Orchestration, Analytics and Policy Engine.

## 2.2. End-to-End Security

While microsegmentation provides significant benefits in terms of reducing the attack surface by limiting lateral movement of malware inside datacenter and cloud, organizations need a comprehensive security model that is enterprise-wide: across hybrid cloud, datacenter and branch network.

Gartner defined a new security approach called [Adaptive Security Architecture](#), one that is beyond traditional prevention and detection **and includes response based on continuous monitoring and analytics**. This adaptive security model suggests organizations to move from “incident response” mindset to a “continuous response” to defend against new wave of security threats.



**Figure 2 - Adaptive Security Model**

### 2.2.1. *Prevent*

This phase of the security implementation, as the name suggests, prevents various attacks at various points and layers in the network.

#### 2.2.1.1. *L3-L4 Stateful Distributed Firewall*

- Limit branch user access to/from internet as well as data center using L3-L4 stateful security
- Validated by 3rd party auditors for PCI-DSS 3.2 network firewall requirements

The Payment Card Industry (PCI) Data Security Standard (DSS) is one of the most wide-reaching standards since virtually every enterprise has individuals or organizations conducting transactions that accept, process or receive payments. Whether safeguarding payment information is an integral part of the core business – as in online retail and financial services – or an important aspect of the core business

(such as internal purchasing departments, consumer payments for services provided in the public and private sector), compliance with PCI DSS standards is essential.

- Logging of ACL actions for compliance and auditing

Logs of ACLs that are configured on the firewall can provide valuable information for auditing and subsequently for compliance. The ACL logs provide information about the packets that match an ACL such as source/destination IP, source/destination port and the protocol (TCP/UDP). The logging of ACL actions provides crucial information for auditing as well.

#### **2.2.1.2. Layer 7 Application Control**

- Restrict branch user access to specific applications using L7 DPI

With the ability to do Deep Packet Inspection (DPI) and to recognize 100s of application signatures, the administrator can define security policies that restrict or allow any of the recognized applications for any user.

#### **2.2.1.3. SaaS Application Control**

The SaaS services have become very prevalent in the enterprise IT environments. These SaaS service definitions should be recognized by the SD-WAN security engine such as Office365, WebEx, Salesforce, GitHub, JIRA, Azure, AWS and Google. These pre-defined SaaS services can be used in application identification, definition and ACL control.

#### **2.2.1.4. Web/URL Filtering**

- Blocks branch user/device access to malware as well as inappropriate content
- DNS based enforcement based on filtering DNS queries to internet sites

An internet site can be blocked or allowed, based on the DNS lookup. If an internet site is to be blocked, DNS query for that site is not returned to the requester. If it is allowed, then an IP address is returned to the requester.

- Supports content/website category based filtering (e.g., block malware, block adult content, block streaming media)

Websites categories are assigned to websites based on their dominant Web content. A website or webpage is categorized into a specific category that is likely to be blocked according to its content. When a website contains elements in different categories, web pages on the site are separately categorized. The categories are defined to be easily manageable and patterned to industry standards. There are many such categories and examples include adult content, dating sites, gambling or even can be categorized based on bandwidth consumption such Internet Radio and TV, Streaming sites, peer-to-peer file sharing etc. SD-WAN Security engine should support many of such website categories with over 32 billion URLs/domains that have been categorized.

- Supports filtering based on custom blacklist/whitelist of websites including wildcard matches

Custom blacklists and whitelists can be created and defined. Filtering of the traffic based on these blacklists and whitelists is supported.

- Supports logging of blocked websites/categories

Any blocked website or category is logged if trying to access that site/category.

### **2.2.2. Detect**

Continuous detection needs to be an ongoing part of security within a cloud-based environment. Typically, the new breed of attacks is more nuanced and sophisticated, unlike a typical Denial of Service (DOS) attack. These attacks could be zero-day attacks with no known signature and designed to permeate and infect laterally (east to west). Using flow analytics, traffic flows for each application need to be tracked throughout the application's lifecycle to anticipate potential threats.

#### **2.2.2.1. Contextual Flow Visualization**

The right solution should also leverage traffic insights from existing installed security measures. By correlating analytics from installed security measures with existing flow analytics (Flow Explorer), further contextual insight into the traffic and potential threats will be unlocked. For example, has this traffic attempted to breach any of the security controls that have been established and if so to what degree? Having access to this information will provide more context and will allow the enterprise to intelligently automate remediation policies.

#### **2.2.2.2. Top Talkers**

Top Talkers gives you topN source and destination pairs that are most talkative. A very handy functionality for identifying who is consuming most bandwidth and potentially attempted attacks.

#### **2.2.2.3. L3-L7 Traffic Visibility**

By setting thresholds and alerts based on those thresholds, allows one to proactively monitor and alerts and alarms if certain thresholds are crossed (rising as well as falling).

### **2.2.3. Respond**

#### **Automated Policy Action**

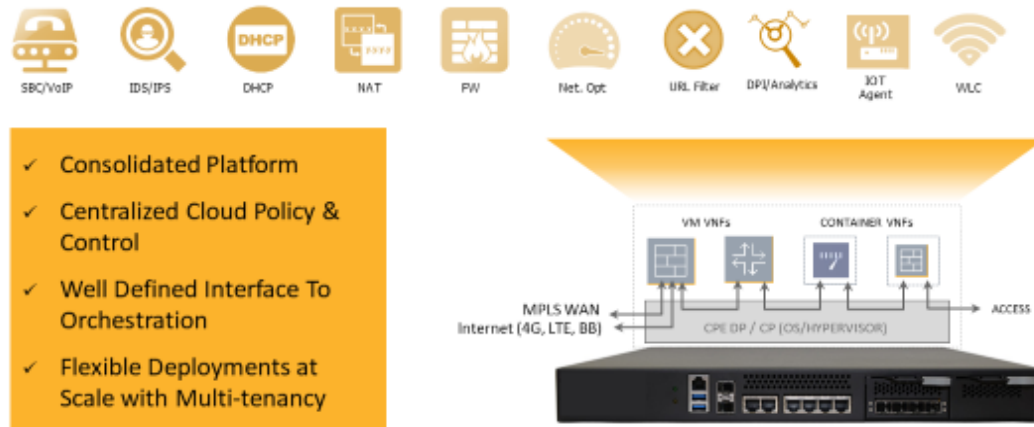
By having a dynamic understanding of the application traffic within the perimeter of the data center, branch, or public cloud, enterprises can then define and implement automated policies that can respond to certain suspicious application traffic flows in real-time. Some examples include:

- Real-time local alerts can be triggered informing the operator of suspicious activity for each application right down to the service-tier level of granularity. For example, an alert can be triggered when a certain TCP port on a virtual DB server is receiving an unexpected amount of ingress traffic
- Suspicious traffic can be steered to an existing SIEM to provide more correlation and analysis on this suspicious traffic flow, or to an IPS or L7 FW to sanitize the traffic
- Suspicious traffic can be quarantined or even blocked by steering traffic into a quarantined zone based on an automated trigger.

### 2.3. Transforming a branch to a Value Added Services platform

#### Key New Requirements Of SD-WAN: Branch-in-a-box

Network Applications



13  
8/5/2019

Figure 3 - Key New Requirements of SD-WAN: Branch-in-a-box

## Conclusion

SD-WAN has been a technology in the recent past that has succeeded and lived up to its hype because of the cost benefits, implementation flexibility and amalgamation of pragmatic technologies. However, with the advent of digital transformation and the relevant cloud transformation and security requirements, SD-WAN 2.0 is an evolution in offering multi-cloud integrations, pervasive security and transforming a branch as a Value Added Services (VAS) platform.

## Abbreviations

ACL	access control list
DNS	Domain Name System
DOS	denial of service
DPI	deep packet inspection
IoT	Internet of Things
IPS	Intrusion Projection System
LTE	Long Term Evolution
MPLS	Multiprotocol Label Switching
SaaS	software as a service
SDN	Software Defined Network
SD-WAN	Software Defined Wide Area Network
SIEM	Security Information and Event Management
URL	uniform resource locator
VoIP	Voice Over IP

# Scaling IP Advertising Using Manifest Manipulation

A Technical Paper prepared for SCTE•ISBE by

**Vipul Patel**

Vice President, Advanced Video Engineering  
Charter Communications  
8560 Upland Drive, Suite B, Englewood, Colorado 80112  
+1 833 267 6097  
Vipul.Patel@charter.com

**Xavier Denis**

Director, Product Management, Advertising Solutions  
CommScope  
1725 NW 167th Place, Beaverton, OR 97006  
+1 503 495 9485  
Xavier.Denis@CommScope.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
Content .....	4
1. Market shift .....	4
1.1. Transition to IP video .....	4
1.2. Impact on video advertising .....	4
2. Technology changes .....	5
2.1. MPEG TS streaming to ABR video .....	5
2.2. Advertising technology changes (splicer to manifest manipulation).....	5
3. Charter case study .....	6
3.1. A review of the IP journey to date and current operation .....	6
3.2. The 10M-session challenge – Taking IP video monetization to the next step .....	7
3.2.1. Capacity expansion and resiliency.....	7
3.2.2. Performance .....	8
3.2.3. Load distribution.....	8
3.2.4. Moving from an appliance model to a Data Center model .....	8
3.2.5. Player compatibility .....	9
3.2.6. Metrics and monitoring .....	9
4. Anatomy of manifest manipulation.....	9
5. The keys to designing for scale .....	12
5.1. Horizontal scaling .....	12
5.2. Performance and load distribution.....	12
5.3. Client Integration .....	12
5.4. Metrics and Monitoring .....	12
6. How did Charter's investment pay off? .....	13
6.1. Leveraging horizontal scaling.....	13
6.2. Load distribution .....	17
6.3. Appliance to Data Center transition: the importance of hardware tuning .....	19
6.4. Multi-ADS support .....	21
6.5. Client diversity and integration .....	22
6.6. Metrics and Monitoring .....	22
Conclusion .....	23
Abbreviations.....	25
Bibliography & References .....	26



## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Charter daily peak concurrent IP video sessions, by quarter.....	7
Figure 2 - Manifest manipulation.....	10
Figure 3 - Manifest lifecycle through delivery pipeline.....	11
Figure 4 - Horizontally scaling manifest manipulator design .....	14
Figure 5 - Distributed session state caching model.....	15
Figure 6 - Manifest requests on two load-sharing sets of manifest manipulators.....	16
Figure 7 - Peak streaming sessions on two load-sharing sets of manifest manipulators.....	17
Figure 8 - Load distribution model in a multi-region deployment .....	19
Figure 9 - Percent of client requests at different bitrate qualities.....	20
Figure 10 - Manifest manipulator routing to multiple Ad Decision Systems .....	21

# Introduction

Charter Communications has pioneered the development of streaming technology to augment their wide portfolio of advanced residential broadband services. One of the challenges Charter faced was to maintain a pristine customer entertainment experience while maximizing revenues through targeted advertising.

In order to effectively scale to millions of IP subscribers with billions of ad impressions, Charter undertook a major effort that included:

- Evaluating and selecting between Server-side and Client-side Ad Insertion
- Architecting a solution that can scale horizontally in a distributed environment
- Designing for high availability and resiliency
- Leveraging virtualization and cloud computing models
- Optimizing routing and load distribution
- Adapting to a diverse array of client platforms and Ad Decision Systems (ADSs)
- Enhancing the platform to support robust monitoring and accurate metrics and analytics

This case study provides an insight into the analysis, implementation, and best practices for scaling IP advertising in a demanding Pay TV environment.

## Content

### 1. Market shift

#### 1.1. Transition to IP video

Much has been written about changing consumer habits when it comes to video consumption. The rise of subscription Video on Demand (VOD) services, “cord-cutting”, short-form Internet video, Over-the-Top (OTT) linear Pay TV services, and others have all had a dramatic impact on the Pay TV industry.

What consumers may not realize is the underlying technology shift in how video is delivered. In order to leverage the Internet for video delivery and easily support multiple devices, video is increasingly delivered over IP as series of segmented files encoded in multiple bit rates. This process is commonly (and somewhat inaccurately) referred to as Adaptive Bitrate (ABR) streaming.

This new technology has not gone unnoticed by traditional Pay TV providers, many of which have long offered IP-based alternatives to their existing delivery models. Originally delivered as companion services such as “TV Everywhere”, operators are increasingly offering fully IP-based services targeting the main screen and other devices, both inside and outside of the home.

#### 1.2. Impact on video advertising

This shift is driving changes in the way that video advertising is delivered, sold and evaluated. Just as importantly, it is also creating several technical and operational challenges, some of which will be the focus of this paper.

While advertising in the IP realm requires new approaches from a technology perspective, it offers many potential advantages, including:

- Normally delivered as a unicast service, IP Video enables targeting ads for each individual stream, meaning that ads are more relevant to consumers and advertisers can target any desired audience segment
- IP Video makes each ad view measurable, which helps create new value metrics to monetize more content, including linear channels without traditional ratings
- Leveraging technologies like manifest manipulation, IP Video can interface to modern dynamic ad workflows for more efficient ad campaign execution and less ad fatigue thanks to frequency capping

While the definition of “TV” may be blurring, traditional linear television advertising was still a massive \$195B business in 2018 and will grow to \$210B by 2023 [1]. The addition of IP technology with its increased targeting capabilities to this powerful platform brings tremendous opportunity. For example, studies have found that ad completion rates on IP-connected TVs were 95%, 32% higher than mobile, and 27% more than desktop [3].

## **2. Technology changes**

### **2.1. MPEG TS streaming to ABR video**

MPEG Transport Stream broadcast delivery over QAM has been the way to reliably deliver video to large audiences for decades. But the advent of streaming technology over IP coupled with intense competition and the large market for streaming devices have disrupted that paradigm. Adaptive Bitrate (ABR) protocols like HTTP Live Streaming (HLS) and Dynamic Adaptive Streaming over HTTP (DASH) have matured to enable robust and scalable delivery while standards efforts have improved the economics of building large streaming systems leveraging operators’ IP networks. New standards like Common Media Application Format (CMAF) are aiming to foster more innovation in the industry by promoting common encoding and encryption formats, and by reducing latency in live streaming to parity with broadcast delivery.

### **2.2. Advertising technology changes (splicer to manifest manipulation)**

In the realm of ABR video streaming, manifest manipulation is the rough equivalent of splicing ads into a broadcast Transport Stream delivered over UDP transport. The differences between the two stem from technological advances and the different problem spaces.

Transport Stream delivery needed to cope with rate clamping in order to pass ads without exceeding the constraints of statistically-multiplexed services to “fit” the largest number of services in a 6-MHz channel. ABR delivery deals with matching content and ad representations across a range of bitrate ladders, formats and codecs, and supporting a large array of streaming devices that often behave inconsistently.

The two approaches to the problem of ad insertion in ABR streams are manifest manipulation (sometimes referred to as Server-Side Ad Insertion because the function is deployed in the operator network) and Client-Side Ad Insertion. Manifest manipulation provides a seamless way to replace or insert video chunks in ABR streams by altering the manifests of those streams independent of the client-side device and player technology. Client-Side Ad Insertion involves implementing this logic in the device itself.

While Client-Side Ad Insertion involves the complex logic of chunk and buffer management, integration with Ad Decision Systems, and other backend ad workflows, manifest manipulation is client-agnostic, and avoids duplicating backend integration costs.

These advantages help explain why manifest manipulation is gaining broad industry adoption for ad insertion and a growing number of applications [4].

### **3. Charter case study**

#### **3.1. A review of the IP journey to date and current operation**

Charter is one of the largest providers of television, broadband internet, and voice services in the world. Known for its commitment to integrating the highest quality service with uniquely differentiated products, Charter is at the intersection of technology and entertainment, facilitating essential communications that connect more than 28 million residential and business customers in 41 states.

Building a scalable IP video network and service is core to Charter's vision to delight its customers with the greatest choice of content and a superior experience available on any device.

To understand how to approach building a scalable IP video service, it is important to step back and see where it all started. Charter initially marketed its add-on content streaming capability as an augmentation to its set-top box (STB) delivered television product. To that end, the strategy was to focus on the building blocks of the system.

When Charter launched its IP video service in 2015, ABR technology was at various stages of maturation. While HLS was scaling remarkably well on the back of strong market demand for the iPhone®, the MPEG-DASH standard was taking longer to yield robust and mature client implementations that could be deployed at scale. Charter used Microsoft® Smooth Streaming in order to support critical platforms like the Xbox® and Samsung® smart TVs.


From a monetization perspective, Charter chose manifest manipulation because it would:

- Ease support of a constantly expanding set of IP-connected devices such as smart TVs, OTT streaming devices, game consoles, PCs, tablets and phones
- Decouple integration with backend systems like Ad Decision Systems from client integrations
- Provide more control over the quality of the user streaming experience by ensuring clean ad insertions and avoiding rebuffering

The HLS streaming protocol inherently lends itself to manifest manipulation for supporting Dynamic Ad Insertion (DAI) and alternate content switching. For the Smooth Streaming protocol, a simple manifest manipulation is not sufficient as the media segments themselves need to be altered. Alternatively, DAI for Smooth Streaming devices is typically implemented using Client-Side Ad insertion, which poses its own set of challenges related to Ad Decision System integration and maintaining media stream buffers. In the near future, the Smooth Streaming devices can switch to the DASH streaming format, which also supports DAI and alternate content switching through server-side manifest manipulation.

Monetization is only possible with robust reporting. Charter initially experimented with server-side ad beaconing during manifest creation but realized that ad tracking events would often be mis-aligned with the occurrence of the playback events. They determined that building a common cross-client ad beaconing framework allowed them to better conform with industry standards for ad metrics collection.

Over the past two years, Charter has seen session usage grow 4-fold (see below diagram), underscoring the reality that as Charter prepared to evolve its offering to target the primary TV screen in the home, it needed to ensure its platform would continue to scale.



**Figure 1 - Charter daily peak concurrent IP video sessions, by quarter**

### 3.2. The 10M-session challenge – Taking IP video monetization to the next step

After the initial successful launch, which extended the primary video services to subscribers on IP-connected devices, Charter focused on updating the platform architecture so that it was scalable and resilient to support IP-based video as the primary service. In doing so, the company was confronted with several challenges.

- Capacity expansion and resiliency
- Performance
- Load distribution
- Moving from an appliance model to a Data Center model
- Player compatibility
- Metrics and monitoring

#### 3.2.1. Capacity expansion and resiliency

As the engine of monetization for IP Video, manifest manipulation was high among Charter's priorities. Charter decided to launch the service with a centralized manifest manipulator design which was a good choice, because it reduced deployment risk, provided for a better ability to contain infrastructure costs, and eased operations.

Since manifests are personalized, they are by nature not cacheable by the Content Delivery Network (CDN), and manifest traffic needs to traverse the network from the manifest manipulator to the clients. As service usage increased, the upward trend in network traffic prompted questions about how to more effectively manage manifest session load at scale. Some have proposed moving manifest manipulation closer to the edge of the network to alleviate this problem [4].

Overall service resiliency, another major factor for Charter, is heavily impacted by where the system is deployed and the type of redundancy it offers. In that regard, a distributed approach provides the benefit of spreading the load across more resources, and isolating the negative impact of any outage. The right solution, however, should ensure the appropriate amount of redundancy and fault-tolerance to minimize revenue impact through fail-over capabilities.

Overall, Charter needed a more flexible way to grow capacity, either centrally or at the edge of the network, with the ability to maintain resilient services.

### **3.2.2. Performance**

Evidence suggests a clear correlation between viewer engagement and video quality issues such as shifting bitrates and video stuttering. Even small playback delays can lead significant portions of a stream's audience to tune out [2]. The ability to monetize IP video and control alternate content events depends in no small part on the manifest manipulator's robustness.

The issues are complex in part because of the different characteristics of live event streaming, scheduled linear television streaming, and VOD streaming experiences. Live event and linear television streaming are highly transactional, with high volumes of simultaneous sessions, and more susceptibility to variations in response times and latency. While VOD is more storage intensive, it imposes less stringent requirements for near real-time responsiveness.

The manifest manipulator is part of an elaborate delivery pipeline of interconnected components (Packager, CDN, and clients). Conditions as varied as packet loss or congestion, clients struggling with slow connections (such as weak Wi-Fi/wireless signals), packager failures, CDN responsiveness, and CDN misconfiguration cause severe quality issues that can be compounded by sensitive clients. Any mechanisms provided by the system to absorb the effect of these conditions will have a positive impact on performance.

### **3.2.3. Load distribution**

Load distribution concerns the imperative of maximizing the use of available resources by evenly distributing sessions. Since we are talking about video, it is clear that we want to also avoid the magnifying impact of requests jumping between different manifest manipulation servers throughout the same session. We will cover the tools and best practices to consider when designing a load distribution model that can scale.

### **3.2.4. Moving from an appliance model to a Data Center model**

As the core video processing functions move to pure software, the Data Center model offers a natural path to scale. While moving away from dedicated hardware toward a common virtualized infrastructure that can be leveraged across applications is a sound approach, advertising for Pay TV video is a specific application with unique requirements. The key challenge is adapting the resource allocation policies to the application to support a constant low latency response to every request on each video stream. Resource sharing among Virtual Machines (VMs) may be possible with certain applications. However, manifest

manipulation is a real-time mission critical application that is heavily dependent on compute and network availability.

### **3.2.5. *Player compatibility***

Charter is very familiar with the challenge of supporting a vast array of IP streaming devices. The problems range from player device compatibility and stability, to varying player behaviors. Some players start at the lowest profile and ramp up, while others start at a middle bitrate. Some can handle chunks of different durations. Differences also exist in how players calculate buffering requirements in order to determine what bitrate to request for future chunks. These are part of a broad array of issues addressed in [4].

When Charter launched its service on the Apple TV® platform in early 2019, it already had many client integrations under its belt, including Roku®, iOS, and Android devices. Those earlier efforts greatly helped reduce risk and time to deployment for the support of DAI and alternate content switching on this new platform.

This serves to underscore another aspect of delivering IP Video at scale. Given the inevitable proliferation of new devices and client technologies, investing in the testing and integration capabilities to qualify these new platforms plays a key role in their successful launch and support.

### **3.2.6. *Metrics and monitoring***

This is a broad topic that covers many aspects of the video service. Addressing these areas impacts revenue directly (failed ad insertions causing ad inventory waste) and indirectly (a bad user experience leading to higher churn rates over time). Some of the relevant questions for Charter included:

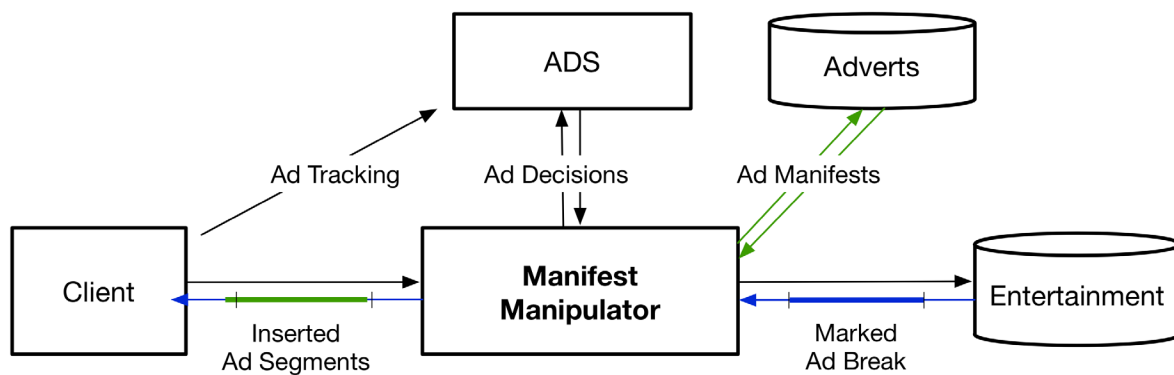
- What are the critical metrics to measure overall video quality and ad performance?
- What tools are needed to effectively troubleshoot at scale?
- Where and how does the data get acquired?

Evidence shows that viewer engagement with ads is heavily dependent on the quality of the streaming experience. It is critical to be able to troubleshoot ad failure points (missing ads, invalid manifests, empty ad responses from ADSs, etc.) and correlate those events to their impact on revenue.

## **4. Anatomy of manifest manipulation**

Live ABR streams (HLS and DASH) generated through a transcoder and packager system are made available to clients through a manifest. The manifest updates on a web server and includes a window of media segments available to be played by all downstream clients. Manifest manipulation can personalize these streams by presenting each client the entertainment content while also altering media segments so the client will also play ads, alternate content, emergency alerts, etc.

Manifest manipulation uses features that exist in both the HLS and MPEG-DASH standards to insert or replace the media segments presented to each client. This may be triggered by markers or event tags in upstream manifests indicating the precise splice point or cue for ad breaks (or alternate content). Inserting ads or alternate content is achieved through defined periods or discontinuity markers in manifests that cause the client to prepare the decoding of the new media for seamless rendering across the transitions.



Input manifests:

E1: Entertainment					E2: Ad Break 90s												E3: Entertainment			
	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	

Output manifests:

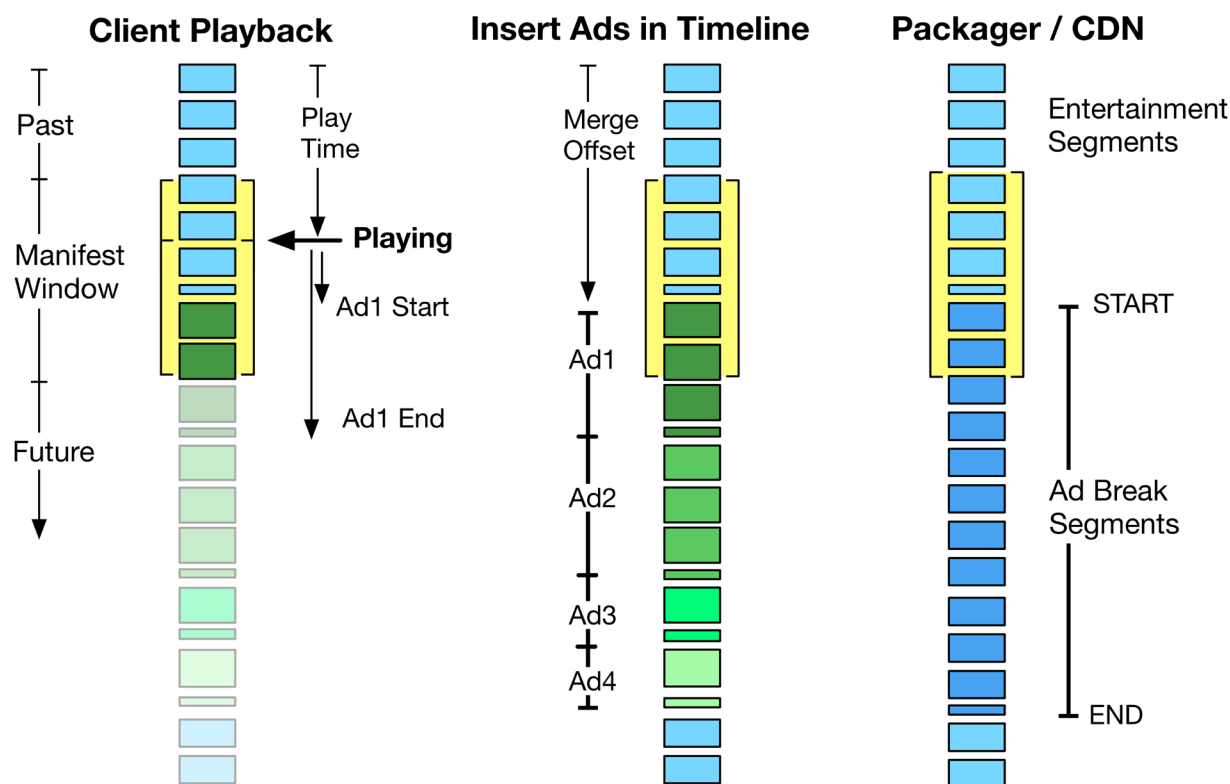
E1: Entertainment					Ad1 30s				Ad2 30s				Ad3 15s		Ad4 15s		E3: Entertainment			
	2	3	4	5	1.1	1.2	1.3	1.4	2.1	2.2	2.3	2.4	3.1	3.2	4.1	4.2	17	18	19	20

**Figure 2 - Manifest manipulation**

Each client makes requests to update its manifest asynchronously to when the media segments are first published by a packager. A client typically starts playback rendering a few segments behind the most recent segment to allow tolerance for bandwidth fluctuations and gain the ability to monitor buffer depth in order to adapt media quality to a bitrate that can be sustained by current bandwidth rates.

For manifest manipulation to insert content into a continuous client timeline, it must maintain an appropriate history of past segments in the manifest window while appending new media segments to that timeline. The application must track the rate of the upstream entertainment media in order to output that same rate of inserted media. In some cases, an upstream ad break may signal an early abort, which needs to immediately return to the correct entertainment segment. Handling such a condition involves negotiating the trade-off between staying close to the input manifest's edge and time-accuracy of return to entertainment.





**Figure 3 - Manifest lifecycle through delivery pipeline**

The manifest manipulator typically executes a workflow defined for a type of session based on attributes of the client and parameters from an initial playback URL. This workflow determines routing policies and settings for Ad Decision Systems (ADS) or Alternate Content Decision Services (ACDS). As each session has signals that determine need for decisions, the initial session attributes can be passed as targeting criteria for these decisions and campaigns. The response from these services determines the content URLs that will be inserted into that session.

A session state object typically caches the session context (ad decisions, alternate content URL, etc.) within the current manifest manipulator instance. That instance can most efficiently build upon the prior output manifests, continue segments from recent placement decision, and maintain tracking of the session variations to deliver seamless playback. The session state also provides opportunity to report on analytics indicating the client targeting attributes, policy workflows, ad opportunities, ad placements, and quality metrics such as shifting quality levels as reported in Figure 2 above.

The manifest manipulation service maintains resiliency of the above session state so, in the rare case the active instance for a session fails or is unreachable, the session can be redistributed and recovered on a neighboring instance.

## **5. The keys to designing for scale**

In response to the challenges described in Section 3.2, Charter identified the following lessons learned in several areas of focus. Section 6 elaborates on each of these topics.

### **5.1. Horizontal scaling**

- A. Scale out using an atomic work engine instance that can be tiled outward with no practical limit
- B. Preserve flexibility between centralized and distributed deployment architectures
- C. Leverage orchestration tools to ease expansions and contractions to meet dynamic loads
- D. Be pragmatic in designing for resiliency and service availability

### **5.2. Performance and load distribution**

- A. The Dos and Don'ts of virtualization and infrastructure planning
- B. Devise a multi-prong strategy to manage load on the network

### **5.3. Client Integration**

- A. Engineering manifest manipulation to adapt to nuanced client behaviors
- B. Perform testing, testing, and more testing

### **5.4. Metrics and Monitoring**

- A. Define the key telemetry data points and strategy to scale data aggregation and storage
- B. Develop the right monitoring tools to find the needle in the haystack

## 6. How did Charter's investment pay off?

For Charter the key design criteria to build the Next Generation IP video system were:

- Optimize traffic load at the edge of the network to minimize client latency
- Ease the ability to dynamically adapt to fluctuating workloads
- Optimize load distribution across all instances
- Build resiliency to enable lights-out, 24x7 service
- Design for geo-redundancy and 100% service availability
- Enable zero-downtime software upgrades and maintenance

While these criteria applied to several components of the video delivery pipeline, this section focuses on how Charter guided the redesign of the manifest manipulation function and the operational impact that resulted from the changes.

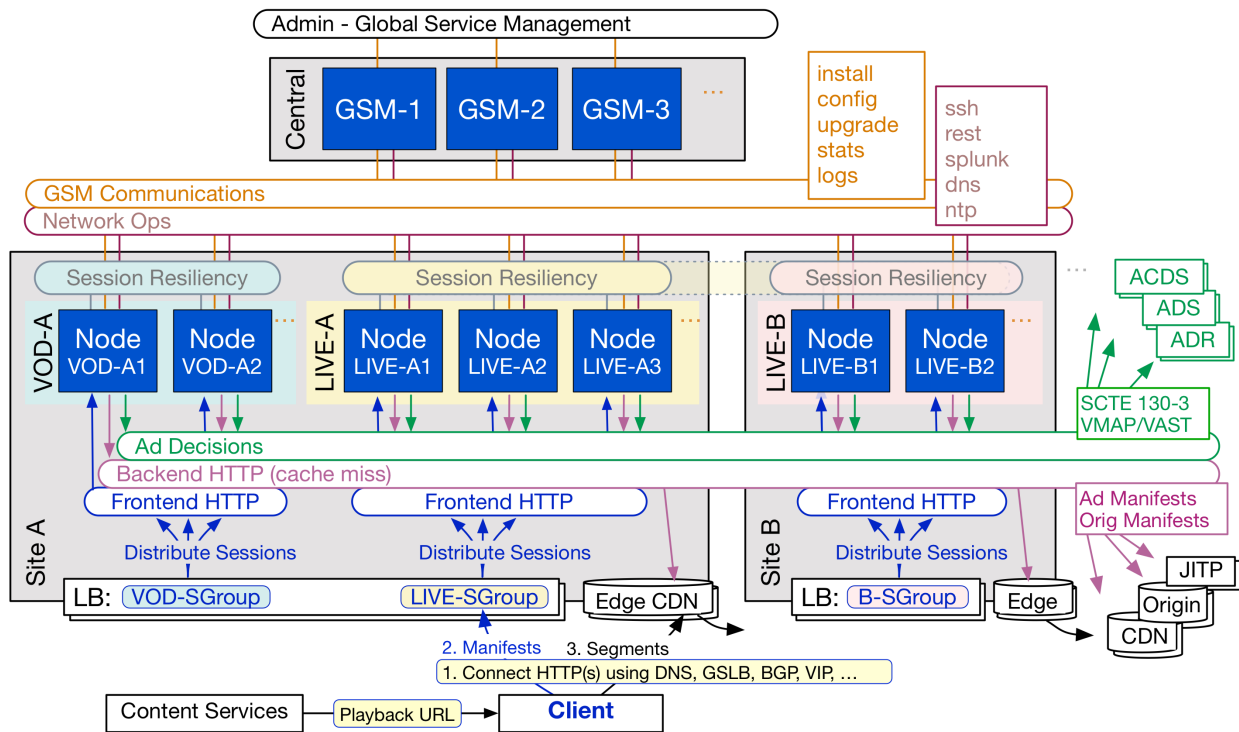
### 6.1. Leveraging horizontal scaling

The goal was to flexibly accommodate any peak load of VOD and live IP-connected client sessions by easily scaling out the number of work engine nodes in the system without any limitations while preserving full session resiliency and service availability. A second goal was to be able to continually roll out new features via software upgrades without disrupting revenue-generating services supported by the application.

That means we require the ability to effectively manage a large collection of work engines (or “nodes”) that operate semi-independently of each other so that it is possible, to:

- Selectively push new software to any subset of nodes
- Turn on or off new nodes on any infrastructure, whether hosted or cloud-based
- Provide session resiliency across instances within a Data Center or across regional ata Centers

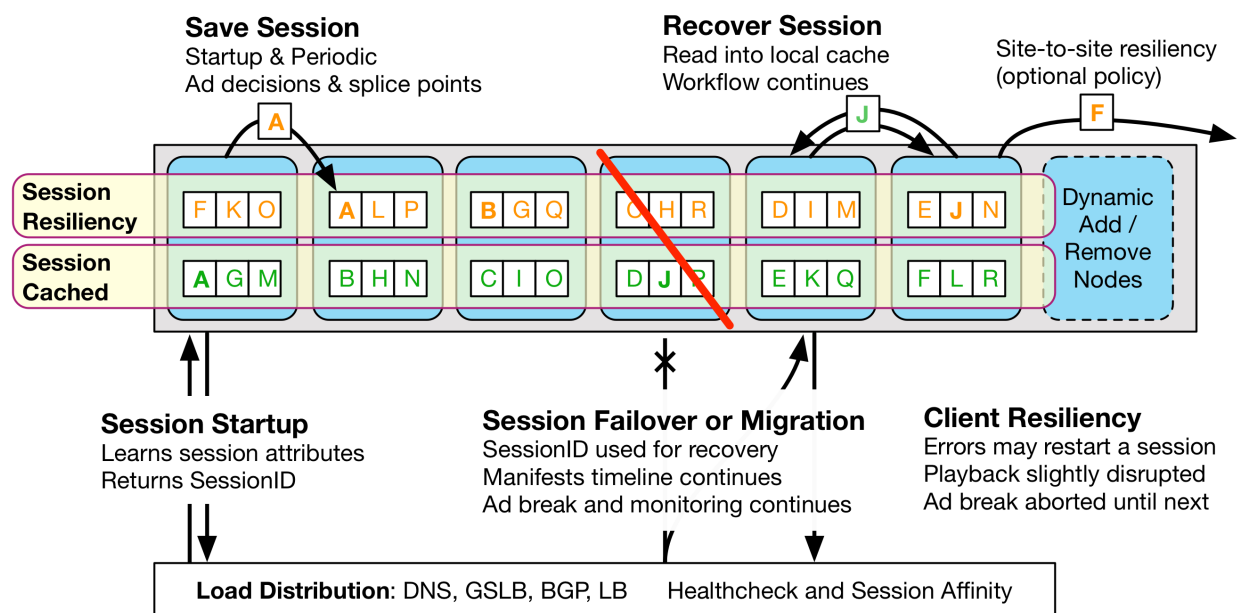
The preferred design was, and continues to be, one where session distribution is built in the node, without reliance on a central database (see figure below).



**Figure 4 - Horizontally scaling manifest manipulator design**

The manifest manipulator operates behind the load balancer (LB), which can support distribution using session affinity such that the node designated at session start keeps getting the requests under normal operations. This ensures maximum efficiency by preserving processing on the node that is in charge of maintaining session state. Session state is also written in a distributed data store that enables “neighboring” nodes to handle requests seamlessly, should the session need to move because of hardware failure, or any cause making the node unresponsive to the load balancer (see Figure 5).

This level of resiliency can be achieved among any number of nodes clustered together as a single set of neighbors within the same site. And it can be extended to support geo-redundancy, wherein two active sets of nodes installed at separate physical locations are sized such that either can handle the peak load as needed, yet load is evenly distributed across both under normal conditions (see Figure 4 illustrating manifest manipulation nodes deployed at sites A and B). The question of which session state resiliency policy is implemented depends on defining the service SLA and evaluating the backend network costs required to meet it.



**Figure 5 - Distributed session state caching model**

With service availability, a device should be able to retain access to the service if one of the sites becomes unresponsive. This can be achieved without requiring session state replication across node sets, thereby saving network costs.

With service resiliency, a session should continue uninterrupted should one site become impaired or unresponsive, and the load be redirected to the other site. This can be achieved by setting replication policies ensuring each session state will be distributed to both node sets. However, this level of resiliency requires the appropriate network bandwidth to be provisioned between sites.

In the horizontally scaling design favored by Charter and depicted in Figure 4, session state replication can be thought of as a multi-level recovery mechanism that optimizes resiliency, service availability and performance:

- L1 – The individual node internally manages active sessions
- L2 – Automatic updates to neighboring nodes in the same site facilitate distributed resiliency within the site
- L3 – Distributed resiliency to nodes in a separate site can be optionally configured


The session state replication mechanisms are self-contained within the node logic, which removes any dependencies and bottleneck issues on a would-be central session state manager. In fact, session replication is independent of the management layer used to provision services on the nodes. Furthermore, a session is not dependent on other nodes, as replication is updated in the background without locking any real-time processing service. Replication is updated at decision points (ad breaks, bit rate switches in live HLS streams, etc.) and mid-window to maintain timeline consistency.

This design enables linear scaling of resources relative to the workload without any limitations. Let  $S$  be the number of sessions each node can handle, and  $N$  the number of nodes. Then:


- The memory in each node holds  $2*S$  sessions, i.e.  $S+(N-1)*S/(N-1)$
- The network interface at each node handles traffic for replica transmission of  $S$  sessions and for ingest of replicas for  $(N-1)*S/(N-1)$  sessions, totaling  $2*S$  sessions
- The network interconnect backbone needs to handle the total load of a site's replication traffic, or  $S*N$
- The option to configure extra cross-site replication doubles traffic requirements, i.e.  $2*S*N$

The two graphs below show the activity recorded during a real-life test that was performed on a 150K-session load served by two sets of load-sharing manifest manipulator nodes, with each set provisioned to handle the full load. At 16:37, one of the two systems went down, causing the sessions to seamlessly move to the other systems without interruption.

Observe that the volume of manifest requests (Figure 6) and peak sessions (Figure 7) went up on the surviving system in the same proportion as the activity went down on the other system. Soon after the down system was put back in service, this action restored an evenly distributed load across systems. At 16:52, the same procedure was repeated by reversing the roles of the systems. Throughout this process, we also measured non-HTTP 200 messages and logged none, further evidence that the clients were unperturbed throughout the whole test.



**Figure 6 - Manifest requests on two load-sharing sets of manifest manipulators**



**Figure 7 - Peak streaming sessions on two load-sharing sets of manifest manipulators**

The scaling model presented in this section meets all the requirements of a modern software application that must support workload fluctuations, service availability and resiliency, and the ability to maintain/upgrade the system with agility and reliability.

This architecture provides Charter more flexibility in growing capacity either centrally for more efficient use of compute resources or at the edge, thereby reducing load on the network and optimizing client response times. As is often the case in solving complex problems, this is not an either or proposition, which makes the flexibility of horizontal scaling very attractive.

## 6.2. Load distribution

Coming up with a well-designed, load balancing system to address the specific needs of ABR video relies on a combination of widely used best practices and optimization techniques.

Equal load distribution is the goal. A new session can be handled at any acceptable node (note: different pools of nodes may be configured to handle different services, such as VOD vs live).

Load balancing should favor the “stickiness” of a session to the resource initially assigned to serve it. That applies to the client requests for manifest updates throughout the session, as well as to client requests for callback and tracking signaling (end-of-session, tracking events when proxied by the manifest manipulator). All this is most easily achieved using a load balancer that distributes load based on client IP address.

Client-IP-address-based distribution has many benefits, including:

1. Greater cache efficiency and lower latency for the session on-going updates (avoids the delay/cost of a session restore)
2. Ability to leverage caching of prior backend manifests being actively spliced into the output
3. Improved diagnostics for a client support issue, because it obviates the need to scan for IP addresses across the system
4. Protection of the client from any issues elsewhere in system (fewer system dependencies)

Client identifiers such as a key or hash can also be used. This method better supports the scenario of clients whose IP address may change during a session, as is the case with mobile clients that may switch between Wi-Fi and cellular networks, as well as clients behind NAT, DHCP, or VPNs.

Load may also be directed by systems, specifically:

- Content Services – Client-provided playback URL with hostname of a site (household region or other attributes)
- Resolution Services – Regional systems such as DNS/GSLB, DHCP, session router/gateway, select the site
- Network – Route to the site with the lowest cost border gateway protocol (BGP), fewest hops, lowest latency, peering bandwidth costs
- Balance – Spread sessions across servers or sites and adapt to available capacity, response times, and resiliency

Or by partitioning services, such as:

- Content Type – Directed to a specific cluster, access point, workflow (premium channels, VOD, PPV)
- Account Attribute – Class of service or assigned region for the subscriber/household/user info
- Device Tag – Provision or assign an endpoint to a regional access point or class of service

Load balancing factors as a key dependency if capacity expansions are realized either by proactive provisioning or elastic expansions. Proactive provisioning is the idea of adding servers or sites for anticipated larger loads without the need to reclaim the resources – which is appropriate when the use of the extra servers can be justified as an investment in anticipation of more customer acquisitions or organic usage growth. Elastic expansion is about quickly deploying more servers or sites for punctual needs, knowing the extra capacity will be reclaimed after the large-volume streaming event.

Finally, the load balancer plays a key role in system resiliency at many levels:

- Manifest manipulator node – The load balancer monitors health check for status / capacity to distribute sessions across nodes
- IP Fail-over – The load balancer's heartbeat fail-over (VIP) or site route health inject (RHI / BGP) mechanism can quickly reroute traffic
- Overflow – A safety valve mechanism that helps protect users' streaming experience by automatically bypassing the manifest manipulator when nodes are saturated or the manifest manipulator is inaccessible
- Network – Networking issues from client (home, service-group) and into the deployed site
- Client – Recovery attempts for retries, reconnects, restarts, user behavior, network roaming

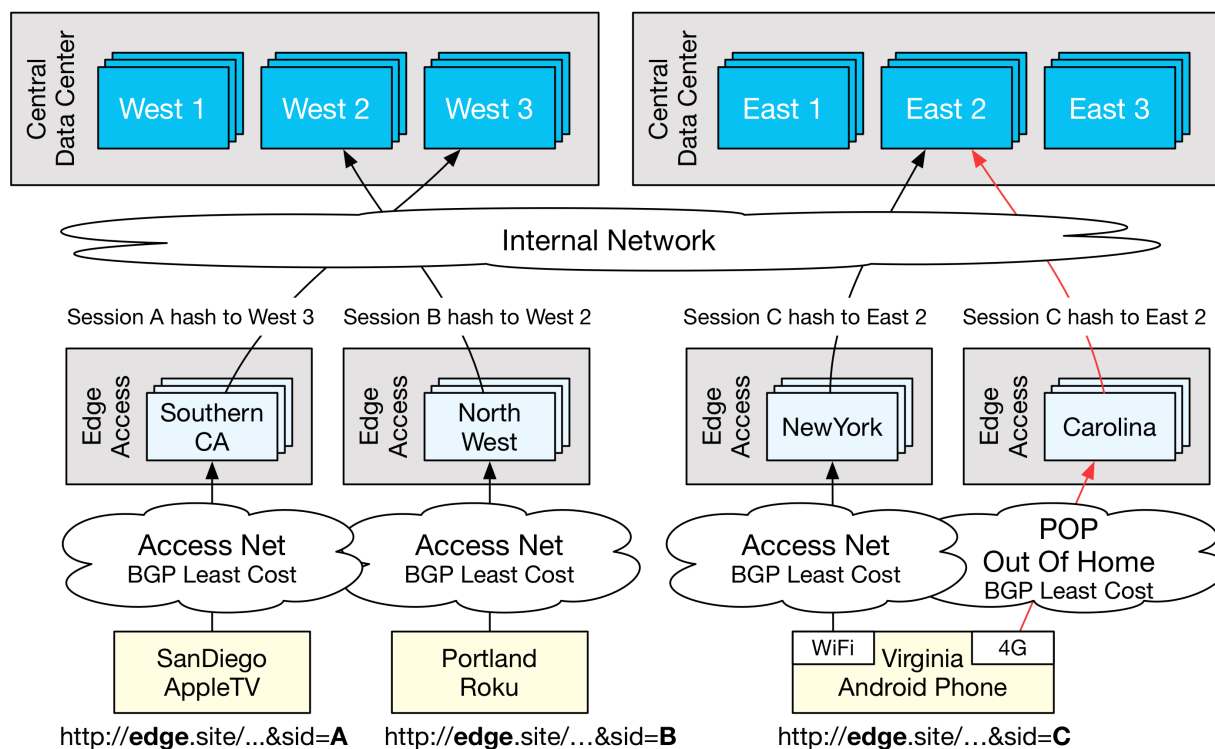
Initially, Charter designed a content-based load balancing scheme for the manifest manipulation function, wherein sessions would be routed to siloed pools of resources based on the channel being requested by the client. At smaller scale, it was relatively simple to implement. As streaming activity grew, load grew more uneven across the siloed pools as the systems serving the more popular channels were taxed more often and sometimes saturated while capacity was unused on other systems.

To scale, Charter borrowed from some of the techniques we just reviewed. It devised a flat load balancing scheme combined with BGP routing that established reachability of the entire pool of manifest manipulation nodes and evened out the distribution of sessions across those resources. That system also



allowed more flexibility in the ability to influence direction of load and traffic. For instance, in order to minimize traffic on the core network, Charter had the ability to move its manifest manipulation resources to the edge and tune its routes such that clients would enter the “closest” edge and be served by the least busy resource at that edge.

Charter also takes care to plan for the adverse impact of “roaming” clients that may change networks and thus may be routed to geographically separate manifest manipulation resources throughout a session. These issues are overcome by leveraging hashing mechanisms using network-invariant parameters associated to the client, combined with more robust session recovery mechanisms between sparse manifest manipulation resources (see Figure 8).



**Figure 8 - Load distribution model in a multi-region deployment**

### 6.3. Appliance to Data Center transition: the importance of hardware tuning

Over the years, Charter has gathered real experience and knowledge in how to leverage hardware virtualization technology to adapt its platform management policies to high performance application needs. An important aspect of manifest manipulation is the management of all the latencies in the critical path to delivering the stream to large volumes of clients.

Virtualization offers hardware independence and enables leveraging continuous improvements in processor, disk, and memory speed and size to offset cost of growth. However, Charter has found that tight policies are required in how the virtual platform is managed to avoid sacrificing performance and service quality.

Manifest manipulation is a mission critical application that directly impacts the user’s quality of experience (QoE). When it works well, the application masks a lot of the complexities and interactions

that happen behind the scenes to deliver a stream uniquely tailored to each connected client. Performance optimization is about ensuring that the application has uninterrupted and reliable access to compute, memory, and storage resources to minimize the delays of processing. The application tolerates latencies with Packager and ADS services within certain limits by design.

Charter evolved its deployment over the years. It started by deploying manifest manipulation on bare metal. As the scale of the deployment quickly grew, Charter moved the application to a virtualized infrastructure to provide better economy of scale and streamline the management and operation overhead of the platform. It took a sustained effort to tune the performance of the application running on top of a virtual environment. The following paragraphs describe the lessons learned along the way.

The initial toolset and virtual stack employed were still maturing to handle video streaming workloads. Tools like tcpdump were required to trace frequent lost SYN packets to the virtual switch. In addition, the system had to be further partitioned and the TCP stack was adjusted to more quickly abort the connection and retry to fetch the manifest from other CDN edges. The manifest manipulator was enhanced to include manifest caching, which greatly reduced downloads from CDN edges and relaxed the application's susceptibility to download latency, while preserving timely responses to client requests.

CPU contention was another challenge as other VMs sharing the host would briefly stall the throughput and responsiveness. While reports showed average resource contention below maximum capacity, they masked punctual spikes of usage by other applications sharing the host, thus starving the manifest manipulator of CPU cycles long enough to cause the potential for bitrate downshifting, which means degradation of the streaming user experience. This was detected using Linux metrics for CPU-stealing and internal health pings that showed gaps in CPU execution.

Charter even changed virtual stacks in the process of growing its infrastructure. The lessons learned from the tuning efforts described above proved valuable in reducing the cost of this transition. However, with more applications running on VMs, other issues needed attention.



**Figure 9 - Percent of client requests at different bitrate qualities**

The graph in Figure 9 plots the percentile of each bitrate in the Charter ladder relative to all client requests recorded over a two-month period after the migration. It shows that clients had a low error rate, but quality levels being played were lower than expected.

This was primarily resolved by increasing capacity at the CDN edges to improve download rates for media segments, and by tuning routing to reduce delay of manifest fetches from the CDN. After these improvements, the average user playback quality was significantly improved. The remaining lower quality levels may be accounted for by mobile clients on Wi-Fi or roaming out of network.


Reviewing the metrics developed as part of the first virtual stack roll-out also helped quickly detect instances of CPU execution gaps on the second virtual stack deployment. The reported metrics allowed teams from video operations, infrastructure, and application vendors to work closely together and isolate cases of automated migration of manifest manipulation instances between hosts. The migration caused a stall with the bit rate downshifting effect previously described. Despite the fact that automated VM migration is common practice in data centers, this case helped shed light on the fact there is no cookie-cutter approach to managing applications in a large scale IP video deployment. Automated migration was ultimately disabled for manifest manipulation.

## 6.4. Multi-ADS support

Another scaling dimension that may be less obvious is measured in the platform's ability to accommodate new ad sales models and relationships. Service providers may open part of their own ad avails to third parties.

The challenge is that opportunities of this type often require working with external and third-party Ad Decision Systems because the parties involved want to avoid having to integrate their backend sales order, ad campaign, and ad preparation workflows with the service provider's own Ad Decision System.

At Charter, the manifest manipulator is equipped with the ability to share ad breaks among multiple Ad Decision Systems, thanks to its ability to route based on ad marker policies, on channel or content, or on parameters conveyed in client URLs. Further support of VAST Wrappers means that Charter retains full control of how an ad request gets fulfilled, because the request first goes to Charter's Ad Decision System for fulfillment of its own inventory before it redirects the manifest manipulator to other third-party ADSs to request ad decisions to fill the remainder of the ad break.



**Figure 10 - Manifest manipulator routing to multiple Ad Decision Systems**

Note the presence of a Retranscode function that serves to automate the adaptation of third-party ads into the operator-specified ABR profiles.

It is then up to the manifest manipulator to properly convey to the device client the tracking events specified by the various ADSs so that ad measurement and verification flow as expected to each party to

the transaction.

Charter engineered its platform to be flexible and expandable to support not just its own avails, but also to enable more creative sales models providing win-win opportunities with other companies in the Television Advertising industry.

## **6.5. Client diversity and integration**

At Charter, the manifest manipulation platform is designed to support DAI and alternate content switching for both HLS and DASH streaming clients. While these two protocols share the same core concept, there are fundamental technology differences. For HLS, the manifest manipulation occurs at the sub-level manifest and involves individual video and audio profile manifest requests. The manifest manipulation in DASH occurs at a single manifest level which encompasses the adaptation set for various video and audio profiles.

In HLS streams, the primary audio stream for the content is most often packaged in the same media segment as the video stream. Some content may involve multiple language versions of the audio experience (e.g. multiple audio streams), which are packaged in secondary tracks fetched separately from the video segments by the client. These multiple manifest requests can drive additional load on the manifest manipulation and should be taken into consideration while designing for scale.

The other key consideration concerns any variation in encoding/packaging profiles between the entertainment content stream and the ad or alternate content stream. As the ads and entertainment streams originate through different sources and the encoding/packaging is managed and operated by different operations groups, there are bound to be variations in how the bitrates, resolutions, and other parameters are specified in the manifest. The manifest manipulation application needs to be robust enough to handle these variations without causing any adverse behavior in the client player.

## **6.6. Metrics and Monitoring**

In order to assure high SLAs and achieve high availability, the manifest manipulation platform plays a key role by providing the hooks and the data to collect all service level and performance metrics and build scalable monitoring. The data provided by the platform address both overall performance of video sessions and performance for ad decisioning and insertion.

Key monitoring metrics include:

- Success rate/Errors for all client requests
- Success rate/Errors for all interactions with the ADS
- Latencies from CDN and ad request processing
- Any CDN errors for fetching the content and ad manifests
- Rogue requests to the application
- Shifts in the bitrate manifests

They help build a robust analytics dashboard to monitor:

- Overall performance of the video streaming platform and any adverse impact to user experience
- Overall performance of the ad placement platform, including measurement of missed monetization opportunities
- Performance of a specific ad campaign or ad creative and success with the placement

Leveraging the telemetry data and log infrastructure instrumented in its manifest manipulation platform, Charter was able to build tools tailored to the specific needs of the Video Operations and Ad Operations teams to enable shared and complementary monitoring of system performance and of individual application process and functions.

Charter leveraged the following dataset from the application:

- 1) HTTP stats with endpoint URLs for all northbound, southbound and ADS requests
- 2) Logs for all ADS requests and responses
- 3) Errors from manifest processing, including under-runs, over-runs, and request timeouts
- 4) Aggregate metrics for ads placed in particular content streams
- 5) Logs for processing time to estimate minimum, maximum and average latencies
- 6) Logs for individual client session context
- 7) Logs for filtering of rogue requests

By harnessing the granularity and depth of the manifest manipulation data, coupled with datasets from other points of the delivery pipeline, it is also possible to more effectively mitigate and resolve complex issues.

## Conclusion

Charter learned valuable lessons and drew major benefits as a result of implementing changes to scale its IP video and advertising system.

Leveraging horizontal scaling for manifest manipulation significantly improved the operator's ability to cost-effectively expand the system on its own infrastructure. As a result, Charter can better adjust to growing loads and deploy temporary instances on cloud infrastructure to absorb occasionally large load spikes.

By taking advantage of the large arsenal of standard routing and load balancing techniques, Charter was able to achieve more efficiency through even distribution of the load.

To scale with more agility, Charter moved manifest manipulation from dedicated appliances to virtualized infrastructure, establishing clear operational policies to ensure the constant availability of compute resources and ensure high quality session handling.

Technological innovations leading to the support of multi-ADS routing helped expand Charter's monetization opportunities.

Leveraging manifest manipulation has helped Charter speed up new client integrations while adapting the manifest delivery to ensure a user experience consistent across an ever-expanding landscape of IP-connected devices.

Finally, Charter leveraged its experience with IP video delivery to define and implement a greater breadth of Quality of Experience metrics and enhanced its data infrastructure in order to more effectively monitor its network and applications at scale.

Today, Charter serves IP video streams in large and growing volumes. With a more scalable architecture and the operational tools to grow, Charter sees few challenges ahead on its path to meeting or exceeding broadcast quality delivery for IP Video. Standardization efforts anchored by the Common Media

Application Format (CMAF) will go a long way to improving scale and reducing content management complexity by allowing common media and encryption across HLS and DASH platforms and many native DRM systems, just as chunked-based encoding and transfer will help reduce ABR latency.

Charter did not make lightly the decision to invest in IP, and for all the efforts made to get this far, there is still the major challenge many operators face to figure out how to transition to the future of IP. The question now is, “How to leverage the investments made in the IP video workflows, a robust and scalable CDN backbone, and advanced Advertising Management Systems, to effectively serve and monetize both IP and QAM platforms over what can be predicted to be a long transition period?”

# Abbreviations

ABR	Adaptive Bit Rate
ADS	Ad Decision Service
ACDS	Alternate Content Decision Service
BGP	Border Gateway Protocol
CDN	Content Delivery Network
CMAF	Common Media Application Format
CPU	Central Processing Unit
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DRM	Digital Rights Management
GSLB	Global Server Load Balancing
HLS	Hypertext Transfer Protocol Live Streaming
HTTP	Hypertext Transfer Protocol
HTTP-MPEG 2 TS	Hypertext Transfer Protocol – MPEG 2 Transport Stream
IP	Internet Protocol
ISBE	International Society of Broadband Experts
LB	Load Balancing
MHz	megahertz
MPEG-DASH	MPEG – Dynamic Adaptive Streaming over Hypertext Transfer Protocol
MPEG TS	MPEG Transport Stream
NAT	Network address translation
OTT	Over the Top (video service)
PPV	Pay per view
QAM	Quadrature amplitude modulation
QoE	Quality of Experience
RGUs	Revenue generating units
(RHI / BGP)	Route health injection / Border Gateway Protocol
SCTE	Society of Cable Telecommunications Engineers
SLA	Service Level Agreement
SSAI	Server-Side Ad Insertion
STB	Set-top box
SYN	Synchronize
TCP/IP	Transmission Control Protocol/Internet Protocol

TTL	Time-to-Live
UDP	User Datagram Protocol
VAST	Video Ad Serving Template
VIP	Virtual IP
VM	Virtual Machine
VOD	Video-on-Demand
VPN	Virtual Private Network
Wi-Fi	Wireless Fidelity (IEEE 802.11X)

## Bibliography & References

- [1] Strategy Analytics Global Advertising Forecast – 2010 -2030, Michael Goodman, August 30, 2018
- [2] *Q1 2019 Conviva's State of the Streaming TV Industry*; Conviva
- [3] Extreme Reach Q3 2018 Video Benchmarks, Mary Vestewig, October 31, 2018
- [4] Industry Voices—Dan Rayburn: Handle Manifest Manipulation at the network edge to personalize video experiences; Fierce Video, Jul 1, 2019



# **2019 Virtualized CPE Services Have Finally Arrived Via Service Delivery Platforms**

## **Home Gateway Feature and Service Flexibility, Avoiding The Monolithic Upgrade Path**

A Technical Paper prepared for SCTE•ISBE by

**Ian Wheelock**

Engineering Fellow, CPE Solutions  
CommScope

4300 Cork Airport Business Park, Kinsale Road, Cork, Ireland  
00353-86-235-2712  
Ian.Wheelock@CommScope.com

**Charles Cheevers,**

CTO CPE Solutions  
CommScope

3871 Lakefield Dr, Suwanee, GA 30024  
678-473-8507  
Charles.Cheevers@CommScope.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
Content .....	5
1. Gateway Platforms.....	5
1.1. Gateway Stack .....	5
1.2. RDK .....	5
1.3. Adding 3 <sup>rd</sup> Party Software .....	6
1.4. Adding Operator User Interfaces.....	7
1.5. Long Term Maintentnace.....	7
Architecture for consideration .....	7
2. Variables.....	8
3. Device type, RAM and flash considerations.....	10
3.1. Software deployment model.....	10
3.2. Compression .....	11
3.3. Compression, Flash and Containers .....	12
Router Stack options .....	12
4. Linux based stacks.....	12
4.1. RDK .....	12
4.2. OpenWrt.....	13
4.3. Proprietary.....	13
5. SDKs, HW integration and HALs .....	13
6. Higher and Lower Layer SW interfaces .....	14
6.1. prpl Higher Layer and Lower Layer API (HL/LL-API) .....	15
6.2. Open vSwitch/OVSB.....	16
6.3. Lua Based Architecture .....	17
6.4. CommScope Container API.....	18
6.5. Life Cycle Management(LCM) / Service Delivery Platform (SDP) .....	18
Container Usage - LXC or Docker/Balena.....	19
7. Container Choice .....	20
7.1. LXC Container.....	21
7.2. Docker Container.....	21
7.3. Balena Container.....	22
Container Experience in Gateways .....	23
8. Containers on Commscope Gateway Platform .....	23
9. Life Cycle Management (LCM) on OpenWrt platform.....	24
10. Nomad POC .....	25
Alternative Virtualization Options .....	26
11. Virtual Machines .....	26
12. Full and Hybrid Cloud Virtualisation options.....	27
Data Plane .....	28
Control Plane.....	29
Concurrency and Orchestration Scalability.....	30
Conclusion .....	32

Abbreviations.....	33
Bibliography & References .....	34

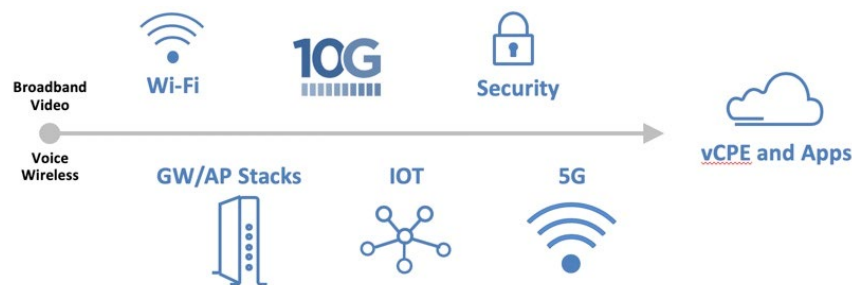
## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Factors Driving New Software Services.....	4
Figure 2 - Traditional Software Development Model .....	5
Figure 3 - Agile Development Model .....	6
Figure 4 - Key HW Elements of Gateway .....	8
Figure 5 - Proposed home gateway, applicable to multiple access types .....	8
Figure 6 - Sample of Potential Services Considered.....	9
Figure 7 - Proposed Router/Services Platform close-up .....	9
Figure 8 - Example GW Memory Trends.....	10
Figure 9 - Flash and RAM Organisation .....	11
Figure 10 - Flash memory Organisation .....	11
Figure 11 - HW/OS/Application Layering.....	12
Figure 12 - Key RDK-B Software Layers and Components.....	13
Figure 13 - prplWrt Organisation .....	13
Figure 14 - Future Router Stack Architecture .....	15
Figure 15 - prpl High and Low Level APIs .....	16
Figure 16 - OVSDDB/OVS Based Architecture .....	17
Figure 17 - NFLua Packet interception and Agent Architecture.....	18
Figure 18 - prpl Service Delivery Platform/Life Cycle Management.....	19
Figure 19 - Native Apps vs Container Apps.....	20
Figure 20 - Docker Ecosystem.....	22
Figure 21 - SDP/LCM and Orchestration Overview .....	24
Figure 22 - SDP/LCM POC Configuration .....	25
Figure 23 - Native, Containers and Virtual Machines .....	27
Figure 24 - vCPE with Cloud Services .....	28
Figure 25 - Potential Service Load over 24hr Period.....	30
Figure 26 - Multitude of Options for Virtualised CPE.....	31

# Introduction

Current models for adding new services and features to the home are highly reliant on upgrading gateway devices with a monolithic firmware image. Typically, lots of effort is required from the Cable Operator, the gateway Original Equipment Manufacturer (OEM), and possibly a 3<sup>rd</sup> party Software supplier to add these new features. This not only involves the specification of how everything should fit together, but the planning, development, and testing of the new feature, as well as the entire monolithic firmware deliverable. As one can imagine the time and effort involved can be considerable. Once the monolithic image is created and deployed, the whole cycle restarts with the next feature or service the Cable Operator would like to deploy.

This model has worked. However, when compared to mobile phones or laptops, adding new software features typically does not require an OS upgrade. Why can't gateways follow this model? or use something a lot more agile that has fewer moving parts to enable faster feature and service delivery to subscribers?



**Figure 1 - Factors Driving New Software Services**

This paper will concentrate on exploring what architectures and platform options exist today to deal with service delivery beyond the monolithic image system and examine the pros and cons of these including how virtualization techniques both in the gateway and in the cloud can be used. Details relating to RAM, flash, and CPU resources will also be covered. The paper will also address aspects of cloud-based applications based on application traffic tunneling and compare these as potential alternatives to thicker gateway hosted services. Its organized as follows and covers the following sections:

- New software delivery options that are beginning to appear in the industry
- What they mean for operators, OEMs and 3<sup>rd</sup> party software/service vendors
- How they might be applied to existing and future gateway platforms
- Impact on RAM/flash/CPU resources
- How to manage or orchestrate these services
- The tradeoff between thick gateway services vs virtualized cloud services
- How gateway traffic filtering and tunneling enables these services

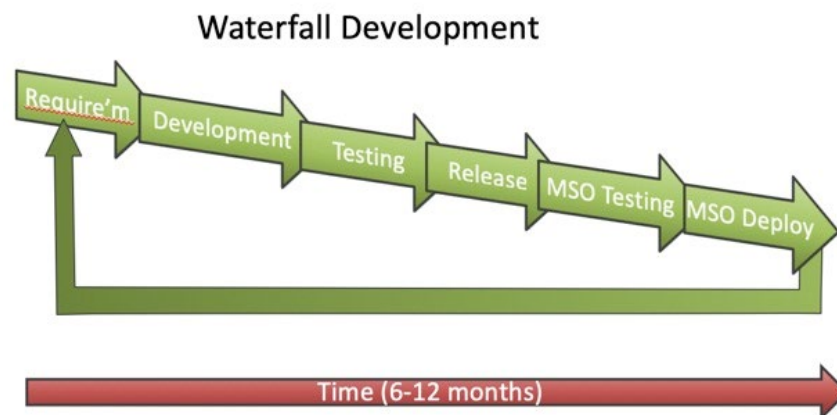
# Content

## 1. Gateway Platforms

Gateway platforms today have not changed much compared to initial gateways in relation to how software is developed and deployed, however the demand for new software and services has increased, particularly the integration of 3<sup>rd</sup> party software and services into existing gateways.

### 1.1. Gateway Stack

The standard approach for a gateway platform is to target a set of routing/networking features typically tied to a release of Linux, and maybe replace some of these with additional or upgraded components depending on the operator requirements for a new release, and then mix in a bunch of management controls and logging options to be able to manage and troubleshoot the platform in the field. The various changes are developed and unit tested, then are all mixed together to produce a monolithic image. This image is then submitted to the OEM test teams for system testing before the final release candidate is made available to the cable operator. The cable operator then takes this new release and applies their own acceptance testing to this image before finally releasing to the field. Finally, the operator can begin to offer the new feature set to their subscriber base. Proprietary, RDK-B, and openWrt gateway platforms all follow this general model.



**Figure 2 - Traditional Software Development Model**

This big bang approach is repeated over and over as new features are requested, or improvements/bug fixes need to be incorporated. In some instances, depending on the type of change involved, more focused testing can be performed, getting the final monolithic upgrade ready for deployment. This lengthy development process include approach is typical within the industry, and has remained in place as a compromise to managing the risk of launching a completely new release out to thousands or millions of deployed gateways.

### 1.2. RDK

In most cases this development process is as streamlined as its going to get. Some new platforms like RDK enable more control around the build environment, and focus on continuous development and integration of new features, working towards constant integration and deployment of these into the field. This model depends on a lot of automated testing, significant logging support, and the ability to move

from development directly into customer deployment on a regular cadence, perhaps once a month or even once a week. Given the exposure of regular updates of software in the field and knowing exactly what has changed from minor release to minor release, this model enables more focused testing and resolution of issues on specific features compared to the big bang omnibus release of features previously described.



**Figure 3 - Agile Development Model**

In most cases where RDK has been deployed with a high release cadence, the operator involved has been tightly coupled with the actual development process, sharing bug tracking and build systems with OEMs they have partnered with, and requiring their own development teams to be able to guide the overall release planning/development of features as well as deal with issues arising from the field (performing triage, collection of logs, etc.). There are definite benefits from this high cadence approach in terms of quick turnaround of new features and bug fixes, but the model does require the cable operator to get down and dirty with the development process, as well as owning the release, system test and deployment processes.

All of this costs money by moving the operator into more of an OEM/development role. In most operator cases, developing software themselves is not how they make money from their business. The RDK codebase/architecture can still be used in the traditional development model, where an operator works with an OEM to release a set of new features and updates at a much lower cadence - maybe once every 6/9/12 months - while building on the stability of known RDK releases from the RDK community.

Both these approaches require either the OEM to do significant development or have the operator get tied into the development process, possibly at an uncomfortable level.

### **1.3. Adding 3<sup>rd</sup> Party Software**

In the case of adding 3<sup>rd</sup> party software, generally there is a need to involve the software provider into the development process, one way or another. Such software maybe supplied in source code format or in binary/library format. For the binary/library format, the 3<sup>rd</sup> party typically requires access to the various code compiler elements to be able to cross compile their source code into a library that can be linked in to the monolithic firmware image.

This approach normally requires the OEM to develop target platform layer interfaces that the 3<sup>rd</sup> party library requires. If source code can be provided, this gets built directly by the OEM themselves, with the OEM still needing to develop the target platform layer, as well as any other management control/logging functions to fit into the existing platform. (The main reason source code is not normally shared is in order to protect any associated Intellectual Property Rights (IPR) from being exposed to OEMs/other parties)

Once the 3<sup>rd</sup> party software is integrated into a monolithic firmware image, it is subject to testing which typically involves the 3<sup>rd</sup> party vendor, the OEM and finally the operator acceptance testing. Again, a somewhat complicated setup.

#### **1.4. Adding Operator User Interfaces**

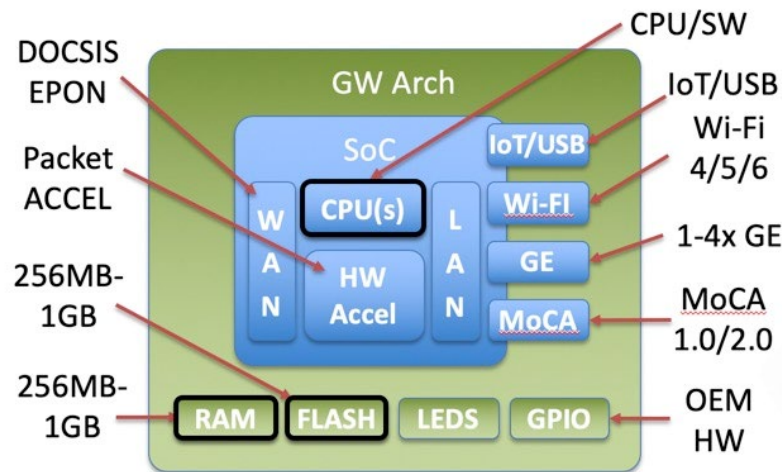
In the case of adding subscriber user interfaces, be they webpage based or mobile app based, the typical approach is for the operator to provide use-cases and some sample screen shots as part of a set of requirements to the OEM to implement. The user interface will typically have extra options included over time depending on what new software features maybe added. Such requirements are typically treated no differently than feature requirements, so being added to this lengthy process that results in a monolithic firmware image being produced. Unfortunately, user interfaces are very subjective due to the interpretation of the look and feel characteristics versus the actual implementation. Another issue with user interfaces is that an operator may have two OEM suppliers of gateways, or indeed have multiple language or countries where the number of OEMs increases but still needs an identical look and feel to apply to all gateways. Given the number of OEMs, expecting to get these independent software developers produce identical look and feel is a challenge.

#### **1.5. Long Term Maintenance**

Overall, adding software features and services to existing gateways is typically complicated and quite involved, and in most cases has to be repeated for every different gateway an operator uses. Once such features and services are deployed, the time to fix issues or update to newer releases of the feature/service is determined more by the overall development and testing cycles rather than the availability of the fix/release, particularly with 6-12 month release cadences in use by a lot of operators. Such an approach can be extremely frustrating and significantly hampers feature velocity. If an operator choses to get involved in the development cycle of RDK-B, then it's possible to accelerate this, but one must remember that the gateway also has to preserve the robustness of all previous software features as well as any new features being developed/added, which is another cost to the development process.

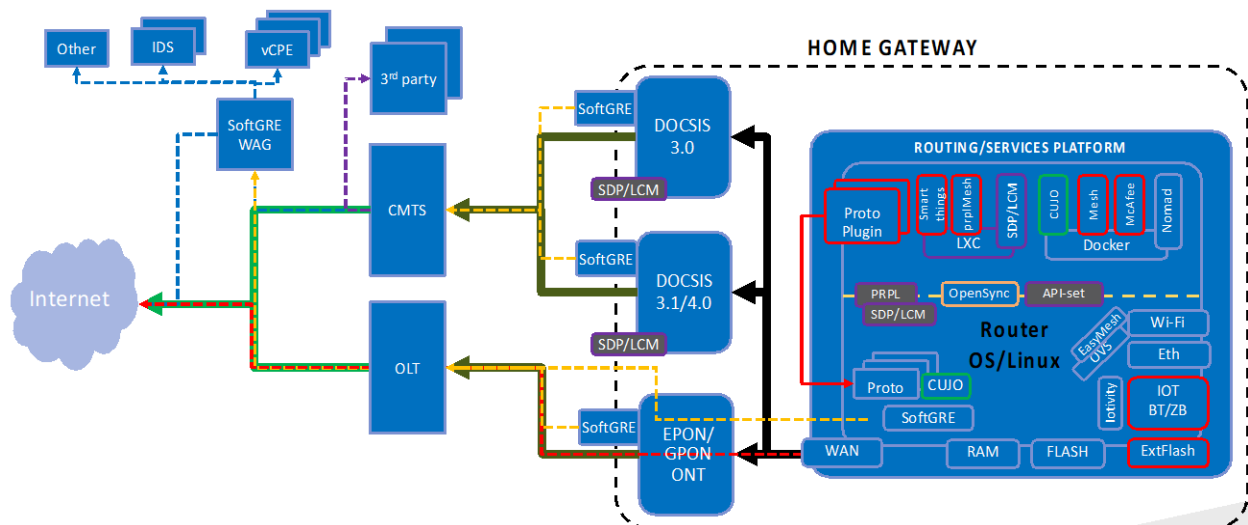
## **Architecture for consideration**

MSOs have a number of different devices that they use for broadband access, ranging from the most recent advanced gateways to older legacy products. Traditionally these devices sole purpose was to provide the networking platform for delivering broadband access to subscribers. In order to stay relevant and offer more than just a dumb internet pipe, MSO's have been evaluating how to enable new services for subscribers in a cost-effective way that can be run on existing devices in the field or augmented via hybrid cloud options. Given the mix of devices in the field, the target architectures need to cover services delivered solely on existing router devices or distributed in multiple places. A big challenge in achieving this is the range of devices in use, all with different capabilities as well as the type of expected service to offer.



**Figure 4 - Key HW Elements of Gateway**

The following system architecture is proposed for MSO's to consider. This uses a mixture of container based orchestrated software services on in-home devices, as well as offering a hybrid option for services provided partially in the home and mostly in the cloud (through tunneling and the use of iptables or ovs on the in-home platform). The architecture also shows some tighter integration of services within the platform itself, for when some software needs higher performance access to networking or other lower layer services. The architecture can apply to most any WAN access, with D3.0, D3.1 and PON all shown.



**Figure 5 - Proposed home gateway, applicable to multiple access types**

## 2. Variables

A lot of variables shape this architecture including: the device type; available RAM; available storage/flash (and read/write ability of same); the type of containers to be used; the number of services that may need to reside on the platform; whether these services are provided by the MSO directly, contracted partners, or third-party developers.

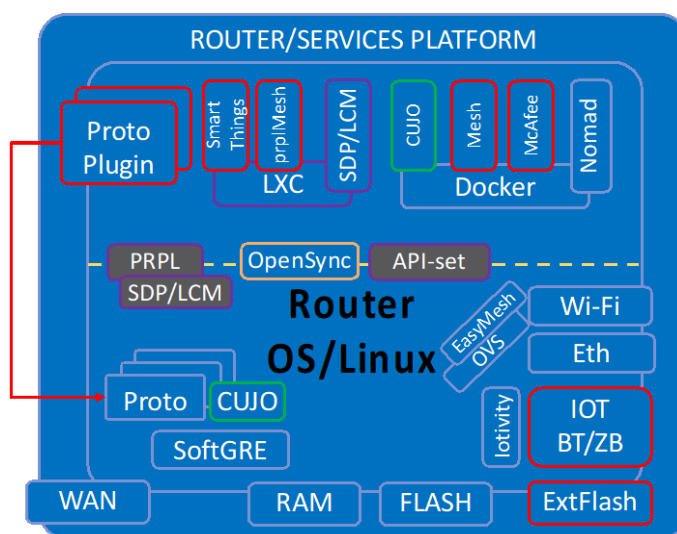


The types of new services are unlimited, with a huge array of options available to be added to the gateway platform



**Figure 6 - Sample of Potential Services Considered**

Other factors include: the supported services for the devices themselves; the types of low-level or high-level APIs available; the infrastructure for hosting services on the device; the types of services being considered (tightly integrated networking applications, or apps that only need IP connectivity); as well as what type of access to hardware or local software stack is needed and what the managed API interfaces to use are, etc.



**Figure 7 - Proposed Router/Services Platform close-up**

Once the overall view of a containerized approach for local applications within devices is agreed, other decisions must be made regarding the types of container infrastructure to use, the type of orchestration

involved, the ability to manage and monitor not only the individual device infrastructure but also the entire network of devices running containers, the overall performance of the system including initial deployment, upgrade and mass reconfiguration across the footprint of devices. There is also a need to consider the mixed management of normal day to day operation and maintenance of the deployed broadband system and these integrated services on the same infrastructure.

### 3. Device type, RAM and flash considerations

The deployed standalone devices have a variety of RAM and flash capabilities. Many platforms are quite limited, only including 512MB of RAM and 128MB of flash, while newer platforms are considering 512MB RAM and 512MB flash, or even 1GB RAM and 1GB flash. Cost is the main driver regarding how much storage to add to a device. A lot of purchasing decisions are made on the basis of the “hear and now”, as opposed to the total cost of ownership of the device and what feature upgradeability might be lost if too little memory is specified. For many years both RAM and flash costs were a large part of a gateway design, and something that could be manipulated with in the design, i.e. I need Wi-Fi, but maybe I could get away with 128MB instead of 256MB of RAM, particularly given price per MB. As a result, if an operator needed price reduction on a new design, RAM and flash were up for the chop. This made sense given 100,000s of device deployed, saving maybe \$2-4 per unit is a big CAPEX save.





Figure 8 - Example GW Memory Trends

Depending on the routing software stack itself, most of this storage may already be consumed. Some firmware stacks using OpenWrt can actually be made to operate within 4MB flash and 32MB RAM, but in reality, need more like 128MB flash and 256MB RAM to fully support operator features.

#### 3.1. Software deployment model

The typical software deployment model for broadband devices is to generate compressed squashFS firmware images (containing the full routing stack and any other software functions) that are distributed and stored in device flash, and during bootup of the device are decompressed from flash into RAM. In most cases the decompressed image itself is a soft copy of the platform Linux root filesystem (rootfs) which is presented in RAM to the OS. The OS will then launch applications from the rootfs.



**Figure 9 - Flash and RAM Organisation**

Another decision made with broadband devices is to use a “dual image” option for firmware image storage, where two complete compressed images are stored in the flash storage, effectively limiting the maximum image size to under 50% of the available flash memory. This is done to have a backup image, in case an image has been corrupted in flash (due to various possible reasons).



**Figure 10 - Flash memory Organisation**

Thankfully RAM and flash pricing has corrected over the last 2 years, meaning prices have come down (different reasons for RAM and flash). However, operator purchasing decisions regarding RAM and flash have had consequences on what feature upgrades may apply to existing deployed device, and in some cases, there just is no space left to factor in any local extensions or alternatives, and alternatives, such as hybrid or virtual cloud services must be considered.

Another aspect of broadband devices is that due to having a single monolithic firmware image containing all the software for the system to operate with, any minor changes requiring a complete replacement of this monolithic image. Even though this appears quite inefficient, there are a lot of operational benefits in knowing that a population of devices are running version #N or version #N-1 of firmware.

Most broadband devices limit flash storage to be READONLY, with only the bootloader or firmware upgrade process being able to write anything to flash. This is a major issue when considering the download and storage of software components separate to the main firmware image. In some platforms, read/write of flash is already supported, but other platforms may need bootloader/code refactoring to accommodate this mode of operation.

New software services packaged in containers (and similar) tend to be overlaid on top of the existing firmware image in some instances taking advantage of features/libraries within the platform image. Other software however may need to be integrated directly with the existing platform image, possibly replacing or adding functionality.

This idea of live patching of the platform itself brings considerable complexity and risk from the point of view of both modifying the actual system properly and ensuring that a patch does not cause any issue to the running system. Also, the management of a mixed population of devices that may have different levels of “patching” applied may present significant operational overhead.

### 3.2. Compression

Given the nature of compression, the firmware image is likely to be much larger in RAM when decompressed. The compressed firmware image is typically CRC/MD5/signature checked before any attempts to decompress/execute code to make sure the image has not suffered any corruption while resident in flash (or due to misprogramming) and that it is a proper cryptographically signed image. The

firmware image includes the Linux kernel, drivers, and complete root filesystem. This system of compressing firmware images is the general approach used on all embedded platforms that use NAND flash memory, as it is not possible to execute directly from NAND, compared to NOR flash.

### 3.3. Compression, Flash and Containers

However, in a platform that may need to offer “container” based services, it may be better to consider separating how container images are stored and accessed in flash compared to the platform firmware image. Like most images, container images are compressed and, once downloaded, are accessed via squashFS. Isolating the container images to a separate flash partition/location will allow the container execution environment access the images directly from flash rather than requiring the complete container root filesystem to be copied to RAM. RAM is still required to load and run the various program files that comprise the container.

In nearly all these cases there is a need to support OverlayFS (a key Linux feature, in mainline since 3.18) to ensure any configuration elements or read-write locations are handled separately to the container read-only space in flash. This approach can reduce the overall amount of expensive RAM required (for storage purposes) on a platform at the cost of adding additional flash, and allows for flexibility in adding extra flash using either onboard eMMC or via plug-in USB/xSD devices.

## Router Stack options

### 4. Linux based stacks

In general, any router stack based on Linux is appropriate to use when considering the addition of new software and services on to a gateway platform. RDK-B/-M, OpenWrt, and proprietary stacks are all candidates for this. In most cases an abstraction or high-level API interface is really important to be able to offer a target layer for 3<sup>rd</sup> party software to work on top of.


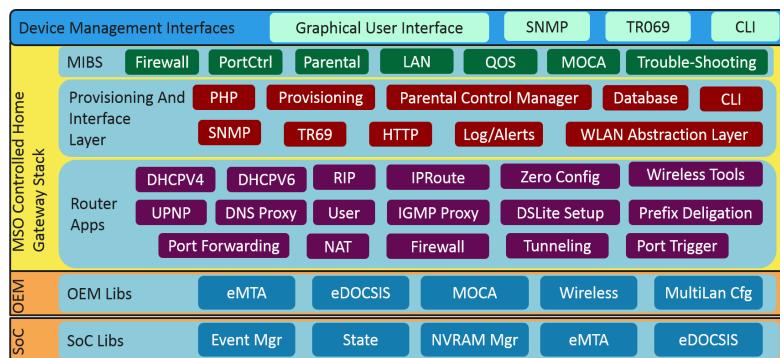


Figure 11 - HW/OS/Application Layering

#### 4.1. RDK

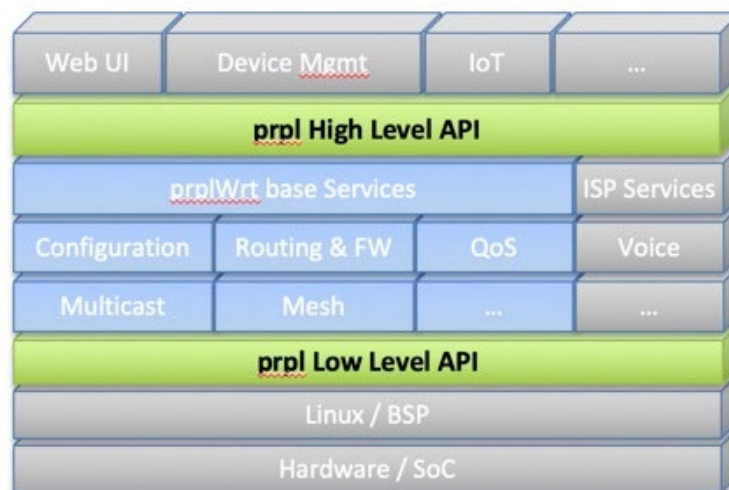
RDK-B/-M itself has an internal CCSP bus (based on DBus) that acts as the backbone of the system, connecting the core RDK-B subsystems together with protocol adapters and software components. New software can be added into this system and have complete first class access to the inner workings of the platform. Support for low level interfaces, such as Wi-Fi HAL, or Cable Modem HAL, or Ethernet/Switch HAL (utopia) is also included, as is support for managing configuration settings/NVRAM. External protocol adapters for TR-069, SNMP and the Comcast developed WebPA interface allow management access to the system.



**Figure 12 - Key RDK-B Software Layers and Components**

## 4.2. OpenWrt

OpenWrt also has an internal bus, uBus, that acts as its backbone for enabling communication and control between all the internal elements that are used for routing and management. It uses “uci” for its configuration management, and offers a lot of equivalent services that RDK-B/-M offers that are typically expected in a gateway stack. prplWrt packages together some new carrier class components into openWrt.



**Figure 13 - prplWrt Organisation**

## 4.3. Proprietary

Proprietary or OEM stacks, such as ARRIS Touchstone or ARRIS 9.x, all offer the same type of functionality as RDK-B/M and OpenWrt. Each stack breaks down the control functions required for each of the underlying subsystems to implement the various software and protocol requirements for a gateway.

## 5. SDKs, HW integration and HALs

In the main, all of these stacks are ported to run on different SOC's through the use of supplied SDKs that provide the base Linux kernel support. The SDKs mostly use Linux defined interfaces, particularly low level interfaces, when possible. When integrating extra hardware with a SOC, new drivers are provided by the 3<sup>rd</sup> party hardware supplier. In an effort to simplify the adoption of different hardware in to the

router stack, say in the case of Wi-Fi, RDK-B has mandated the use of a Hardware Abstraction Layer (HAL) with a view of managing and controlling each Wi-Fi system in much the same way. This requires the hardware supplier to adapt their drivers to support the HAL. In the case of OpenWrt, the approach taken is to leverage existing Linux layers for Wi-Fi, such as cfg80211 or hostapd/wpa\_supplicant, and require the hardware suppliers deliver this interface, with most of them doing so.

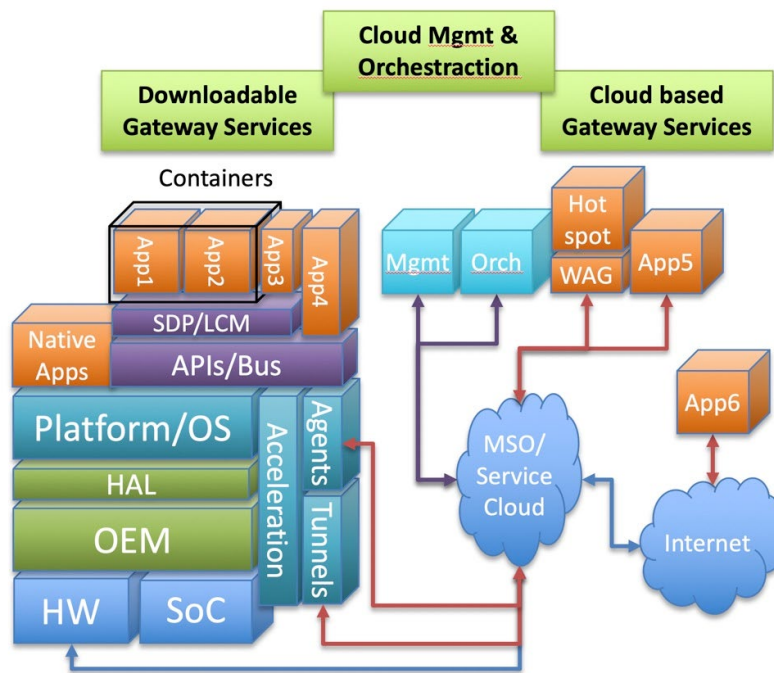
Other subsystems in the SOC, such as the low-level packet acceleration and switching functions are harder to get standard drivers for, as each hardware supplier does things differently matching their HW architecture. However, with advances such as Open vSwitch (where switching is performed in software), Switch Abstraction Interface (SAI) and “switchdev” it is possible to tie in these low-level hardware features in an abstract and performant way into the chosen router stack (as long as the SOC provider supports these features!).

## 6. Higher and Lower Layer SW interfaces

Given the advances in the hardware integration efforts, it could be assumed that the higher layer software interfaces are just as advanced. Unfortunately, this is not the case, although a lot of work is ongoing in this space. In most gateway cases, there was no real need to expose “standard” software interfaces, as no one apart from the OEM vendor was developing software for the platform. The accepted interface into a gateway was typically the network management layer, namely SNMP or TR-069, or alternatively a local HTML/web interface.

The various stacks mentioned all have internal buses for connecting their various HALs and adapters/components together. One straightforward way of exposing software interfaces to 3<sup>rd</sup> party software is to simply provide access to the internal bus. In a number of instances (say high performance network interfacing software), this is exactly how 3<sup>rd</sup> party software is integrated, using the internal bus as well as tight integration with low level driver interfaces. Such integrations can be challenging (requiring legal agreements for source code sharing, engineering access/etc.) and because two or more codebases become so tightly coupled, the only option of releasing bug fixes or enhancements is to release a completely new firmware load (going against the need for speedy releases). Such tight integration may also require more software development resources to achieve the final deliverables.

Using this model for delivering the majority of new software and services cannot scale. Such a model would also threaten the security and robustness of the stack itself, something to be avoided. What is needed are a set of defined interfaces that can be supplied to 3<sup>rd</sup> party/Independent Software Vendors (ISV) to allow them work somewhat independently of the detailed underpinnings of the firmware stack, and they will still likely need the platform tool chain to enable them build software. The following sections outline the different APIs that are available with PRPL, OVS/OVSDB, NFLua and an internal CommScope API. These interfaces are not only critical for so-called “native” software integration where software is built into the monolithic firmware image but also critical for container based options. OpenSync is also described and offers an OVSDB interface that allows a hybrid model of native code developed for the gateway that also interfaces with a remote/cloud system that may be running additional cloud applications.



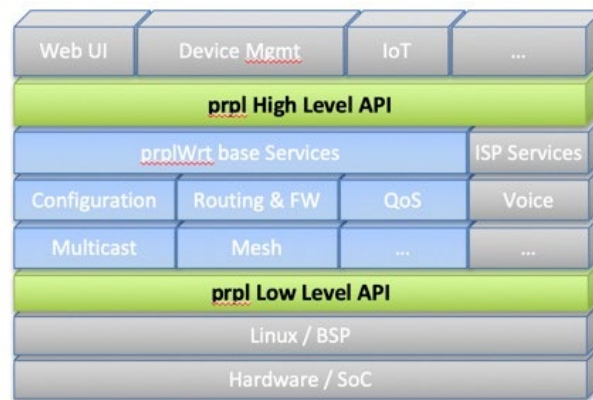
**Figure 14 - Future Router Stack Architecture**

### 6.1. prpl Higher Layer and Lower Layer API (HL/LL-API)

CommScope knows that some MSOs have been very successful with OpenWrt and would like to continue to use this stack into the future. The flexibility of OpenWrt provides for a complete stack covering most features needed for some MSOs. The prpl Foundation are also working on a carrier-grade definition for OpenWrt called prplWrt, as well as other work trying to standardize on higher level (HL) APIs and low-level (LL) APIs (e.g. cfg80211), and pushing Wi-Fi silicon vendors to standardize on the use of Linux Wi-Fi control layers to avoid proprietary drivers.

The prpl High Level API has been considered from the ground up as a platform abstraction layer to enable the delivery of new services to be easily integrated to GW devices. The HL-API consists of a definition of 30+ primary features typically used in a GW as well as a model on how this can be integrated into multiple industry stacks, including OpenWrt and RDK-B. CommScope is currently reviewing the use of the HL-API on RDK-B, and what it will take to work over D-Bus\*. The HL-API is not limited to higher layer services being added to the device, it also supports the idea of new underlying system components being added to a platform that can increase system functionality (and having this available to other software layers). The HL-API and prplAdapter also support features critical to enable 3<sup>rd</sup> party software to be added to platforms, particularly in the areas of access control and “user management”. These areas are fundamental to enabling and restricting what elements of the gateway platform can be interacted with or controlled by software services.





**Figure 15 - prpl High and Low Level APIs**

A key part of prplWrt and the higher-level APIs is to provide a so-called “prplAdapter” component that provides access as well as access-control to the inner operation of the routing/platform stack. This interface approach is meant to help the development of services required by operators, as well as exposing certain APIs to 3<sup>rd</sup> party application developers. Even though prpl has focused on OpenWrt, the major effort on the higher- and lower-level APIs is considered stack agnostic, and the expectation is that these interfaces will be available on RDK-B and other router stacks. CommScope is currently involved in an exercise to identify the work effort for mapping prpl High-level API to RDK-B, while prpl is also pushing the use of certain APIs into the RDK-B community for Wi-Fi management.

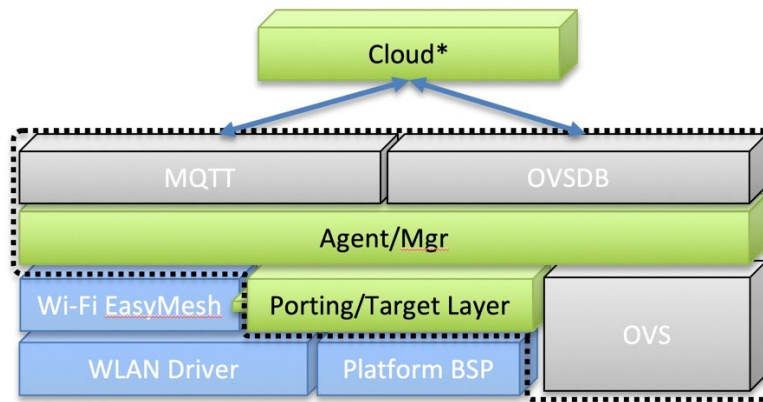
One example of software using both the prpl HL-API and low-level API is the prplMesh implementation, where a software platform exposes control of the EasyMesh controller using the HL-API while also using the prpl LL-API (namely cfg80211) for the EasyMesh agent, interacting with the low-level control and management functions of the Wi-Fi chipset. The portable prplMesh implementation for EasyMesh will run on any platform that supports the LL and HL prpl APIs, as well as exposing the necessary interfaces to allow 3<sup>rd</sup> party Wi-Fi optimization systems interact with the prplMesh EasyMesh controller function.

The set of APIs provide a significant abstraction layer to support development of both applications on the system, as well as exposing stack information (including status and monitoring information) to remote management platforms.

## 6.2. Open vSwitch/OVSB

New “OpenSync” software has been developed that relies upon OVS and OVSDDB to expose internal operation of a gateway to a remote cloud controller. The software currently supports Wi-Fi management and has extensions for local tunneling. Current implementations use a MQTT service for actively monitoring the status of the home Wi-Fi environment back to the proprietary cloud. Open vSwitch is also a key part of this architecture, where the majority of its configuration and monitoring system had been developed. OVSDDB is used in conjunction with Open vSwitch to provide a distributed database solution managed from the backoffice that controls pretty much all of the functionality of the OpenSync home deployment. The model is quite distinct from the existing/ traditional network management model used by operators be it TR-069/098/-181 or SNMP.





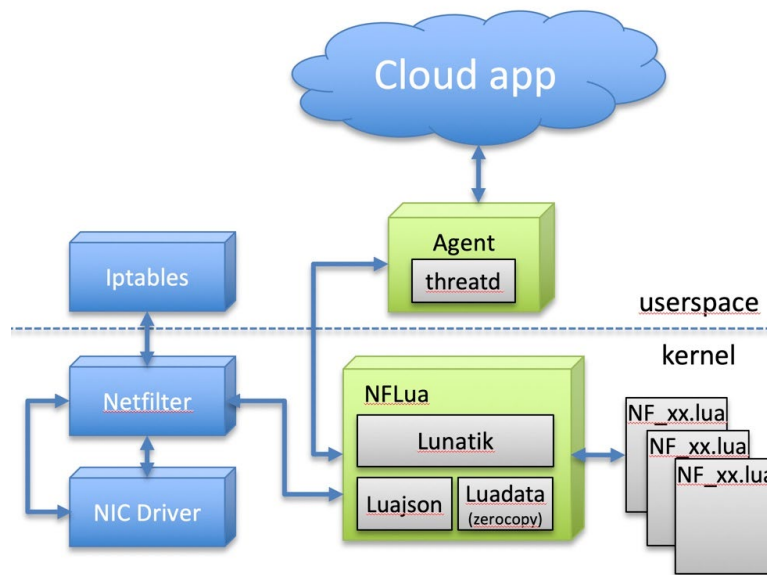
**Figure 16 - OVSDB/OVS Based Architecture**

The OpenSync solution is currently limited to Wi-Fi (on supported platforms) and some visibility into L2 switching. Some additional features such as basic device identification, basic speed test, and QoS control exist along with tunneling support of home network traffic between proprietary Plume Wi-Fi PODs back to the home gateway. The use of OVS for complete switch management is being considered on multiple platforms. Retrofitting it on older SoC platforms may have some challenges due to existing SoC supplied slow-path/fast-path handling and having to deal with very specific WAN access handling. However, where it has been ported, there is an option of dealing with everything relating to packet handling directly in software in the Linux kernel.

The use of Open vSwitch/OVSDB in the OpenSync has the potential to bring an SDN control plane to the operator subscriber network, and could in theory be coupled with hybrid cloud applications where traffic is selected in the home, and delivered using GRE tunnels to cloud applications that provide various software and networking functions, similar to how a WAG works today, but dealing with much more than just Wi-Fi hotspot related traffic.

### **6.3. Lua Based Architecture**

A Netfilter/Lua based architecture has also gained some traction in the industry, with its primary software layers combining netfilter and lua extensions within the kernel to simplify the interception and handing over of packet to user space within a Linux platform. The architecture enables a scalable approach for higher layer applications for interacting with packet flows passing through the gateway.



**Figure 17 - NFLua Packet interception and Agent Architecture**

The solution is a potential alternative to OVS. The NFLua kernel component integrates with the existing Linux netfilter and iptables for packet inspection. The model enables userspace agents to interact directly with NFLua packets that are intercepted, with the ability to operate on these packets locally (using various protocol plugins for different protocols) or to act as an agent to a cloud entity that can process these packets remotely, possibly using more complex or capable functionality not possible in the gateway footprint.

#### **6.4. CommScope Container API**

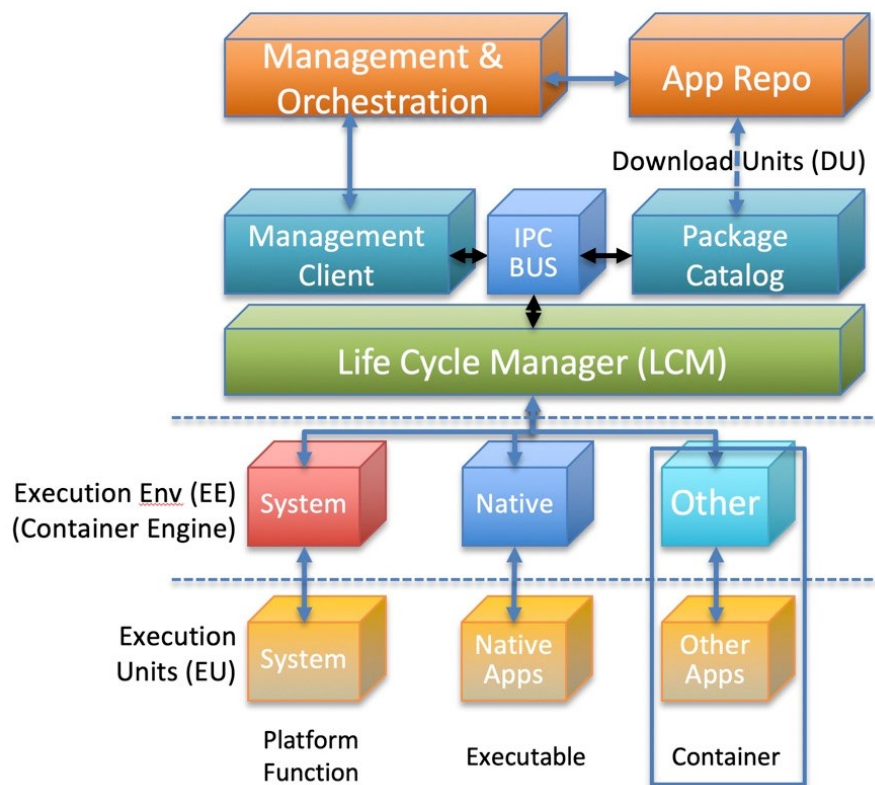
The CommScope container interface is based on providing a controlled API for ISVs to use that enables access to elements of the internal router bus. The APIs expose objects that software can use, while an access control system marshals which software can access what objects. The API interface relies on a form of object based loosely on TR-181 data definitions as well as extensions for interacting with low-level layers within the stack. The interface has already been used for several container applications.

#### **6.5. Life Cycle Management(LCM) / Service Delivery Platform (SDP)**

In addition to the need for the aforementioned interfaces both data plane and control plane for implementing software, there is also a need for software interfaces or a subsystem to manage these new software components that can operate on gateway platforms. This is somewhat equivalent to what Docker provides to manage interactions on a platform as well as interacting with a remote Docker repository hosting available applications, but one key difference is the target platform, in this case embedded gateway platforms.

Typically, Docker is used on extremely capable hardware platforms with plenty of RAM and Flash as well as large CPU resources, something quite different to embedded platforms like broadband gateways. As such, companies have been investigating more lightweight options to achieve equivalent function for gateways. Broadband Forum developed the TR-157 approach many years ago, including a key element known as Software Module Management (SMM). The SMM system provides the basis for a new Life

Cycle Management (LMC) and Service Delivery Platform (SDP) that Vodafone has created over the last number of years.



**Figure 18 - prpl Service Delivery Platform/Life Cycle Management**

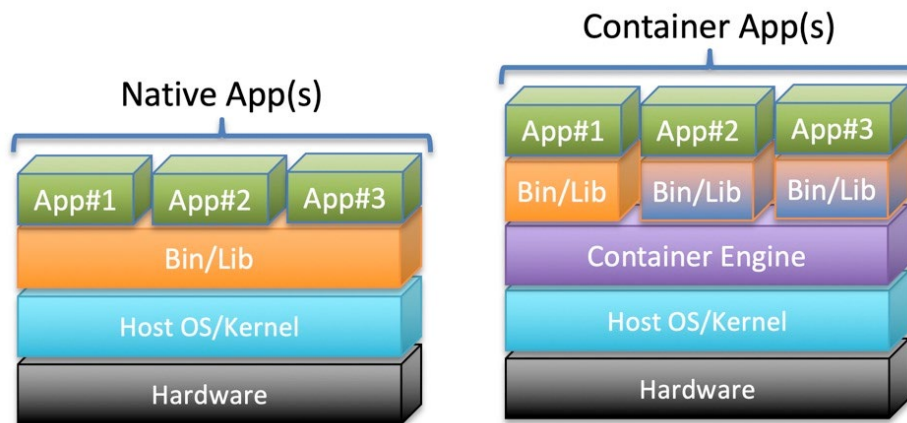
This system deals with the full lifecycle management of software components, from arranging the download, to the provisioning and running/monitoring, and eventual removal of the software within a gateway platform. LXC is used for application containers. The solution operates with existing ACS's relying on TR-069 (or in the future USP), requiring some additional capability in the ACS to help with orchestration of what SW components are position on what gateways, etc. The system provides a complete solution and is planned to be open sourced into the prpl Foundation, providing an option that can be ported to any router platform for managing software components in a consistent way. Other companies have also been working on similar approaches and the hope within the community is that we can bring multiple parties together to create a common solution.

## Container Usage - LXC or Docker/Balena

The main reason to look at containerized applications is to try and provide abstraction from the main monolithic firmware image. As mentioned, the release cadence of the main firmware images maybe too slow compared to new software feature needs. Having the ability to develop and test software that can then be deployed on top of an existing image can resolve this cadence issue. It also allows for such software to be tested against a fixed target release in the field, ensuring more confidence in the new feature when they are deployed in the field. Having independence from the monolithic image also means it's possible to upgrade such software quickly in the field in the event of issues arising, without having to perform a complete system test of the main monolithic firmware image again.

## 7. Container Choice

The choice of container system for gateway platforms needs to consider the overall Service Provide and Supplier model being used (at arms-length or tightly integrated), the available RAM/flash, and the long-term maintenance requirements. LXC, Docker, and Balena offer a nicely packaged system to manage applications, and in some cases also handle the application repository aspects as well. One other option that avoids application containers relies on the original Linux primitives that can limit resource used by applications/processes running in a system using chroot() and cgroups/namespaces. This alternative approach works as well as containers, and some operators are considering this more lightweight approach for managing applications in order to reduce the overhead associated with the other options.



**Figure 19 - Native Apps vs Container Apps**

The choice of container option will likely have an impact on the build system for the devices being used. In terms of LXC containers and cgroups/namespaces approach, it is possible to get much smaller container images as a result of reusing the available dynamic libraries within the primary firmware image root filesystem. A challenge with this approach however is the tight coupling required as a result of having to build the LXC container applications as part of the overall firmware image process. When working with internal SW teams, this is not a major issue, but there may be the usual “sharing problems” if 3rd party software companies need access to this build system.

Challenges such as the overall version of firmware image and versions of libraries contained in the root filesystem may change due to upgrades, fixes, new features, etc. and any LXC container application may be incompatible with the changes results in the need to have very careful feature and change planning in order to avoid a permanent state of development.

In addition to resource management and resource limits for new software and services, a key requirement to consider is how to interface with the main routing platform. In some cases, the integration requirements for new software can be limited to an IP and TCP/UDP port mapping, whereas other integrations need to directly interact with the local platform. Clearly defined interfaces (like all those described earlier) are a **must** to ensure coordinated access to the platform is maintained. Such interfaces enable 3<sup>rd</sup> party software providers understand how to interact with the platform, while the same interfaces provide a defined bridge that the platform software can marshal in terms of access control rights (what application can interact with what subsystem), and abstraction (allowing underlying systems to be modified while maintaining consistent northbound interface).

In terms of Docker or Balena containers, as they create their own root filesystems the resultant container sizes can be much larger compared to LXC. However, 3rd party developers need far less access to the internal build system or the firmware internal libraries, relying mostly on the target toolchain to build their applications. The approach enables independence between such containers and the system platform, but at the cost of RAM and flash resources.

In most cases the target platform for containerisation will require at least Linux 3.18, and preferably the latest kernel available

The choice of container comes down to the following options

- LXC Container
- Docker Container
- Balena Container

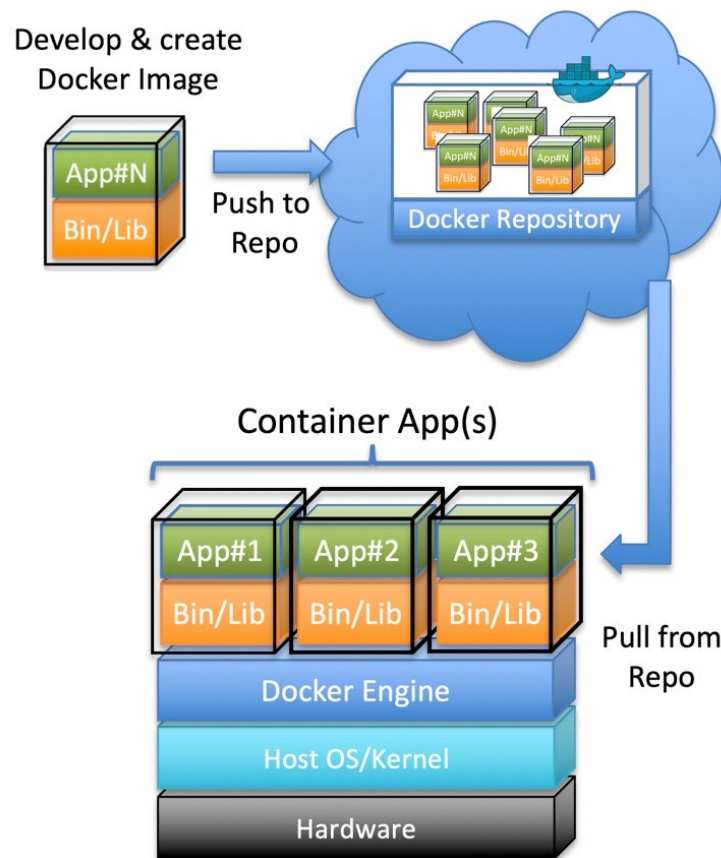
## **7.1. LXC Container**

LXC Containers provide the closest coupling and reuse of existing resources in a GW platform, and as such are definitely being considered on platforms that need some flexibility to deal with very constrained environments. They are the pre-cursor to nearly all other container systems, being a packaging up of the previously developed kernel tools developed to “contain” processes. They tend to be light weight in terms of RAM and flash, and fit into an existing platform without much overhead in terms of “container execution environments”. However, these benefits come at a cost of tight integration and reliance on a flexible firmware build system. LXC also does not have a native “orchestration” option that can be used directly, resulting in the need to create a suitable environment (the previously mentioned LCM/SDP addresses this need).

Multiple industry efforts are underway to add LXC containers to embedded GW platforms, with options being discussed with most of the Tier-1 operators.

## **7.2. Docker Container**

Docker is an open platform for developing, shipping, and running applications. Docker enables the separation of applications from a local platform to enable speedy software delivery. Docker execution environment provides isolation and security allowing multiple containers to simultaneously on a platform. Compared to virtual machines, Docker containers are much lighter, but these containers and the execution environment can be a lot heavier in resource usage compared to LXC.



**Figure 20 - Docker Ecosystem**

A major benefit of Docker is how it creates containers, where every file/library/application required is packaged into a single container image enabling it to be distributed in a highly portable fashion. As a result a Docker container can be deployed on a gateway, local laptop, physical, or virtual machine in a datacenter or in a cloud provider environment. The portability of the Docker container means that many more software providers can develop their applications to run on Docker, enabling a very rich and vibrant market space. A Docker container is a runnable instance of an application image. Like LXC it can be started and stopped using a Docker API or a CLI. Docker relies on Linux services (either natively in a Linux kernel, or through “Linuxkit”) and uses namespaces in the same way as LXC does to provide the required workspace isolation for the container to operate within. Namespaces offer process, networking, inter-process communication, mount/filesystem, and some kernel isolation.

In addition to these fundamental features, Docker introduces a whole host of extra functionality to be able to manage and interact with containers, enabling eco-systems to be built to fully manage and orchestrate the operation of large numbers of Docker images/containers over vast “fleets” of compute resources.

### 7.3. Balena Container

Balena containers are very similar to Docker, having been developed as a cut-down version of Docker. The following features have been removed from Docker Container support to create the lightweight Balena container platform, resulting in a 3.5x reduction in size:

- Docker Swarm

- Cloud logging drivers
- Plugin support
- Overlay networking drivers
- Non-boltdb discovery backends (consul, zookeeper, etcd, etc.)

Balena concentrates on using RAM and storage more conservatively and focuses on atomicity and durability of container pulling. These facets are ideal in the context of embedded systems, compared to the more traditional cloud systems that Docker is targeted at).

## Container Experience in Gateways

### 8. Containers on Commscope Gateway Platform

The Docker ecosystem is being used on some CommScope gateway devices, with a view of enabling common applications to be deployed across a range of mixed capability devices (concerning CPU, RAM and flash resources). The current management/orchestration of the Docker system on these platforms relies on the TR-157 (SMM) functionality previously mentioned as well as more explicit Docker controls. The SMM system has some roots in the Home Gateway Initiative NERG as well as previous efforts that tried to add OSGi to gateways. The TR-157 defines platform attributes as well as lifecycle management.

The CommScope gateway platform relies on the open source Docker Engine to provide the framework for hosting containers. A Docker Client is added to the gateway to manage and control access to the Docker environment. Remote Docker clients are also supported to assist with the installation of containers as well as querying status/etc.

SMM depends on Execution Environments (EE), Deployment Units (DU) and Execution Units (EU). The Docker Engine is equivalent to the EE, providing a platform for hosting applications that are effectively sandboxed to the rest of the gateway/host system. Docker Images are equivalent to the DU, providing a way of managing the specific files/etc. associated with the application being downloaded. The EU is the active running Docker Container executing within the Docker Engine/EE environment.

Docker containers are either pre-downloaded or downloaded from the Docker Registry. Interactions with the Docker Registry, including authenticating access, are all logged to ensure diagnostic information can be reviewed in the event of issues.

The current model is to use the CommScope Container API for Docker Container applications control objects on the gateway platform itself. Extensions such as providing access control to local Linux services and Dbus access are also provided. From an operational perspective, as some platforms are flash limited, the Docker Engine itself is run time installed into RAM, as are the other Docker Container images.

The running of the Docker Engine on the platform requires allocation of resources from the gateway for any Docker Applications being deployed. The current support in the gateway based Docker support is for installing, enabling, uninstalling, and disabling using either an External Docker Client or using the Docker Configuration file. The main features of the SMM are provided to report on status/etc. of the Docker Engine and running applications.

Some of the Docker applications include McAfee security gateway as well as SamKnows. Other applications are considered as well as internally developed features.

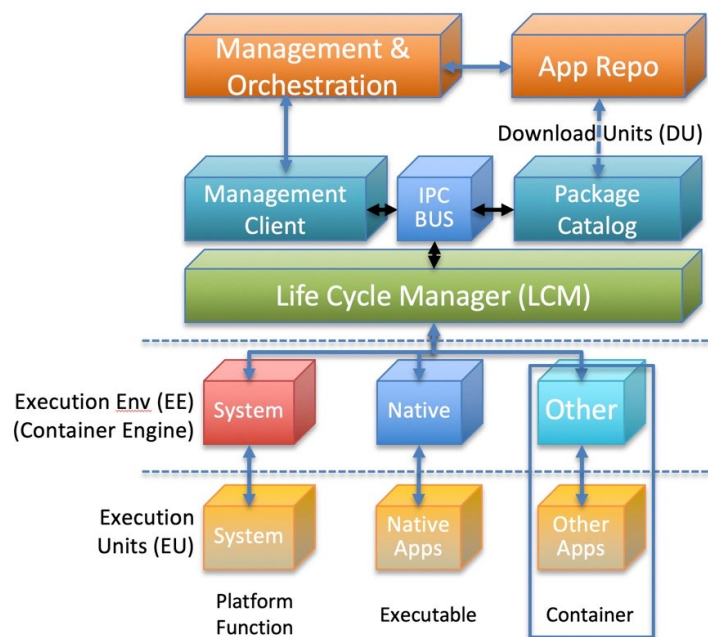


## 9. Life Cycle Management (LCM) on OpenWrt platform

A previously mentioned the LCM/SDP platform developed by Vodafone enables the management of downloadable software components in LXC containers. This was demonstrated through a POC that Vodafone developed and demonstrated at openWrt 2018.

This POC chose to use the primary features of the TR-157 Software Module Management (SMM) specification for this, providing a generic interface for this interface, allowing it to be mapped into TR-069 for ACS management or made available for other agents to use with other orchestrators. The LCM component provided external access to execute the available Life Cycle Management API methods, while also being responsible for fetching packages/containers, retrieval of information about packages from the local filesystem, as well as delivering the required applications to the Execution Environment to run.

The POC demonstrated the use of multiple Execution Environments (EE) allowing for mixed service operation. A Base EE was used to allow upgrading of specific components into the main root filesystem, that did not require any separation, such as new native images. A key feature of the Base EE was to allow direct patching of the main OpenWrt system, enabling the installation of a new native package directly into the running system. The use of the Base EE also allowed for a bit more package information to be included to be able to authenticate the packages, etc. A so-called Native EE was added to enable root filesystem separation, meaning that a new Native package would not overwrite anything in the base root filesystem, enabling isolation from the running system. The final EE was the Container EE, where new 3rd party applications needed isolation from the main system, and would be have limits placed upon all resource usage, as well as preserve system stability.

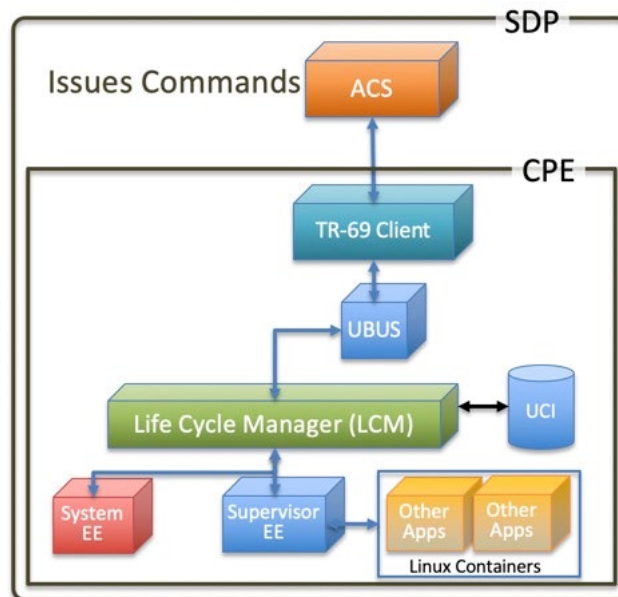


**Figure 21 - SDP/LCM and Orchestration Overview**

In all cases the LCM was responsible for managing the different EE, where it would perform actions on the EE, and deal with the returned status. Operations such as install, uninstall, start and stop were all supported.

The following diagram shows the model used in the POC:





**Figure 22 - SDP/LCM POC Configuration**

The ACS platform was used to issue commands to the CPE platform, where they were handed over to the LCM to perform all actions related to running services within the platform. The System EE listed above allowed for the ACS to request package updates (OpenWrt) to be applied to the running system, while the Supervisor EE was used to actually run the isolated applications. The Supervisor EE is responsible for handling the environment that applications run with

The Supervisor EE supported features such as package verification and install/remove, service startup/shutdown, as well as isolation (including limiting namespace, RAM and CPU). The OpenWrt Summit 2018 demonstration showed some basic containers running, as well as a more complex setup that involved Samsung SmartThings integrated in a container, downloaded into the system and using a local Zigbee USB dongle to interact with an external Zigbee lightbulb. Other aspects such as CPU resource limitation were also demonstrated. All of the interactions in the demonstration were controlled using the SMM functionality on the connected ACS.

The SDP/LCM system as currently defined delivers a complete solution for managing containers and even native applications on embedded gateways. It offers orchestration through the connected ACS (although ACS platforms probably require custom extensions to really hope to act as orchestration systems), and works on OpenWrt. Work is ongoing to get this functionality working on RDK and hopefully the overall SDP/LCM solution software will be opensourced at some stage.

## 10. Nomad POC

Nomad (from Hashicorp) is a highly scalable orchestration system that has been deployed to deal with launching 100,000s of container applications for various purposes. It manages clusters of machines and runs different types of applications on top of them, integrating with another Hashicorp product called Consul (a service discovery and configuration tool). Its primary function is to manage microservices efficiently over clusters. A Nomad POC was developed to demonstrate the management of Docker based

containers being orchestrated using the Nomad system ([www.nomadproject.io](http://www.nomadproject.io)). This POC relied on integrating Docker CE onto a platform, running a Nomad Agent container on the platform and using Nomad server to interact with the agent to orchestrate the setting up of “IOT” application container and some other basic containers.

The work began on a platform with 512MB RAM and 4GB flash. The POC team already had extensive experience with Nomad Server and were using this to understand the client side and how this would scale. The Nomad Agent (that runs on a client device, such as a gateway) includes support for so-called “Task Drivers”, allowing it to manage multiple types of execution environments, including, Docker, Isolated Fork/Exec, Raw Fork/Exec, LXC, Java, QEMU, Rkt, Custom. For the purpose of the POC, Docker was the chosen environment.

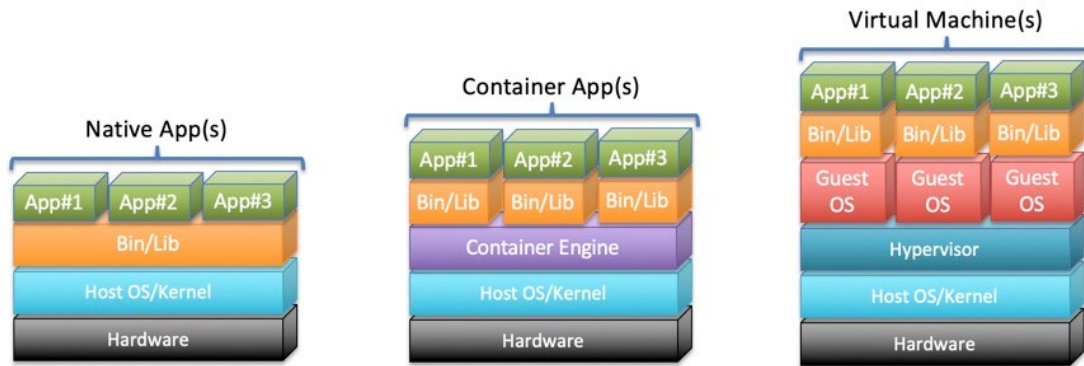
The POC demonstrated that the platform was well capable of delivering the required services. The size of available flash used was considerably more than on most embedded platforms today, where for example a gateway might only be designed with 128MB flash, which is only **3.2%** of that available in the POC platform. The POC was capable of demonstrating the use of Nomad for orchestrating the local services, as well as showing that an IOT container application was able to function and access a local Zigbee interface on the gateway.

## Alternative Virtualization Options

Containers are not the only virtualization option that exists. The following sections explore virtual machines as well as full/hybrid cloud applications.

### 11. Virtual Machines

Virtual Machines (VM) are another way of isolating software functionality to run on a platform. They require a complete running platform of software, including a complete OS. When a VM runs on a platform there are different ways it can be positioned. It can run on top of a so-called “baremetal” hypervisor (the software to manage VM access to the local HW platform, also known as a ‘type-1’ hypervisor). Systems such as VMware ESXI, or Microsoft Hyper-V server or open source KVM are all examples of a type-1 hypervisor. Alternatively, a VM can run on top of a ‘type-2’ hypervisor running in an OS on the HW platform. VMware Workstation, Oracle VM VirtualBox, and Parallels are all type-2 hypervisors. VMs provide a complete environment, meaning that in most cases they require massive amounts of RAM and storage to operate, one example of a ubuntu 18.04 desktop VM requiring 25GB of storage and 2GB of RAM. As such, VMs really don’t suit when trying to add some basic software features on top of an existing embedded system.



**Figure 23 - Native, Containers and Virtual Machines**

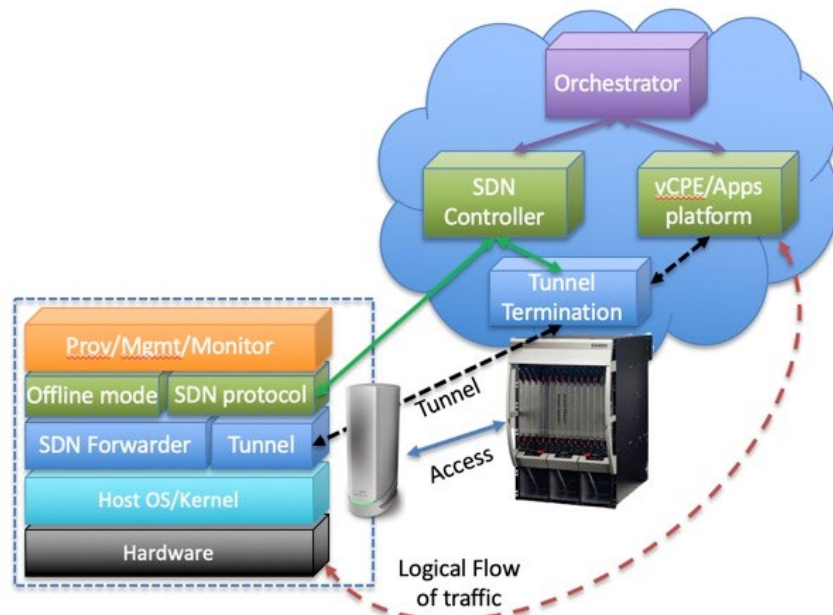
VMs are not discussed any further in this paper.

## 12. Full and Hybrid Cloud Virtualisation options

The traditional model of Virtual CPE has talked about moving the entire routing platform software out of the home and into the cloud, and having all of the network traffic hairpin through this remote virtualization platform that replicated all of the home networking functions. The hope of vCPE was to reduce the cost of the gateway device and move the SW complexity to the cloud. The approach was to enable different SW instances in the cloud that would support centralized feature development and versions, with the ability to rapidly cutover gateway devices from one version of software to another to get new features, as well as to offer new software features and services to customers, regardless of the type of gateway hardware was in each home.

This model has not really succeeded. The costs of offering the vCPE platform in the cloud while also offering gateway HW in the home never quite added up to something that was more economic than a dedicated gateway in the home. The truth is that vCPE hardware platforms, especially with Wi-Fi, are not too different in costs compared to equivalent home gateway platforms, in most cases the actual SOC is the same and offers the same MIPS processing power. The main difference would be RAM/flash costs, with a vCPE platform requiring less of both (but ironically, could be forced to buy more than required just to hit RAM/Flash price sweet spot).

So, is vCPE dead? The answer is no, as some very good pieces of vCPE can be used. The idea of isolating certain traffic flows and certain Virtual Network Function (VNF) software to the cloud is an idea that has persisted and been demonstrated to work well. In this case a traffic tunnel connects the home gateway to the remote cloud VNF, where all the hard work is performed. One of the main examples of this is “Wi-Fi Public HotSpot” services. The traffic to be tunneled is simply that traffic that operates on one of the Wi-Fi SSID that the gateway offers. Every data packet is received from the SSID and tunneled using a ‘softGRE’ tunnel to the cloud VNF. The cloud VNF terminates the tunnel, extracts the traffic and operates a Wireless Access Gateway (WAG) function, that deals with AAA and all the required traffic management (DHCP, etc.) and encapsulation/decapsulation, before dispatching the traffic off to the internet. This model is one of the first real examples of vCPE and has been widely adopted.



**Figure 24 - vCPE with Cloud Services**

However, it's a very basic option, using a coarse traffic filter (the entire SSID) to isolate traffic. The main function the gateway must provide is the ability to isolate such traffic and pack it into a SoftGRE tunnel connected to the cloud VNF, so it's definitely minimized SW complexity in the gateway.

More advanced versions of vCPE have started to be developed, using more fine-grained data plane filtering options. In a lot of cases, traffic that is filtered must be transported out to a remote cloud VNF where the actual software processing occurs, typically through a SoftGRE or equivalent tunnel. Alternative options also exist where this traffic could be handed over to a local container or software component, mixing up the different models (where it makes sense).

## Data Plane

Traditionally, Linux network tools, such as iptables have been used to manipulate traffic flows, providing low level filtering and redirection/etc. These tools are used by some of the key networking functions within the routing platform, but typically are not open to higher layer software components, as they have the potential (if used incorrectly) of wrecking the network packet forwarding of a system. No real programmatic API has been developed to expose this interface to 3<sup>rd</sup> parties. However, Software Defined Networking (SDN) does offer some new ways around this.

The basic tools of SDN, such as openFlow and Open vSwitch, have offered the ability to isolate incoming traffic flows on a platform and modify or redirect such flows for additional software processing, including forcing a flow to be sent out an interface that happens to be a tunnel or another local interface, possibly connected to a container. New software approaches for gateways are starting to reuse this type of processing.

The benefit of this model is that once the software agent is enabled on the gateway platform, then any interesting traffic flows can be dispatched via a tunnel interface to a remote cloud VNF, without requiring

new software to be added to the gateway (filtering instructions would simply be configured into the gateway depending on what traffic flows had to be isolated).

Integrating the opensource OVSDDB and Open vSwitch (OVS) into a gateway has enabled OpenSync to exert very fine-grained control over traffic passing through a platform, with the possibility of redirecting such traffic to a tunnel interface for carriage to a cloud VNF. The benefit of open source OVS is that it is possible for 3<sup>rd</sup> party software to also use the same infrastructure if required.

A similar packet interception model that embeds a NFlua component linking to the Linux Kernel network packet handling has also been developed. This has been used to provide sophisticated AI driven cybersecurity and network intelligence features for network operators. The ability to deploy an agent and then dynamically reconfigure its basic rules provides a very powerful model that allows for independent upgrades/etc. without having to involve an operator at all. Such an agent module could also be repurposed to provide a packet filtering option, like OpenFlow, to redirect traffic to a remote cloud VNF.

In these traffic interception/filtering/redirection cases, the traffic is either hair-pinned out to the cloud VNF and sent back to the gateway, is completely consumed by the remote service, or dispatched to a local agent present in the gateway that also performs processing or other software handling. Using these tools enables easier manipulation of the data plane than ever before and offers more organized control about how to isolated traffic and direct to software components (local or remote). More effort is being put in by SoC providers to ensure that hardware acceleration can also be applied to this traffic manipulation, ensuring that software can access the high speeds expected from gateways.

## Control Plane

As mentioned a lot of container systems have their own proprietary backends for controlling how containers are deployed and operated on compute platforms (e.g. gateways in the case). These tools are more concerned with treating the containers as black boxes and satisfying the “label” of resource requirements that come with the container.

In the case of the ARRIS Docker Container POC, additional supports were provided to allow the manipulation of the Docker system from a remote ACS by using TR-069 extensions mapped into the TR-157 SMM system, allowing some more native (from an operator perspective) management to be employed. Kubernetes was not used to provide orchestration in this instance.

The SDP/LCM system that Vodafone has created also uses a similar model to ARRIS, relying on TR-157 EE, EU and DUs to enable a very flexible control system for managing sophisticated software delivery options and life cycle management. This system also relies on the use of TR-069 to assist with orchestration/etc.

Nomad is another orchestration system capable of flexibly managing many different images (via nomad agent). It is capable of dealing with Balena, Docker and LXC container images, as well as many other image options. Like Kubernetes, Nomad can scale very well in a data center setting, coping with very high container deployment scenarios. However, Nomad and Kubernetes may not be able to scale to the required number of containers when deployed in an operator environment with thousands or millions of devices with multiple containers per device.

Existing ACS platforms may be able to cope with the scale of unique devices, but need additional “orchestration” extensions to be added to them. ACS platforms already deal with firmware image

management/etc., and the TR-157 SMM option extensions provide a defined model for managing the EE, EU, and DU options in a gateway.


## Concurrency and Orchestration Scalability

Orchestration must deal with a number of constraints within the embedded gateway ecosystem when used to manage the deployment of multiple applications in containers across the footprint. For one, any orchestration system must be capable of dealing with the massive numbers of deployed devices. It must also be able to cope with the potential of many different applications and application types per gateway, as well as different gateway types (varying in SoC supplier, CPU, RAM, and flash at a minimum). This just cries out for sophisticated orchestration systems that can address the multi-dimensional complexity.

What currently is not understood is the level of concurrency of applications running within gateways. What this means is whether the limited resources in a gateway are going to be under pressure if multiple applications are deployed, and if some clever orchestration technique will be required to constantly add and remove applications on demand or on a timed basis.

What is also not understood today is if operators will only allow their own curated container based software and services to run on these gateway platforms, or will decide to open up and potentially monetize the platform, allowing 3<sup>rd</sup> party applications to run, similar to the Android Play Store or IOS App Store. Given the high level interfaces and various access controls available with these, it does appear as a possibility, and may allow for hybrid mobile applications and other software services (such as IOT systems) to be developed that rely on an “always on presence” in the home rather than having to pay for high latency cloud based servers.

In terms of concurrency and high application counts, one of the easiest ways of addressing this is to basically ensure sufficient storage and memory is available in the platform. Such an approach means applications are rarely removed and replaced with other applications, thus avoiding a never ending game of Tetris that the orchestration system must play – constantly trying to fit apps into available space. This does at a slight cost of extra storage (the RAM can be freed up if an app is no longer active) but removes the need for a complex orchestrator.



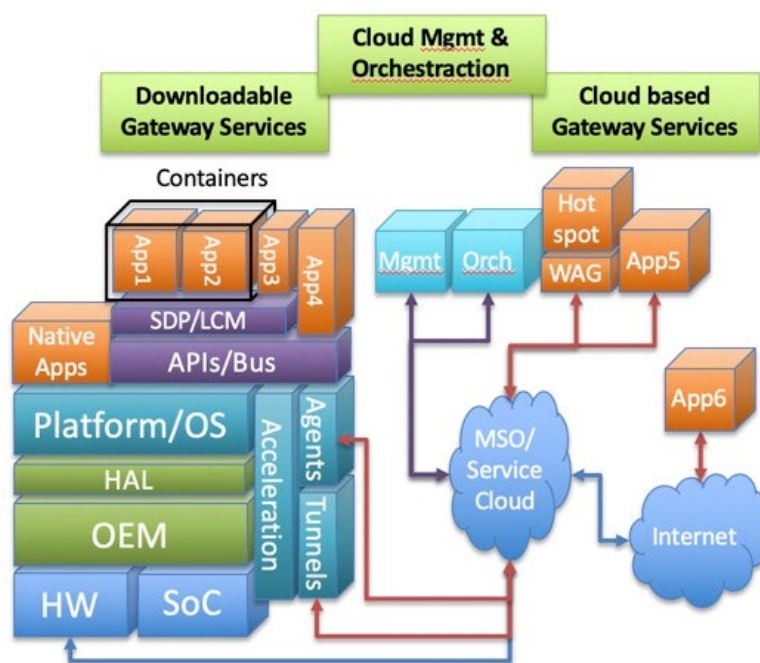
**Figure 25 - Potential Service Load over 24hr Period**

If the orchestration complexity outlined can be removed (through extra storage/etc.), then it's quite feasible to believe that the existing NMS/ACS systems that already manage vast numbers of broadband devices should be capable of supporting the required orchestration function. Existing TR-069 systems maybe usable, but the upgraded Universal Services Platform (USP/TR-369) protocol from the Broadband



Forum maybe be better suited, given the new features it brings to managing broadband devices, as well as its backward compatibility to existing data models like TR-181. USP modifies the transport protocol in use, providing a faster and more scalable link between the ACS and device populations. CommScope has recently opensourced a complete USP agent implementation that is available for integration with existing gateways to enable this new functionality.

The days of all software being delivered as a monolithic firmware image are numbered. The availability of all the required elements to create new portable software is very encouraging. The new dataplane and control plane options enable application developers (including ISVs, open source developers and the MSO community) the option of creating new applications not considered before. Along with the system high level and low-level APIs, developers are able to bundle all their required libraries and executables within a container based system (be it LXC, Docker, Balena or others), and have these orchestrated on to gateway platforms. The addition of LCM/SDP as well as reuse of Docker/Kubernetes, Nomad, or TR-069/USP based orchestration systems will enable cable operators more control over what to deploy and when/how to deploy.



**Figure 26 - Multitude of Options for Virtualised CPE**

Right now there are a multitude of Docker based container applications, while only a few 3<sup>rd</sup> party container applications have been totally focused on embedded gateways. Expect this to change very soon as hardware profiles change and the various software layers and interfaces are developed and adopted in the multiple routing platforms that exist in the embedded broadband gateway world.

# Conclusion

The constant change in the broadband gateway space is driving demand for newer software services at an unprecedented rate. As a result, the complex gateway platforms need to innovate faster at the software architecture level and hardware must keep in step to match the software needs.

Native application and container applications need to take advantage of new APIs, HALs and Service Delivery Platforms that are emerging to ensure fast adoption on to gateway platforms.

SDP/LCM and Docker are good options to consider for container deployments, with other platforms like Nomad also to be considered in this space. However, orchestration systems that can manage the scale of broadband gateway deployments and mixed deployed services have not been realized yet, resulting in the potential use of existing or future ACS (TR-069/USP based) to handle this workload.

Getting these new software services into gateways is essential for MSOs to entice and retain subscribers



# Abbreviations

ACS	Auto Configuration Server
AP	access point
API	Application Programming Interface
AR	augmented reality
BSP	board support package
CDN	Content Delivery Network
CLI	Command Line Interface
CPE	Customer Premise Equipment
CPU	Central Processing Unit
DOCSIS	Data Over Cable Service Interface Specification
EE	execution environment
EPON	Ethernet Passive Optical Network
EU	execution unit
GPIO	General Purpose IO
GW	gateway
HAL	Hardware Abstraction Layer
ISBE	International Society of Broadband Experts
IoT	Internet of Things
LCM	life cycle management
LED	light emitting diode
LXC	linux containers
MQTT	message queuing telemetry transport
MSO	Multiple System Operator
MoCA	Multimedia over Coax Alliance
NFV	Network Function Virtualization
NFVO	NFV Orchestration
OEM	Original Equipment Manufacturer
OS	Operating System
OVS	Open vSwitch
OVSDDB	Open vSwitch Database
PRPL	prpl Foundation
QoE	Quality of Experience
QoS	Quality of Service
RAM	Random Access Memory
SCTE	Society of Cable Telecommunications Engineers
SDK	Software Development Kit
SDN	Software Defined Networking
SDP	Service Delivery Platform
SOC	System on Chip
UBUS	OpenWrt micro bus
UCI	Unified Configuration Interface
USB	universal serial bus
USP	user services platform
VM	virtual machine
VR	virtual reality
WAG	Wireless Application Gateway
vCPE	Virtual CPE

## Bibliography & References

Docker Information; <https://www.docker.com/>

RDKit-B Architecture; <https://wiki.rdkcentral.com/download/attachments/23593288/arch.png>

Prpl Foundation; <https://prplfoundation.org/working-groups/prplwrt-carrier-feed/>

OpenSync; <https://www.opensync.io/documentation>

Safe browsing using Lua; <https://www.lua.org/wshop17/Lourival.pdf>

prpl Service Delivery Platform; [https://openwrtsummit.files.wordpress.com/2018/11/sdp-openwrtsummit2018\\_1-3-1.pdf](https://openwrtsummit.files.wordpress.com/2018/11/sdp-openwrtsummit2018_1-3-1.pdf)

# Detecting Video Piracy with Machine Learning

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

**Thomas Belford**

Software Engineer – Technology Department  
NCTA – The Internet & Television Association  
Washington, DC  
thomasbelford32@gmail.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Content .....	5
1. Deeper Look at Video Piracy Traffic .....	5
2. Machine Learning .....	9
2.1. Overview .....	9
2.2. Machine Learning Algorithm .....	10
2.3. Data Features.....	11
2.4. Data Sets .....	12
2.5. Machine Learning Model Performance.....	12
2.5.1. Machine Learning Metrics .....	13
2.5.2. Machine Learning Model Performance Results.....	14
2.6. Machine Learning with NetFlow .....	15
3. Case Study .....	16
3.1. Implementation.....	16
3.2. Case Study: Enterprise.....	17
3.3. Case Study: Cable Operator.....	18
4. Applications .....	20
Conclusion .....	21
Abbreviations.....	22
Bibliography .....	23
Endnotes.....	25

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 Common Pirated Video Flow Sequence .....	6
Figure 2 Example Netflix(Blue) vs Pirate IPTV (Red) Feature Comparison .....	7
Figure 3 Packet Sizes per Flow Comparison of Netflix(Blue) vs Video Piracy(Red).....	8
Figure 4 Feature Comparison of Pirate (Red) vs Internet (Blue) .....	9
Figure 5 Concept of Machine Learning.....	10
Figure 6 Machine Learning Video Piracy Detection System.....	16
Figure 7 Video Piracy Traffic Classification System.....	21

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 Pirate IPTV Providers Analyzed .....	5
Table 2 IP Flow Data Features of Interest.....	6
Table 3 Features .....	11
Table 4 - Data Sets.....	12

Table 5 Machine Learning Algorithms Performance Results.....	15
Table 6 Enterprise Capture File Statistics .....	17
Table 7 Enterprise Traffic with Full Feature Set.....	17
Table 8 Enterprise Traffic with Meta-Data Only Feature Set.....	17
Table 9 Top IP-Flows Identified with Random Forest and Full Feature Set.....	18
Table 10 Cable Operator Capture File Statistics.....	19
Table 11 Cable Operator Traffic with Meta-Data Only Feature Set .....	19
Table 12 Top Labeled Pirate Hosts for Cable Operator Case Study .....	19

# Introduction

The broad adoption of broadband internet and growth in average internet speed [1] has fueled the streaming video industry. In turn, the growth and popularity of streaming video has also fueled the growth of video piracy [2].

Video piracy is a form of copyright infringement and refers to the use of works protected by copyright law without permission for usage where such permission is required. There are two primary forms of video piracy. The first form commonly referred to as “video-on-demand” (VOD) uses a file sharing distribution model and is commonly used by applications such as Kodi, Titanium TV, TVZion and BitTorrent based applications [3].

The popularity of streaming video has resulted in the creation of illegal virtual cable operators selling subscription based over-the-top IPTV, complete with electronic programming guides, that stream multiple channels of linear video. This second form is known as “pirated linear streaming”.

Pirated linear streaming is a business threat to the pay-TV industry as the pirated linear streaming product is a good substitute for legitimate pay-TV services. For the pay-TV industry, one of the issues is understanding the true scope of the problem. There are some industry reports [4] that estimate that 5.5% of North American households are accessing pirated content. The pay-TV industry has been trying to better quantify the problem, as part of determining what actions to take to mitigate it.

To truly understand the scope and scale of video piracy, operators need to measure the volume, frequency and scope of traffic on their networks that is associated with pirated linear streams. Pirated streams use the same technologies and streaming protocols (HLS and MPEG/DASH) as legal linear streams making it difficult to distinguish the two without the use of deep packet inspection (DPI). Even with DPI, it is still difficult due to multi-tenant hosts, content delivery networks, multiple IP addresses being associated with the content sources, and the diverse demographics across the footprint of the network.

Due to a number of reasons including cost and privacy concerns, operators typically have only equipped a small portion (e.g. < 10%) of their network with DPI, if at all. In addition, collecting video piracy data using DPI from a small number of points on the network can lead to a selection bias due to the demographic makeup of the network footprint.

To effectively measure video piracy on broadband networks requires something other than DPI. An approach using available IPFIX/NetFlow data, which is embedded in most carrier-grade routers and switches, provides a cost-effective approach to measuring traffic across an entire network.

In 2016 Cisco [5] showed that by using IP flow data fields it was possible to create a feature set for machine learning that used an L1-logistic regression model with an accuracy of 99.978% at 0.00% false discovery rate (FDR) to identify malware – encrypted and non-encrypted. In 2018, Cisco [6] introduced an enhanced version of NetFlow, Encrypted Traffic Analytics (ETA), that included these additional IP flow data fields to a number of its products as part of a cybersecurity solution and open-sourced the code<sup>1</sup> that captures, extracted, and analyzes network flow data and interflow data that includes the additional IP flow data fields.

In this paper, we look at applying a similar supervised machine learning process using IP flow data to assess the viability of using machine learning and IP flow data to detect pirated linear streaming traffic on broadband networks.

# Content

## 1. Deeper Look at Video Piracy Traffic

Building upon the work by Anderson and McGrew, which showed machine learning's capability to detect the unique signatures of malware, we began by inspecting both legal and illegal IP video streams. Using machine learning, we tried to identify the features that have the most discriminatory power, knowing that these features will be able to uniquely identify pirated from non-pirated traffic.

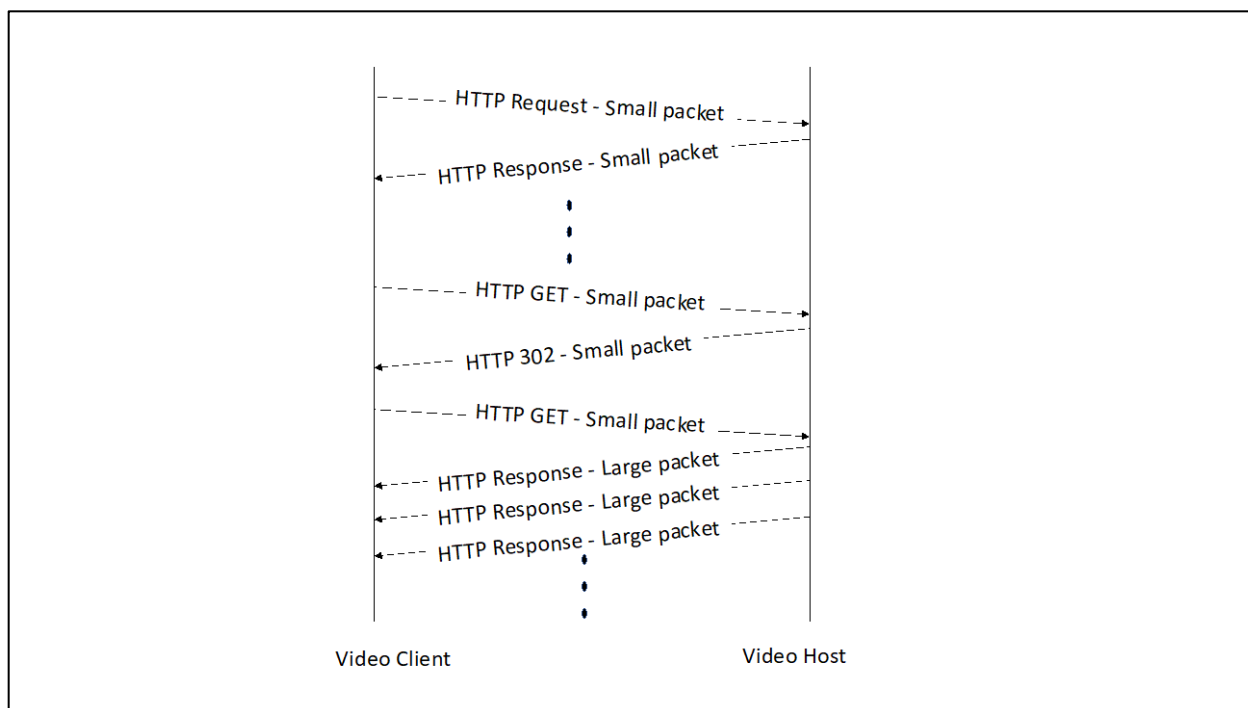
In today's video pirate ecosystem, most providers are using a client/server software pair from a small number of providers.<sup>2</sup> This fact coupled with the long-tailed nature of linear streaming video makes it possible to identify a set of features with strong discriminatory power.

For this study we captured packet capture files from a set of known pirate subscription sites as listed in Table 1.

**Table 1 Pirate IPTV Providers Analyzed**

Pirate IPTV Provider	Homepage URL
<b>IPTV Shop</b>	<a href="https://iptv.shop">https://iptv.shop</a>
<b>Excursion TV – Premium USA IPTV</b>	<a href="https://www.excursion-tv.com">https://www.excursion-tv.com</a>
<b>IPTV Choice</b>	<a href="https://iptvchoice.com">https://iptvchoice.com</a>
<b>Easy Expat IPTV</b>	<a href="https://easyexpatiptv.org">https://easyexpatiptv.org</a>
<b>Gears TV HD</b>	<a href="https://www.gearstvhd.com">https://www.gearstvhd.com</a>
<b>Necro IPTV: IPTV</b>	<a href="https://necroiptv.com">https://necroiptv.com</a>
<b>Nitro IPTV</b>	<a href="https://www.iptvnitro.com">https://www.iptvnitro.com</a>
<b>SoftIPTV</b>	<a href="https://www.softiptv.com">https://www.softiptv.com</a>

For all the pirate subscription services a common flow session is as shown in Figure 1.



**Figure 1 Common Pirated Video Flow Sequence**

Video streaming typically includes one or more long-tailed flows, that are initially preceded by a series of small HTTP transactions or short-tailed flows where the video client is logging into the back-end, followed by a the video client selecting a channel and the backend server sending an HTTP redirect to redirect the video client to the location of the video which may either be on a dedicated server or on a content delivery network (CDN).

The linear streamed video is delivered using either the HTTP Live Streaming (HLS) [7] protocol or the MPEG DASH [8].

For each of the packet captures, we looked at the flows. A flow is defined as the traffic between the 4-tuple (source & destination address, source & destination ports) and their data features. We compared the data features of the pirate linear streaming traffic to the data features of other forms of streaming video and benign internet traffic to evaluate which should have the most discriminatory power when used in a machine learning model. The data features we studied are listed in Table 2.

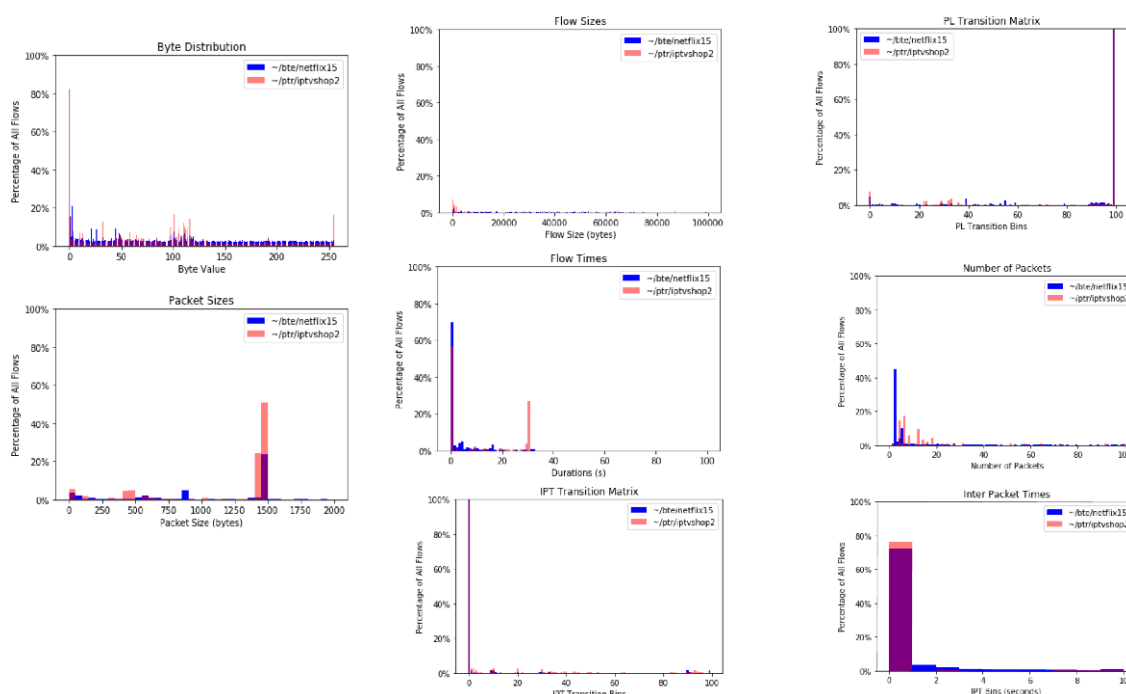
**Table 2 IP Flow Data Features of Interest**

Data Feature	Description
Packet sizes per flow	The size of each packet in the flow
Number of packets per flow	The number of packets in the flows
Number of bytes per flow	Total number of payload bytes in the flow




<b>Flow duration</b>	Time in seconds of the flow from the TCP SYN to the TCP FIN
<b>Inter-packet time per flow</b>	The number of milliseconds between each packet in the flow
<b>Source Port</b>	The source port of the flow
<b>Destination Port</b>	The destination port of the flow
<b>Byte Distribution</b>	The frequency of occurrence of the byte values in the first “n” packet payload of the first packet of the flow

We first looked at how the pirated linear streaming services compared to some of the more popular over-the-top (OTT) video services – Youtube, Netflix and Twitch. Figure 2 shows the histograms for the features listed in Table 2, and shows that the video piracy has a number of features that have distributions that differ from Youtube. Figure 3 shows a larger version of the histogram for the packet sizes in the IP flow. Even though both are forms of long-tail video, as can be seen in the figures, there are still distinct data features that emerge. The same is true when video piracy is compared to other forms OTT video.



**Figure 2 Example Netflix(Blue) vs Pirate IPTV (Red) Feature Comparison**



**Figure 3 Packet Sizes per Flow Comparison of Netflix(Blue) vs Video Piracy(Red)**

Next, we compared the pirate video traffic to a collection of internet traffic. The internet traffic collection contains captures for multiple forms of short-tail internet traffic including web browsing, webmail, mobile phone, and cloud storage. As shown in the Figure 4, just as when the pirate video traffic is compared to the OTT video, the pirate video traffic has unique characteristics that make its data features unique when compared to the internet traffic collection.

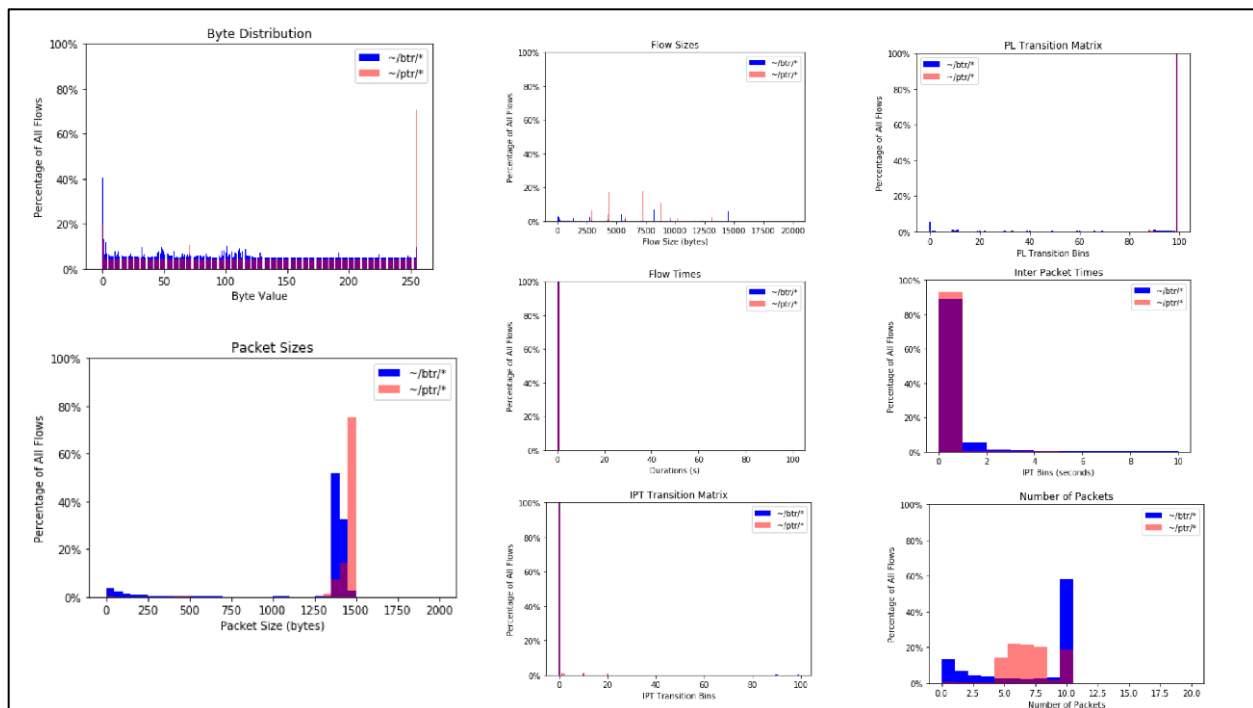


Figure 4 Feature Comparison of Pirate (Red) vs Internet (Blue)

## 2. Machine Learning

### 2.1. Overview

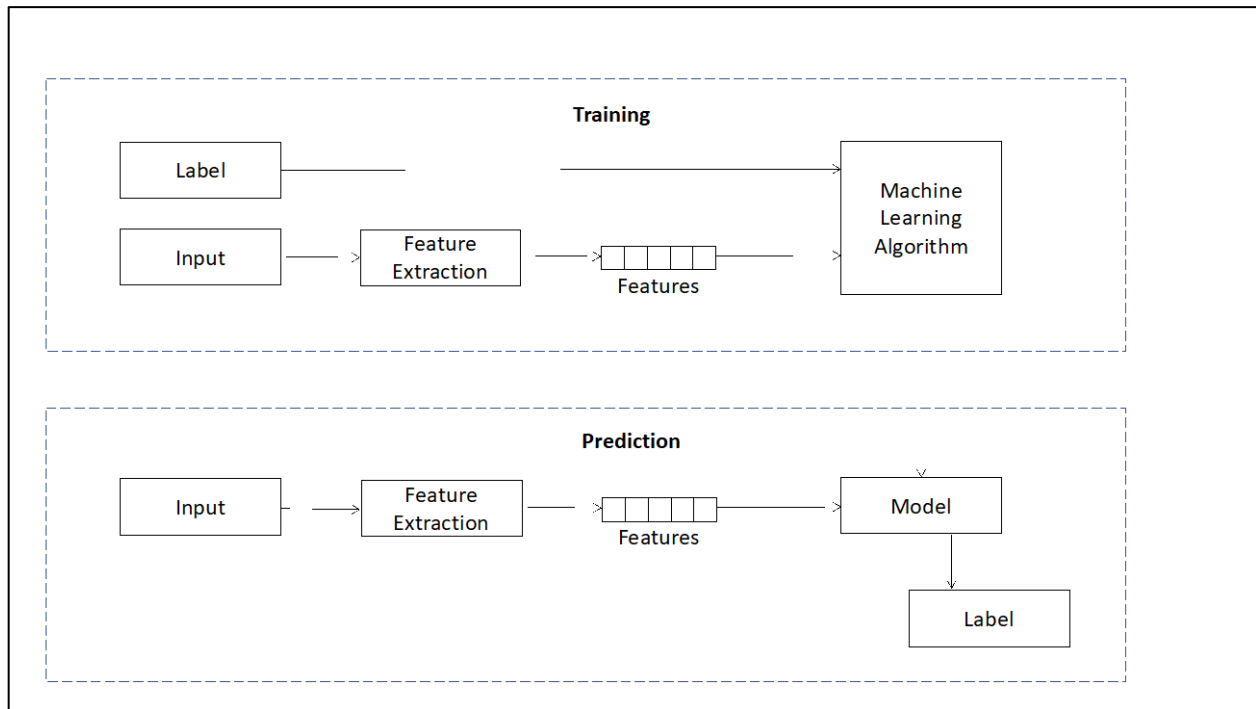
Machine learning is a subset of Artificial Intelligence (AI) and uses algorithms to discover patterns in data and constructs mathematical models using these discoveries. The models can then be used to make predictions on future data. Using machine learning to perform traffic classification is not new [9]; however, the use of IP flow data and NetFlow as the transport mechanism and synthesizing features from the flow data is new. Machine learning can either be supervised or unsupervised. Supervised machine learning trains or teaches the machine using data that is labeled with the correct answer, e.g. pirated or not, while unsupervised machine learning trains the machine using information that is neither classified nor labeled and allows the algorithm to act on the information without guidance.

Because of this, we embraced supervised machine learning as the best way to use previously observed video piracy to detect video piracy. Further, a supervised machine learning classifier provides the most direct way to build a detector, and it can also provide a probability estimate.

Machine learning makes use of the following terminology:

- **Machine Learning Algorithm** – Machine learning algorithms build the mathematical model based on the ‘training data’. There are a number of machine learning algorithms, including Logistic Regression, Decision Trees, and Random Forest.
- **Model** – A model is a specific representation learned from data by applying some machine learning algorithm.
- **Feature** – A feature is an individual measurable property of data. A set of numeric features can be conveniently described by a **feature vector**. Feature vectors are fed as input to the model.

- **Feature Extraction** – Feature extraction refers to the method and process of constructing data from the initial raw set of data to a more manageable group for processing.
- **Label** – A label is the value to be predicted by the model.
- **Training** – The process of creating a model or classifier using a machine learning algorithm.
- **Prediction** – Predicted output from a set of inputs



**Figure 5 Concept of Machine Learning**

## 2.2. Machine Learning Algorithm

Detecting video piracy is a classification problem, and we are treating it as a binary classification problem as we want to predict if the input is pirated video or not. We looked at three different machine learning algorithms to assess which model performed the best for classifying video piracy using flow data.

- **Logistic Regression** – Logistic Regression is a predictive analysis classification algorithm and based on the concept of probability. It is used to assign observations to a discrete set of classes and transforms its output using a logistic sigmoid function to return a probability value.
- **Decision Tree** – Decision tree builds classification into a form of a tree structure. It breaks down a dataset into smaller and smaller subsets while at the same time an associated decision tree is incrementally developed. Decision tree uses entropy and information gain to construct a decision tree.
- **Random Forest Classification** – Random Forest Classifier is an ensemble-based algorithm which is comprised of  $n$  collections of de-correlated decision trees [10]. Random Forest uses multiple trees to compute majority votes in the terminal leaf nodes when making a prediction.

We also briefly looked into Neural Networks and Support Vector Machines, however, the size of our data and feature sets were not optimal for these models, so we decided to disclude them from the study. Of the

three we studied in depth, Logistic Regression is computationally the most efficient, while Decision Tree and Random Forest are often more accurate while being computationally more intense.

## 2.3. Data Features

Feature selection is most important step in machine learning. For the feature selection, we looked at the data features available in the flow data as shown in Table 3. The flow data we looked at could either be extracted from IPFIX/NetFlow v9 or from NetFlow enriched with a set of user defined fields that are available in a proprietary form of NetFlow from Cisco or from open source flow data feature extractor called “Joy”.

**Table 3 Features**

Field	Source	Description
<b>IP Protocol</b>	NetFlow v9	TCP or UDP
<b>Source Port</b>	NetFlow v9	Source port in the IP header
<b>Destination Port</b>	NetFlow v9	Destination port in the IP header
<b>Flow Length, Bytes</b>	NetFlow v9	Number of bytes for the TCP connection
<b>Flow Duration, Seconds</b>	NetFlow v9	Duration of the flow from TCP SYNC to TCP FIN
<b>Packets/Flow</b>	Enhanced NetFlow	Number of the packets for the IP flow
<b>Sequence of Packet Lengths and Times (SPLT)</b>	Enhanced NetFlow	An array of LENGTH values followed by an array of INTER-ARRIVAL TIME values describing the first N packets of a flow that carry application payload. Each LENGTH is encoded as a 16-bit integer to form a 20-byte array. Immediately following this, each INTER-ARRIVAL TIME is encoded as a 16-bit integer to form another 20-byte array.
<b>Byte Distribution</b>	Enhanced NetFlow	A histogram giving the frequency of occurrence for each byte value or (range of values) in the first N bytes of application payload for a flow. Each “frequency of occurrence” is represented as a 16-bit integer.
<b>Initial Packet Data (IPD)</b>	Enhanced NetFlow	The content of the first packet of this flow that contains actual payload data, starting at the beginning of the IP header.

Joy<sup>3</sup> was used to extract the flow data from either from packet captures or collected live as flow data records (i.e. NetFlow) into a JSON format as part of the feature extraction. Joy models **packet lengths**

and **packet inter-arrival times** as Markov chains (a sequence of possible events) and excludes TCP retransmissions. The packet length is the payload length of the packet. Inter-arrival times have milli-second resolution.

For both the lengths and times, the values are discretized into equally sized bins. The length data Markov chain has 10 bins of 150 bytes. The timing data Markov chain uses 50 millisecond bins and 10 bins for 100 total features. The Markov chains are transformed into their transition probabilities are used as the features.

Joy extracts the **byte distribution data** from the flow data and generates a 256-byte distribution probability from the 256-byte distribution array for the feature.

Additional **IP meta-data** is also included in the features. The IP meta-data includes source port, destination port, number of packets in, number of packets out, number of bytes in, number of bytes out, and flow duration.

## 2.4. Data Sets

To analyze the models, we created four data sets. A pair of training data sets and a pair of test sets were created using a packet capture program. To train the model on the pirate video traffic that we wanted to identify, we used a packet capture with a set of long-tail flows between two video pirate hosts. To train the model on traffic we did not want to classify or label as benign traffic, we used a packet capture file with a mix of internet browsing, legal OTT video, email, and other enterprise traffic. To test the model we create a second, unique pair of packet capture files. For the benign traffic we simply used a packet capture from an enterprise network that was known not to include video piracy. For the video piracy test file, we performed a packet capture that included a video piracy session between the pirate video client and the pirate video server.

**Table 4 - Data Sets**

Data Set	Number of Flows	Description
<b>Benign Training</b>	54,726	Enterprise traffic, cloudfront, twitchtv, webmail, web browsing, akamai CDN traffic
<b>Benign Test</b>	14,768	Google search, Netflix, TwitchTV15
<b>Piracy Training</b>	94,742	Expat IPTV channel 3; Gears IPTV HBO; IPTVChoice NBC, PrimeAtlantic, ESPN; IPTVShop 3 &4; Unlock
<b>Piracy Test</b>	3,611	Gears IPTV, Necro, and Vaderstreams

## 2.5. Machine Learning Model Performance

We used the four data sets to calculate the classification metrics for different combinations of machine learning algorithms and feature sets. To evaluate the performance of the machine learning models we used six standard machine learning metrics. Our goal was to minimize false positives while maximizing

accuracy and precision since it would be better for a cable operator to miss a few pirated flows as opposed to classifying all pirated flows while also misclassifying a lot of benign flows as pirated.

### **2.5.1. Machine Learning Metrics**

For each algorithm we used the classification metrics – accuracy, true and false positive, precision, F1, and Log Loss – to determine which algorithm worked best for classifying video piracy.

#### **2.5.1.1. Accuracy**

Accuracy is the ratio of correct predictions to the total number of input samples.

$$Accuracy = \frac{\text{Number Correct Pirated \& Benign Predictions}}{\text{Total Number of Predictions Model}}$$

Accuracy works well if there are an equal number of samples belonging to each class.

#### **2.5.1.2. Precision**

Precision is the number of correct positive results divided by the number of positive results predicted by the classifier and is intuitively the ability of the classifier not to label as positive a sample that is negative.

$$Precision = \frac{\text{True Pirated}}{\text{Actual Results}} = \frac{\text{True Pirated}}{\text{True Pirated} + \text{False Pirated}}$$

#### **2.5.1.3. True Positive Rate (Sensitivity)**

True Positive Rate, also known as Recall, corresponds to the proportion of positive data points that are correctly considered as positive, with respect to all data points. True Positive Rate provides a measure of how sensitive the classifier is, and how well it is at not missing actual positives.

$$\text{True Positive Rate or Recall} = \frac{\text{True Pirated}}{\text{Predicted Results}} = \frac{\text{True Pirated}}{\text{False Benign} + \text{True Pirated}}$$

where,

True Positive is the number of cases in which the model predicted YES and the actual output was also YES.

False Negative is the number of cases in which the model predicted NO and the output was YES.

#### **2.5.1.4. False Positive Rate (Fall-Out)**

False Positive Rate or Fall-Out corresponds to the proportion of negative data points that are mistakenly considered as positive, with respect to all negative data points. False Positive provides a measure of the classifier's probability of falsely rejecting the null hypothesis.

$$\text{False Positive Rate} = \frac{\text{False Pirated}}{\text{False Pirated} + \text{True Benign}}$$

where,

False Positive is the number of cases in which the model predicted YES and the actual output was NO.

True Negative is the number of cases in which the model predicted NO and the actual output was NO.

#### **2.5.1.5. F1 Score**

F1 is used to measure a test's accuracy and is the harmonic average between precision and recall. It tells how precise the models classifier is, as well as how robust it is. The greater the F1 score the better the performance of the model.

$$F1 = 2 * \frac{1}{\frac{1}{\text{precision}} + \frac{1}{\text{recall}}}$$

#### **2.5.1.6. Logarithmic Loss**

Logarithmic Loss, works by penalizing false classifications.

$$\text{Logarithmic Loss} = -\frac{1}{N} * \sum_{i=1}^N \sum_{j=1}^M y_{ij} + \text{Log}(p_{ij})$$

Where,

$y_{ij}$ , indicates whether the sample I belongs to class j or not

$p_{ij}$ , indicates the probability of sample I belonging to class j.

### **2.5.2. Machine Learning Model Performance Results**

We analyzed the three machine learning algorithms and the generated models using three different combinations of feature sets to determine the model with the best performance. We used the same training and test data sets as described in Table 4 with each. Table 5 shows the results. As shown in table 6, it can be seen that the Random Forest using the largest feature set has the best overall performance. The Random Forest with the full feature set has both a high accuracy and low false positive. The results also show that Random Forest with just meta-data will still identify a large percentage of the video piracy, but may have a high false positive rate that needs to be filtered out with additional post processing.



**Table 5 Machine Learning Algorithms Performance Results**

Feature Set – Flow Meta Data						
	Accuracy	Precision	F1 Score	Log Loss	True Positive Rate	False Positive Rate
<b>Logistic Regression</b>	21%	20%	33%	27	98.4%	80%
<b>Decision Trees</b>	82%	52%	67%	6.36	98.8%	48%
<b>Random Forest</b>	81%	51%	67%	6.6	98.8%	49%

Feature Set – Flow Meta Data + Byte Distribution						
	Accuracy	Precision	F1 Score	Log Loss	True Positive Rate	False Positive Rate
<b>Logistic Regression</b>	21%	20%	33%	27	98.4%	80%
<b>Decision Trees</b>	99%	97%	97%	.45	99%	3%
<b>Random Forest</b>	99.5%	99.5%	98%	.15	99.5%	0.5%

Feature Set – Flow Meta Data + Byte Distribution + Packet Lengths + Packet Timing						
	Accuracy	Precision	F1 Score	Log Loss	True Positive Rate	False Positive Rate
<b>Logistic Regression</b>	21%	19%	33%	27	98.4%	80%
<b>Decision Trees</b>	96%	96%	88%	1.5	96%	3.7%
<b>Random Forest</b>	97%	99.8%	93%	.90	97%	0.19%

## 2.6. Machine Learning with NetFlow

On large networks, such as those operated by cable operators, NetFlow is configured to operate in sampled mode. Sampled NetFlow means the NetFlow exporter (i.e. router or switch) is sampling every  $n^{\text{th}}$  packet to update the flow state tables. Large operators configure  $n$  to be anywhere from 1,000 to 4,000. In other words, the NetFlow exporter is sampling every 1,000<sup>th</sup> packet or every 4,000<sup>th</sup> packet.

As the goal here was to determine the feasibility of using machine learning with flow data for piracy detection on large networks, we needed to analyze the impact of using sampled NetFlow for the flow data and feature extraction; however, due to time constraints, we weren't able to calculate the performance analysis with sampled NetFlow.

Despite the fact that unsampled NetFlow provides only the meta-data features, performance metrics indicate that by using the meta-data the model can identify and classify video piracy, but with a high false positive rate. We expect that reducing the flow data rate from every packet to every  $n^{\text{th}}$  (i.e. 1,000<sup>th</sup> or

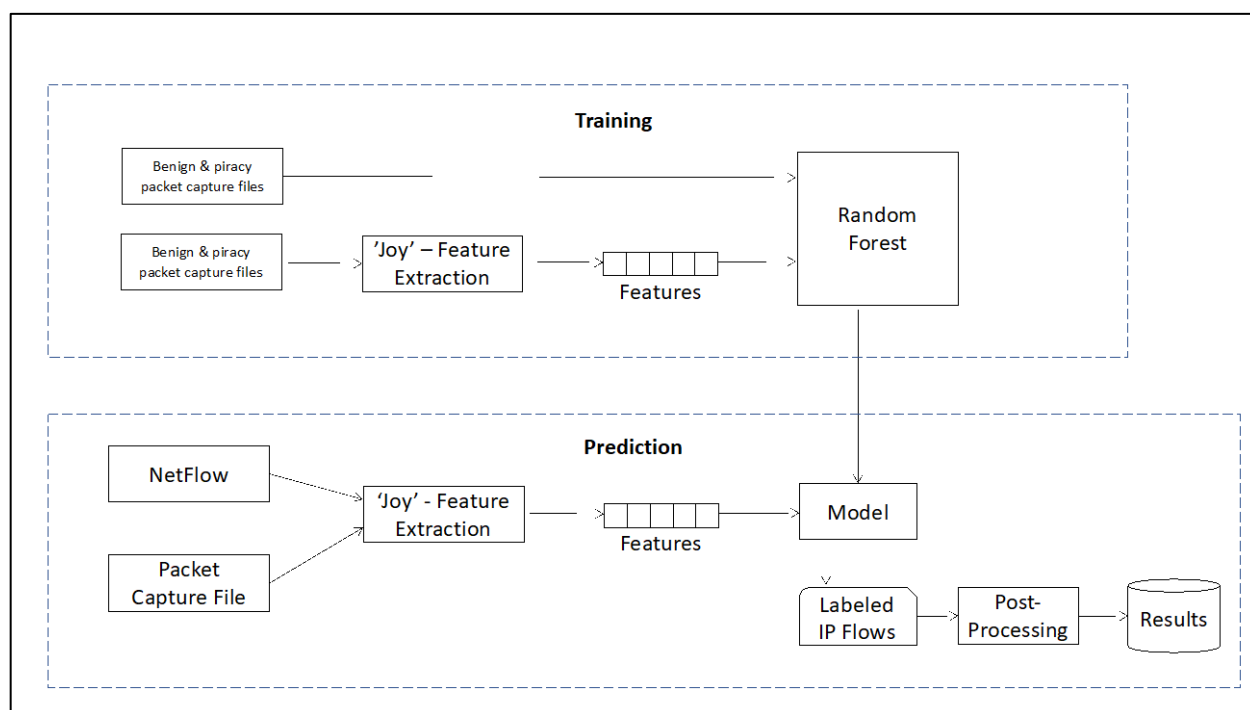
4,000<sup>th</sup>) will further degrade performance, causing the feature extractor to take longer to assemble a flow with enough information associated with it to be positively classified.

### 3. Case Study

We performed two case studies to evaluate how well the machine learning classifier works in the real world. The first case study used data collected from an enterprise with about 100 employees. The second case study used NetFlow data from a North American cable operator.

#### 3.1. Implementation

For the case studies, we implemented the system shown in Figure 6 Machine Learning Video Piracy Detection System.



**Figure 6 Machine Learning Video Piracy Detection System**

For the system we used Joy for the flow data feature extraction. Joy supports processing both packet capture files for offline processing and NetFlow flow records for on-line or live data processing. We implemented the machine learning algorithm and model using SciKit [11]. Scikit provides a number of tools that simplify the data analysis and include pre-built machine learning algorithms including Logistic Regression, Decision Trees, and Random Forest. The machine learning model labels the flows with the probability that the IP flow is video piracy. The labeled results are post-processed or filtered using a whitelist to remove any false-positives. The final result is stored in a database for report generation. The whitelists contains hosts that are well known sites that aren't sourcing video piracy (e.g. Netflix, Amazon, YouTube)

### 3.2. Case Study: Enterprise

An hour long packet capture file was captured from a small enterprise network that was known not to have any video piracy traffic. While performing the packet capture we introduced both pirated video and legitimate OTT video. The pirate video traffic included traffic from a IPTV set-top box connected to a pirate IPTV provider and multiple free IPTV sites that restream linear video. The legitimate video traffic included streamed video from NetFlix and ESPN.com. The capture file had the characteristics shown in Table 6.

**Table 6 Enterprise Capture File Statistics**

<b>Capture File Size, bytes</b>	16 GBytes
<b>Number of Flows</b>	13,503
<b>Time Duration</b>	69 minutes

We performed two tests. We ran the packet capture file through the model we generated with the Random Forest algorithm the first time using the full feature set (meta-data, byte distribution, packets lengths, and packet timing) and then a second time using just the meta-data feature set to compare the two results and to give us a some kind of baseline for how the model should work with the sampled NetFlow in the second case study.

**Table 7 Enterprise Traffic with Full Feature Set**

	<b>Labeled Piracy</b>		<b>Labeled Benign</b>		<b>Total</b>
<b>Unique end-points</b>	249	3%	8453	97%	8,702
<b>Number of flows</b>	2,269	1%	345,596	99%	347,865
<b>Number of bytes</b>	7,353,501	0.47%	1,542,729,091	100%	1,550,082,592

**Table 8 Enterprise Traffic with Meta-Data Only Feature Set**

	<b>Labeled Piracy</b>		<b>Labeled Benign</b>		<b>Total</b>
<b>Unique end-points</b>	663	7%	8388	93%	9,501
<b>Number of flows</b>	13,503	4%	334,362	96%	347,865
<b>Number of bytes</b>	115,862,972	7.47%	1,434,219,620	93%	1,550,082,592

Consistent with our performance metrics, the Random Forest model generated using the full feature set performed better than the Random Forest model generated with the smaller meta-data feature set.

**Table 9 Top IP-Flows Identified with Random Forest and Full Feature Set**

Hostname	Flow Bytes	Probability
<b>Belgacom.be</b>	6,499,379	0.95
<b>Lofanga</b>	263,839	0.93
<b>Ip-streaming.net</b>	219,701	0.94
<b>Mivitec.net</b>	135,214	0.93
<b>Ucom.am</b>	126,228	0.89
<b>Worldstream.nl</b>	117,343	0.98

We expected to find in the results worldstream.net as that was the host for the pirate IPTV service we used with the IPTV set-top box. In addition to worldstream.nl, the model identified the hosts associated with other free pirate IPTV services – Belgacom.be, ucom.am, lofanga, and ip-streaming.net. We inspected the packet capture file and validated that IP flows associated with video and the hosts on the network that we used to view pirated video.

In addition, the model did NOT label the Netflix and ESPN.com video traffic as video piracy that we had running.

The model proved to be efficient at identifying pirate video flows that had the characteristics of the pirate IPTV service that we trained the model to look for. Later inspection of the Belgacom.be flows in the packet capture file, revealed that the IP address was associated with a residential ADSL modem. Further illustrating the performance of the Random Forest model with the full feature set.

As expected, the meta-data only feature set did not perform as well and falsely labeled a higher percentage of the traffic as pirated video. This was consistent with the lower accuracy, precision and higher false positive rate.

### **3.3. Case Study: Cable Operator**

To further test how well the machine learning model performed, we tested the model on flow data from a cable operator residential broadband network that provided a NetFlow feed with samples every 1000 packets. The flow data was formatted in NetFlow v9 [12] format and included only the IP flow meta-data fields.

**Table 10 Cable Operator Capture File Statistics**

<b>Capture File Size, Bytes</b>	100 GBytes
<b>Number of Flows</b>	82, 955,586
<b>Time Duration</b>	100 minutes

**Table 11 Cable Operator Traffic with Meta-Data Only Feature Set**

	Flow Pairs		Bytes		End-points	
<b>Labeled Piracy</b>	96,765	5.9%	21,406,672,027	0.4%	82,809	10.1%
<b>Labeled Benign</b>	152,087	94.1%	5,324,495,204,054	99.6%	735,344	89.9%
<b>Total</b>	1,617,602	100%	5,345,901,876,081		818,153	100%

The byte count is the total number of bytes measured by the sampled NetFlow and therefore is lower than the actual number of bytes seen.

To reduce the false positives, we post-processed the results by running them through a whitelist filter as the labeled data included traffic that was either non-routable traffic (i.e. 0.0.0.0) or was from well-known sources such as ISPs, web hosts, and multi-tenant CDNs such as Amazon's CloudFront, Akamai, and Google and updated the labels to benign. The post-processing also included labeling traffic from well-known ISPs as benign. The post-processed results are shown in Table 11.

As we did with the enterprise data results, we looked at the end-points labeled piracy and ranked the top hosts by volume and are shown in Table 12. The top hosts labeled pirates includes at least one multi-tenant host, CDN-77, that is known to host video piracy. The results also included a number of false positives such as blizzard.com that is online gaming site.

**Table 12 Top Labeled Pirate Hosts for Cable Operator Case Study**

<b>Hostname</b>	<b>Flow Bytes</b>	<b>Probability</b>
<b>Blizzard.com</b>	2,677,441,935	1.0
<b>Cdn77.com</b>	551,847,384	1.0
<b>Valvesoftware.com</b>	489,138,456	1.0
<b>Servers.com</b>	167,763,386	1.0

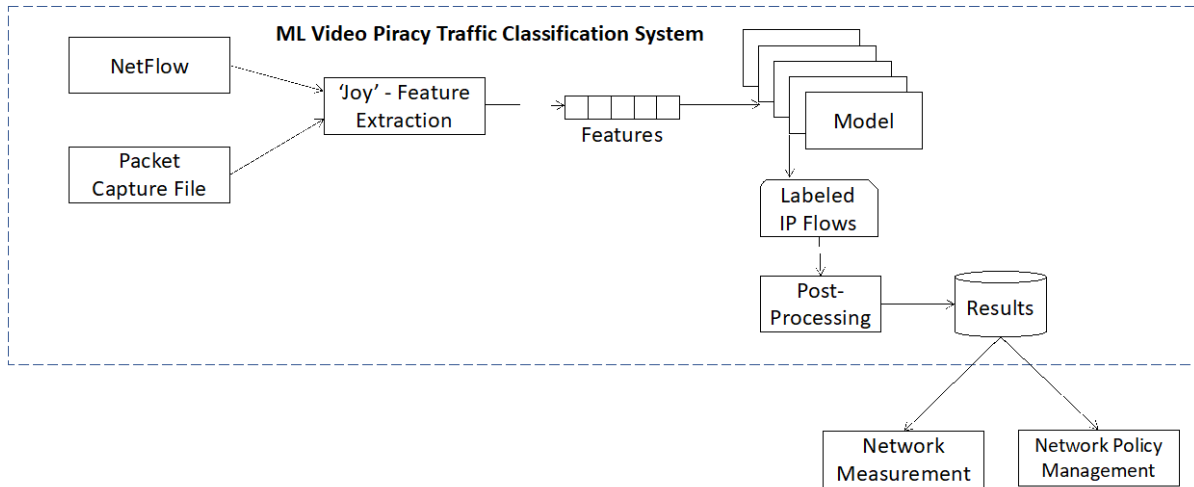
<b>Net.br</b>	252,875,736	1.0
<b>Bungie.net</b>	199,601,235	1.0

Overall, the results were consistent with our machine learning performance metrics when using the Random Forest model and the meta-data only feature set in Table 5. The results were also consistent with the findings from the enterprise case study when using only the meta-data feature set. The results could be improved with by including additional benign traffic samples in the training data and with further post-processing by expanding the whitelisted hosts.

## 4. Applications

The supervised machine learning model can be applied to operators networks to classify traffic in a number of ways in the measurement and mitigation of video piracy [13] [14]. (Note, patent has been applied for some of the methods and processes as applied to video piracy described in this paper.) One application is the classification of traffic as part of a traffic measurement system. A second application is to use the results of the machine learning system to mitigate piracy traffic with network policy enforcement systems such as PacketCable Multimedia [15] or as input to the Policy Charge Rule Function (PCRF) in the Evolved Packet Core of 4G and 5G network architectures.

Figure 7 shows a system schematic for an implantation of a system using flow data as input to a machine learning system for identifying video piracy. The system utilizes multiple machine learning models or classifiers, both for identifying multiple forms of video piracy and for reducing false positives by identifying forms of legitimate OTT video traffic. The output of the system may then be fed to an operators network measurement and/or policy enforcement system.



**Figure 7 Video Piracy Traffic Classification System**

## Conclusion

In this paper, we showed that some forms of pirated video delivered as OTT IPTV can be efficiently identified using a machine learning based traffic classification system. Further we showed that the Random Forest model out performs other machine learning models such as Logistic Regression and Decision Trees when used in this fashion.

We also showed that an efficient machine learning model can be built to classify traffic using IP flow data that can be ingested directly from a packet capture file or from an IPFIX/NetFlow feed. And finally, we showed that using flow data that is enhanced with byte distribution, packet size, and packet timing information such as the proprietary fields included in NetFlow for Cisco's Encrypted Traffic Analytics can be used to build a machine learning model that has a high accuracy and a low false positive rate.

# Abbreviations

ADSL	Asymmetric Digital Subscriber Line
AI	Artificial Intelligence
CDN	Content Delivery Network
DASH	Dynamic Adaptive Streaming over HTTP
DPI	Deep Packet Inspection
ESPN	Entertainment and Sports Programming Network
HBO	Home Box Office
HLS	HTTP Live Streaming
HTTP	Hyper Text Transfer Protocol
IPD	Initial Packet Data
IP	Internet Protocol
IPFIX	Internet Protocol Flow Information Export
IPTV	Internet Protocol Television
ISP	Internet Service Provider
MPEG	The Moving Picture Experts Group set standard for encoding and compressing video images
NBC	National Broadcasting Company
OTT	Over The Top
PCRF	Policy Charge Rules Function
SCTE	Society of Cable Telecommunications Engineers
SPLT	Sequence of Packet Lengths and Times
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VOD	Video On Demand



# Bibliography

- [1] NCTA, "Broadband by the Numbers," [Online]. Available: <https://www.ncta.com/broadband-by-the-numbers>. [Accessed 19 July 2019].
- [2] Parks Associates, "Parks Associates Forecasts \$12.5B in Lost Revenue in 2024 Due to Pay-TV and OTT Piracy and Account Sharing," [Online]. Available: <https://www.parkassociates.com/blog/article/pr-07162019>. [Accessed 19 July 2019].
- [3] D. Jones and K. Foo, "'Analyzing the Modern OTT Piracy Video Ecosystem - NCTA Technical Papers" Tech Paper Database," 2018. [Online]. Available: <http://www.nctatechnicalpapers.com/Paper/2018/2018-analyzing-the-modern-ott-piracy-video-ecosystem>. [Accessed 19 July 2019].
- [4] V. Mihajlovic, "'Global Internet Phenomena Spotlight: Video Piracy in North America." Sandvine," 13 December 2019. [Online]. Available: <https://www.sandvine.com/blog/global-internet-phenomena-spotlight-video-piracy-in-north-america>. [Accessed 19 July 2019].
- [5] B. Anderson and D. McGrew, "Identifying Encrypted Malware Traffic with Contextual Flow Data," in *Proceedings of the 2016 ACM Workshop on Artificial Intelligence - AISec 16, 2016*, 2016.
- [6] Cisco, "Cisco Encrypted Traffic Analytics," 2019. [Online]. Available: <https://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise-networks/enterprise-network-security/nb-09-encrytd-traf-anlytcs-wp-cte-en.pdf>. [Accessed 19 July 2019].
- [7] R. Pantos and W. May, "HTTP Live Streaming," August 2017. [Online]. Available: <https://tools.ietf.org/html/rfc8216>. [Accessed 19 July 2019].
- [8] International Organization for Standardization, "ISO/IEC 23009-1:2014 Preview Information technology -- Dynamic adaptive streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats," May 2014. [Online]. Available: <https://www.iso.org/standard/65274.html>. [Accessed 19 July 2019].
- [9] T. T. Nguyen and G. Armitage, "A Survey of Techniques for Internet Traffic Classification using Machine Learning," *IEEE Communications Surveys & Tutorials*, vol. 10, no. 4, Fourth Quarter 2008.
- [10] T. Hastie, R. Tibshirani and J. Friedman, "The elements of statistical learning: data mining, inference and prediction," Springer, 2009.
- [11] Scikit-learn, "Scikit-learn Machine Learning in Python," [Online]. Available: <https://scikit-learn.org>. [Accessed 19 July 2019].
- [12] Cisco, "NetFlow Version 9 Flow-Record Format," May 2011. [Online]. Available: [https://www.cisco.com/en/US/technologies/tk648/tk362/technologies\\_white\\_paper09186a00800a3db9.html](https://www.cisco.com/en/US/technologies/tk648/tk362/technologies_white_paper09186a00800a3db9.html). [Accessed 19 July 2019].

- [13] M. Tooley and W. Check, "Method and System for Detecting Pirated Video Network Traffic". United States of America Patent 62/740,569, 8 October 2018.
- [14] M. Tooley and W. Check, "Method and System for Detecting Pirated Video Network Traffic". United States of America Patent 16/381,571, 11 April 2019.
- [15] CableLabs, "PacketCable Specification - Multimedia Specification," 11 November 2015. [Online]. Available: <https://specification-search.cablelabs.com/packetcable-multimedia-specification>. [Accessed September 2019].

# Endnotes

- 
- 1 <https://github.com/cisco/joy>
  - 2 Flusonic, TVHeadend, Xtremecodes
  - 3 <https://github.com/cisco/joy>

# **The Importance Of Wi-Fi 6 Technology For Delivery Of gbps Internet Service**

A Technical Paper prepared for SCTE•ISBE by

**David John Urban**  
Distinguished Engineer  
Comcast  
One Comcast Center  
Philadelphia PA 19103  
610-476-2596  
david\_urban@cable.comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Traffic Demand.....	4
Background on Wi-Fi 6 Technology .....	8
FFT forms the OFDM Symbol .....	11
Modulation and Coding.....	13
Data, Pilot, Null Subcarriers.....	20
Spatial Streams .....	21
PHY Rate and Speed Test Relationship.....	23
Resource Units and OFDMA.....	24
Speed Tests and File Downloads.....	29
Conclusion .....	32
Abbreviations.....	32

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Data rate measured for video device UHD streaming .....	5
Figure 2 – Data rate measured of UHD video streaming device over full show .....	6
Figure 3 – Download 4.4 GB file in 30 seconds with Wi-Fi 6 2400 Mbps PHY.....	6
Figure 4 – Wi-Fi 6 2400 Mbps PHY speed test with 1 Gbps Internet service in room above.....	7
Figure 5 – Wi-Fi 6 2400 Mbps PHY download progression during speed test .....	8
Figure 6 – Illustration of extra path length of a reflected signal .....	9
Figure 7 – Primary channel 100 with channel width 160 MHz center frequency 5570 MHz .....	10
Figure 8 – FFT size is 2048 for a 160 MHz channel width 802.11ax signal .....	11
Figure 9 – BPSK modulation has two constellation points sending 1 bit for each data subcarrier in the 20 MHz OFDM symbol.....	16
Figure 10 – MCS 11 uses 1024-QAM modulation carrying 10 bits per subcarrier with measured MER and constellation diagram.....	17
Figure 11 – Measured Receiver sensitivity shows lower MCS levels work at lower received signal levels.....	18
Figure 12 – Diagram of channel matrix with 4 transmitters and 2 receivers.....	21
Figure 13 – Back to the drawing board, 802.11ax task group formed resource units, RU, to match original 802.11a signal .....	25
Figure 14 - 802.11a signal are still used in Wi-Fi 6, here for Block ACK.....	26
Figure 15 - 802.11a constellation for 24 Mbps Block ACK 16-QAM .....	27
Figure 16 - Time duration of 24 Mbps 802.11a signal, 1/8 of the total ax Block ACK.....	27

Figure 17 - Screen shot of notebook computer with 2400 Mbps Wi-Fi 6 2x2 160 MHz station.....	30
Figure 18 - Speed test with 2400 Mbps PHY notebook computer 2x2 160 MHz Wi-Fi 6.....	31
Figure 19 - Speed test with phone having 1200 Mbps Wi-Fi 6 2x2 80 MHz station .....	32

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Bits Modulation and SNR .....	13
Table 2 - 20 MHz 802.11ax Wi-Fi 6 Modulation and Coding .....	14
Table 3 - Measured levels, MCS, PHY rate, data rate moving throughout a residential home 2x2 160 MHz Wi-Fi 6 station.....	19
Table 4 - Measured results in each room of a residential home with 2x2 160 MHz Wi-Fi 6 station. ....	19
Table 5 - OFDMA tones and pilots for each resource unit, RU.....	24
Table 6 - OFDMA resource units RU, for 2x2 STA .....	28

# Introduction

Wi-Fi 6 based upon 802.11ax standard changes the OFDM carrier spacing and the multiple access method for the first time in 16 years.

The subcarrier spacing is narrowed by one fourth resulting in symbol times four times longer allowing the spectrum to be divided into many resource units to be shared by multiple users at the same time.

A home may have several video set top boxes and security cameras and each member of the family may actively be using a phone and notebook computer with messaging and video streaming applications.

Thus, twenty active stations consuming less than 5 Mbps each is an important use case.

The devices can be served efficiently with OFDMA.

Each device transmits and receives a narrow portion of the 160 MHz channel width.

This is more robust for each station while still efficient since many stations are served at the same time.

In addition to many devices in the home consuming small average data rates, the cable operator is expected to deliver the Gbps service tier that the customer pays for.

Serving many stations simultaneously with OFDMA frees up air time for devices that occasionally demand much higher peak data rates.

Wi-Fi 6 PHY rates for 2x2 80 MHz phones and tablets will be 1200 Mbps, for 2x2 160 MHz notebooks will be 2400 Mbps, for 4x4 160 MHz desktop computers with media bridge adapters 4800 Mbps.

These PHY rates are critical for the stable, reliable, consistent delivery of Gbps service in the presence of spectrum sharing both within the home and with neighbors.

The phone with 1200 Mbps maximum PHY rate can be expected to deliver speeds of 700-900 Mbps in the same room and adjacent rooms as well as directly above and below.

The notebook with 2400 Mbps maximum PHY rate can be expected to deliver equivalent speeds of a 1 Gbps Ethernet NIC with whole home coverage.


This paper describes the technical details of Wi-Fi 6 802.11ax high efficiency and why the technology is critical for cable operators to deliver consistent, stable, reliable Gbps service.

# Traffic Demand

Like Ethernet before it, Wi-Fi has been successful in part by using simple low cost technology with raw data rate an order of magnitude greater than traffic demand. 10Base-T for LAN was faster and simpler than 1.544 Mbps T1 lines. Wi-Fi works in shared spectrum governed with CSMA-CA and other distributed coordination methods such as RTS/CTS with very low cost stations so that Wi-Fi connectivity is ubiquitous in phones, tablets, video streaming boxes, televisions and notebook and desktop computers.

The speeds of Wi-Fi have been able to stay ahead of traffic demand over the years, from 11 Mbps, to 54 Mbps, to 300 Mbps, to 1300 Mbps, and with Wi-Fi 6 to 2400 Mbps. Having an understanding of the traffic demand from customers is essential for the service provider to size up the speeds needed in a cable modem wireless router.


As an illustrative example, the data rate was measured for a popular UHD video streaming app. The video was displayed on a UHD television set with the app running on an HDMI dongle with external power supply. The STA connected to the AP in the 2.4 GHz band. The STA had a maximum PHY rate of 130 Mbps, 2x2, 802.11n mode, 20 MHz channel width. The AP and STA were separated by one floor and located in opposite corners of the house. The data rate in Mbps was measured by the AP every second. The mean data rate measured was 3.928 Mbps with a peak data rate of 29.5 Mbps. The minimum and median value of the data rate was 0. The first quantile was 0 and the third quantile was 7.9 Mbps.



**Figure 1 – Data rate measured for video device UHD streaming**

Figure 1 shows a plot of the download rate measured in 1 second intervals during video streaming. With 12 bursts in 60 seconds, the traffic demand is satisfied with 20 Mbps one second data bursts followed by four seconds without any download demand. Figure 2 plots the measured downloaded data rate over a 37 minute time interval. While the traffic is constant over the full video viewing period, the traffic demand overall is less than 5 Mbps consisting of bursts lasting 1 second followed by no traffic demand for four seconds. Bursty traffic demand such as measured in this example is typical of video streaming and can serve as a traffic model for one of the most common and popular Internet consumer applications.





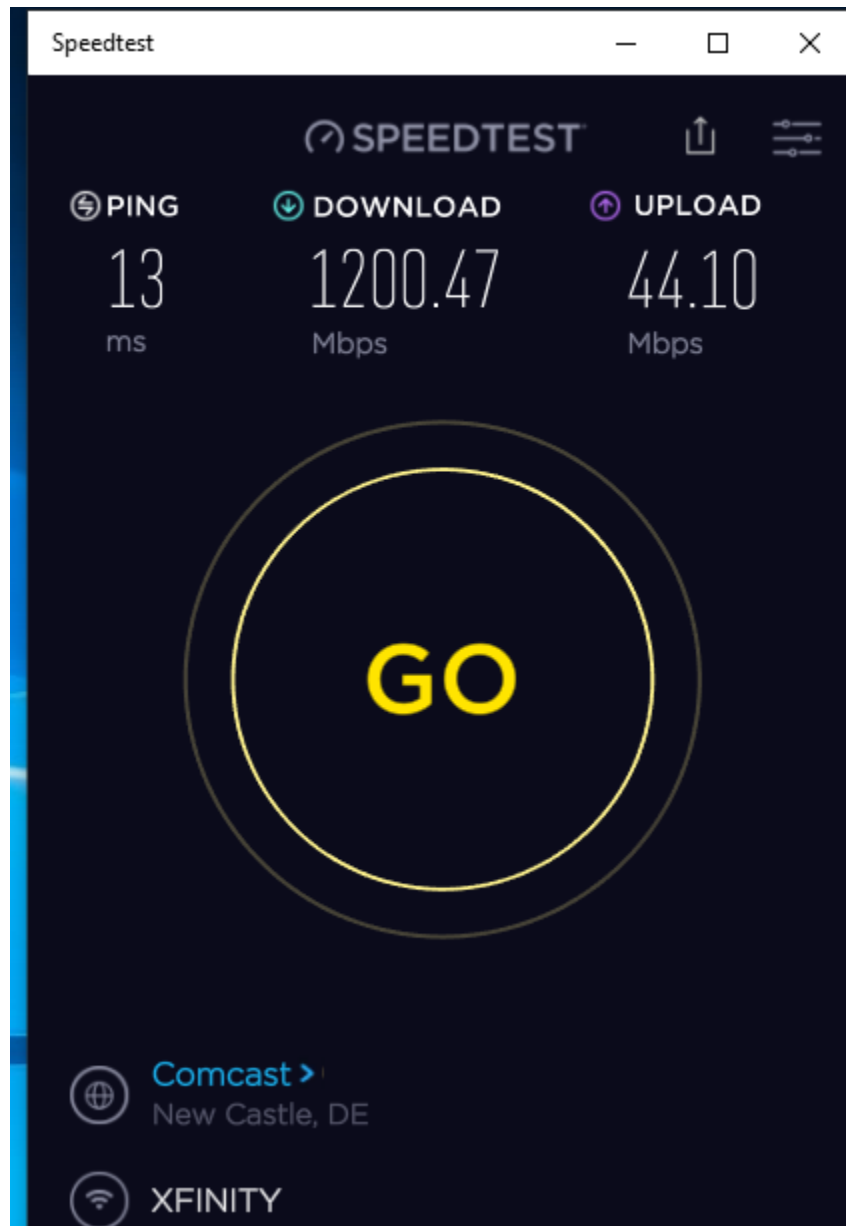
**Figure 2 – Data rate measured of UHD video streaming device over full show**

Downloading large files, including video downloads of many Gbytes, is an important application for high speed data service. Files can be very large and when a customer wants to download a video for offline viewing before running out the door appreciates the fastest speeds possible. As an illustrative example, a 4.4 GB file was downloaded from a popular website with a Wi-Fi 6 2400 Mbps STA over the Internet with a Gbps high speed data service. Even when downloading a large file, it still takes time to navigate the websites to set up the download. The peak download speed measured during the download of the 4.4 GB file was 1.2 Gbps and the average transfer rate was 235 Mbps. Once the download started, it took about 30 seconds before the download was finished as shown in Figure 3.




**Figure 3 – Download 4.4 GB file in 30 seconds with Wi-Fi 6 2400 Mbps PHY**

Traffic demand must be much lower than the delivered capacity by the service provider in order to ensure a satisfactory customer experience. It is important for the customer to know that the service promised is being delivered. A speed test is one method for customers to determine how their Internet service is working. Many factors impact speed test results, interpreting speed test results is as much an art as a science. Still, when understood correctly speed tests available to customers are a useful tool for the service provider and the customer. Figure 4 show a speed test measured with a Wi-Fi 6 2400 Mbps PHY station with a CMTS max-tr-rate of 1250 Mbps with the STA one floor above the AP.



**Figure 4 – Wi-Fi 6 2400 Mbps PHY speed test with 1 Gbps Internet service in room above**

Figure 5 shows the progression of the download rate during the speed test. Not all speed test download progressions look like this, which accounts for variations in speed test results from run to run. Getting the best speed test results from DOCSIS to Wi-Fi requires a steady increase in download rate from 0 to 1250 Mbps within the first 10 seconds of the download portion of the speed test followed by 5 seconds of steady 1250 Mbps download rate. When speed tests go wrong and report less than the provisioned speed of the cable modem and the throughput capability of the Wi-Fi, it is typically due to the download rate taking too long to ramp up to 1250 Mbps or hitting the 1250 Mbps but dropping to lower rates, failing to maintain the peak rate for the full 5 seconds.




**Figure 5 – Wi-Fi 6 2400 Mbps PHY download progression during speed test**

## Background on Wi-Fi 6 Technology

To help us understand the new 802.11ax he mode Wi-Fi 6 protocol, let's put ourselves in the shoes of the developers of the standard. How can we improve upon the very successful 802.11 Wi-Fi system? Let's start with a very simple channel model. If we can estimate the delay spread, then we can determine the necessary guard interval between OFDM symbols needed to prevent inter-symbol interference. Once the guard interval has been established a more efficient FFT size can be selected to form the OFDM symbol.

The delay of a ray of light or other radio frequency wave that reflects off an object 400 feet away is 800 ns since light travels one foot each ns and the reflected ray must travel an additional 800 feet. This is illustrated in Figure 6. The average home size in the United States is 2700 square feet. A rectangle with sides of 37 feet and 73 feet has an area of 2700 square feet. The distance between the AP and the STA both inside the average home will be less than 70 feet. Reflections from signals with delay greater than 800 ns are due to objects outside the home or many long reflections inside the home. This indicates that the delay spread of the WLAN channel for indoor residential homes is well below 800 ns. This has been born out in practice since the introduction of 802.11a in 1999, the last century. The normal guard interval is 800 ns for 802.11a, 802.11g, 802.11n, 802.11ac. All these standards have proved well suited for wireless connectivity for indoor residential channels as well as office building and short-range outdoor applications.



**Figure 6 – Illustration of extra path length of a reflected signal**

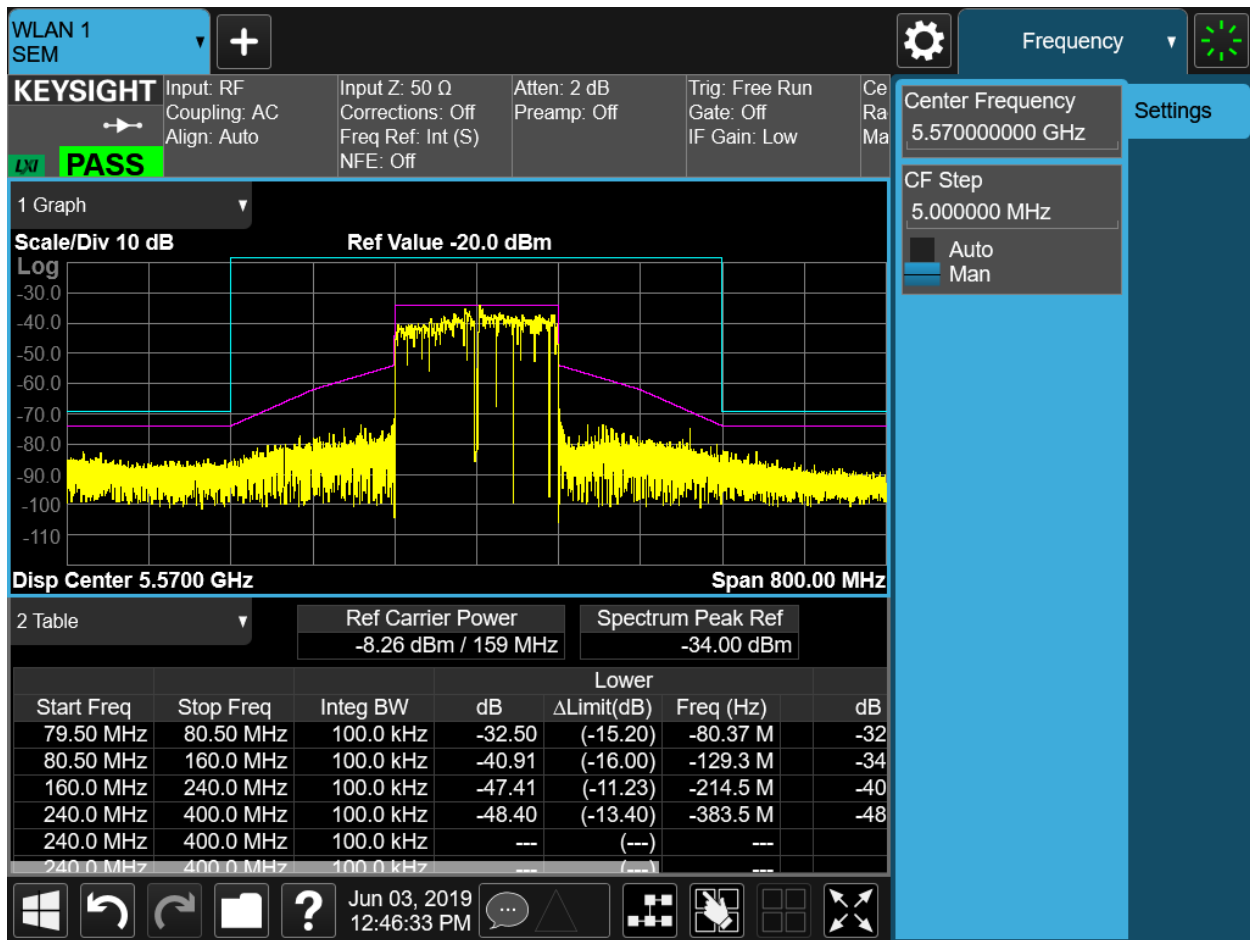
Channel 100 with 160 MHz channel width in the 5 GHz band has a center frequency of 5570 MHz as shown in Figure 7. An OFDM symbol with a 160 MHz channel width has a sampling frequency of 160 MHz and a sampling period of 6.25 ns. Equation 1 shows the formula and the calculation where  $T_s$  is the sampling period and  $R_s$  is the sampling rate; the sampling rate is expressed as 0.160 GHz so that the sampling period is calculated in ns. 800 ns guard interval between OFDM symbols requires 128 FFT time samples. The calculation of 128 FFT time samples for the guard interval is shown in equation 2.  $N_{CP}$  is the number of FFT time samples in the guard interval,  $T_g$  is the guard interval of 800 ns, and  $T_s$  is the FFT sampling rate of 6.25 ns.

*Equation 1*

$$T_s = \frac{1}{R_s} = \frac{1}{0.160} = 6.25$$

*Equation 2*

$$N_{CP} = \frac{T_g}{T_s} = \frac{800}{6.25} = 128$$



**Figure 7 – Primary channel 100 with channel width 160 MHz center frequency 5570 MHz**

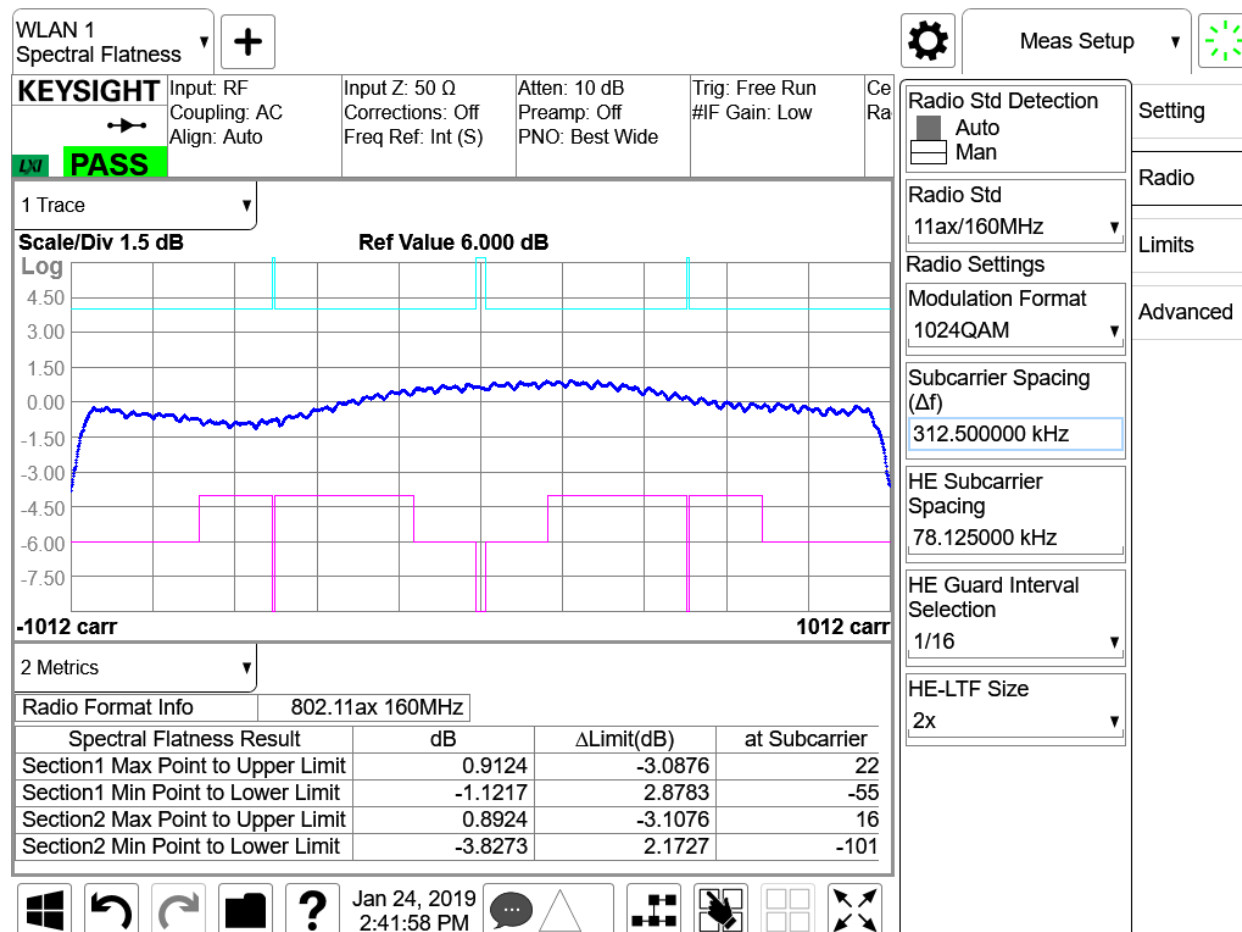
Wi-Fi 6 802.11ax he mode is 4 times more efficient than older modes. The guard interval is 1/16 of the symbol period rather than 1/4 of the symbol period. The FFT time samples used to create the guard interval between symbols are called the cyclic prefix or CP for short. Knowing the guard interval will help us determine the FFT size. To date Wi-Fi signals have had a guard interval that is 1/4 the symbol period. With a CP of 1/16 the FFT size and 128 time sample CP the FFT size is 2048 as calculated in equation 3.

Equation 3

$$N_{FFT} = \frac{N_{CP}}{CP} = \frac{128}{\frac{1}{16}} = 2048$$

# FFT forms the OFDM Symbol

The spectral flatness of a 160 MHz channel width Wi-Fi 6 is shown in Figure 8. Notice that the radio standard is 802.11ax 160 MHz, the HE subcarrier spacing 78.125 kHz and the HE guard interval 1/16. The FFT index of the lowest frequency data subcarrier is -1012 and the FFT index of the highest frequency subcarrier is 1012. For a 2048 point FFT the subcarrier indices range from -1024 to 1023. The 12 lowest frequency subcarriers are set to zero and the 11 highest frequency subcarriers are set to zero to establish a frequency guard band between lower adjacent and upper adjacent channels.



**Figure 8 – FFT size is 2048 for a 160 MHz channel width 802.11ax signal**

The shortest in time guard interval is 800 ns for the mode of 802.11. Longer in time guard intervals of the mode are 1600 ns and 3200 ns. The WLAN standard 802.11ax is designed for high efficiency, so the mode is called HE. The Wi-Fi alliance Wi-Fi 6 is based upon the HE mode IEEE 802.11ax standard. The guard interval provides a space between consecutive symbols. The guard interval is implemented as a cyclic prefix, abbreviated CP.

The term cyclic prefix is used since the first time samples of the IFFT of the OFDM symbol are repeated at the end of the symbol before the next symbol is sent. Since light travels one foot in one ns an 800 ns guard interval between symbols allows for reflections from objects 400 feet away to have a delay that is not long enough to result in inter-symbol interference.

800 ns guard interval is enough to prevent inter-symbol interference for most indoor WLAN channel conditions and even most short-range outdoor applications. Since the distance between AP and STA is less than 30 meters for indoor applications and less than 100 meters for outdoor applications reflected rays travelling greater than 800 meters are highly likely to encounter obstructions such as walls and trees. Long delayed multi-path rays suffer high attenuation after reflection, diffraction, passing through walls and trees. Due to attenuation from scattering objects rays with delay greater than 800 ns have negligible level relative to the rays taking a more direct path from AP to STA having delays less than 80 ns. A rectangular space with area of 5000 square feet has sides of 70 feet. With an AP placed in the middle of the rectangle, an STA in the middle of one side would have a direct path of 35 feet for a delay of 35 ns.

The guard interval of 1600 and 3200 ns, allowing for reflections 800 feet and 1600 feet away, benefit in longer range applications where the AP is higher in elevation and the STA is further away from the AP. The guard interval is selected based upon the delay spread of the channel. The Wi-Fi 6 802.11ax he mode guard interval of 800, 1600, and 3200 ns makes the protocol suitable for channels with delay spread less than the chosen guard interval.

Indoor residential, indoor office buildings, and short range outdoor channels in the 2.4 and 5 GHz band meet the delay spread limits of the he mode. These are the channels of interest for a service provider delivering Gbps Internet access to residential and business customers over a hybrid fiber coaxial cable HFC network. The same is true for service providers delivering Gbps service to residential and business customers over other access architectures or a mix of access architectures.

An OFDM symbol with channel width of 160 MHz has an FFT sampling period of 6.25 ns. With an FFT size of 2048 the FFT duration is 12.8  $\mu$ s as calculated in equation 4. The FFT is sometimes called the useful symbol time and that is why the symbol used for it is often  $T_u$ . Some folks object to the term useful symbol period since it implies that the guard interval is not useful, perhaps hurting the guard intervals feelings.

*Equation 4*

$$T_u = T_s \cdot N_{FFT} = 6.25 \cdot 2048 \cdot 10^{-3} = 12.8 \mu s$$

The symbol period is the sum of the FFT duration  $T_u$  and the guard interval  $T_g$  as calculated in equation 5. The FFT duration is 12.8  $\mu$ s and the guard interval is 800 ns = 0.8  $\mu$ s so that the symbol period is 13.6  $\mu$ s. The symbol period for older 802.11 OFDM symbols is 4  $\mu$ s for normal guard interval. The FFT duration for older 802.11 OFDM modes is 3.2  $\mu$ s. The FFT duration of the he mode is 4 times longer than the FFT duration of older OFDM modes, 3.2  $\mu$ s for older modes and 12.8  $\mu$ s for the he mode. Remember the he mode symbol period for normal guard interval of 13.6  $\mu$ s, this is what we will use in the denominator to calculate the data rate of the he mode symbol for various PHY settings.

*Equation 5*

$$T_{symbol} = T_u + T_g = 12.8 + 0.8 = 13.6 \mu s$$

The subcarrier spacing can be calculated from the inverse of the FFT duration. The subcarrier spacing of an he mode signal is 78.125 kHz as shown in equation 6. 1000 is used in the numerator dividing by 12.8  $\mu$ s so that the resulting frequency spacing has units of kHz. The subcarrier spacing of older 802.11 OFDM modes is four times greater than the he mode; 312.5 kHz for older modes compared to 78.125 kHz for the he mode.

Equation 6

$$\Delta f = \frac{1}{T_u} = \frac{1000}{12.8} = 78.125 \text{ kHz}$$

Now we know the characteristics of the FFT that creates the he mode OFDM symbol. We know we need a 160 MHz channel width in the 5 GHz band. We are confident that the 800 ns guard interval will prevent inter-symbol interference in our customers homes. We want to know if we can deliver Gbps Internet service to our customers. We will need to calculate the data rate of the he mode symbol. We know the denominator, the symbol period. We need to know the numerator. For this we need to understand the number of spatial streams, the number of data subcarriers, and the modulation and coding of the signal.

## Modulation and Coding

Each data subcarrier of the FFT that forms an OFDM symbol is modulated. The modulation is adaptive, adjusting to the channel conditions to get the highest data rate possible for the signal to noise ratio at the time of transmission. The constellations used for modulating the data subcarriers in Wi-Fi 6 he mode are BPSK, QPSK, 16-QAM, 64-QAM, 256-QAM and 1024-QAM. Table 1 lists the bit loading of the subcarriers, the name of the modulation corresponding to the bit loading and the SNR needed for  $10^{-6}$  bit error rate without error correction coding.

**Table 1 - Bits Modulation and SNR**

bits	modulation	SNR in dB
1	BPSK	6
2	QPSK	12
4	16-QAM	18
6	64-QAM	24
8	256-QAM	30
10	1024-QAM	36

The measured constellation of a demodulated BPSK signal is shown in Figure 4. There are two constellation points, one point can represent a 0 and the other a 1. That is why BPSK modulation transmits 1 bit each symbol for the data subcarriers of the OFDM symbol.

The modulation options for he mode are the same as vht mode with the exception of the addition of 1024-QAM. There have been proprietary implementations of 1024-QAM but Wi-Fi 6 introduces 1024-QAM modulation as part of the standard allowing better device interoperability. The constellation diagram is shown in Figure 9 for an 802.11ax 160 MHz channel width signal. Since  $2^{10}=1024$  with 1024 constellation points each point can represent 10 bits. 1024-QAM modulation conveys 10 bits for each data subcarrier in an OFDM symbol.

Adding LDPC coding to the modulation allows for more bits per subcarriers for a given signal to noise ratio. The bits that load subcarriers for modulation come from LDPC codewords. Low density parity check coding, take large blocks of information bits and form large codewords. When decoded the LDPC



codewords use message passing algorithms to find the codeword based on the log likelihood ratio of the received signal. There are three different LDPC codeword block lengths; 648, 1296, and 1944 bits. The more bits in a LDPC codeword the better the coding gain, at the expense of latency and complexity. Smaller traffic demand may not be able to fill a 1944 bit codeword so a smaller codeword is used. There are four different coderates; 1/2, 2/3, 3/4, 5/6. For a coderate of 1/2 and a code word length of 1944 bits, the LDPC information block length is 972 bits. 972 bits are fed to the LDPC coder to form a 1944 bit codeword. The 1944 bits of the codeword bit load the data subcarriers of the OFDM symbol.

<b>coderate</b>	<b>Codeword</b>	<b>Information</b>
<b>ratio</b>	<b>bits</b>	<b>bits</b>
1/2	1944	972
2/3	1944	1296
3/4	1944	1458
5/6	1944	1620
1/2	1296	648
2/3	1296	864
3/4	1296	972
5/6	1296	1080
1/2	648	324
2/3	648	432
3/4	648	486
5/6	648	540

The receiver demodulates the signal with an FFT to determine the magnitude and phase of each data subcarrier. With the magnitude and phase measurement of enough data subcarriers for a codeword, the receiver calculates the log likelihood ratio of the 1944 bits of a codeword. The log likelihood ratio for each bit is calculated by measuring the distance from the received subcarrier vector and each of the possible transmitted constellation points for the modulated signal and making an estimate of the probability density function.

For a service provider, the important point to understand about LDPC codes are that they are 1) very powerful, 2) very long so complexity is high and latency and traffic demand need to be considered, and 3) the LDPC code rate must be known to determine the speed that customers will experience using Wi-Fi.

Table 3 shows the modulation, code rate, receiver level for each of the 12 modulation and coding schemes of 802.11ax Wi-Fi 6 he mode.

**Table 2 - 20 MHz 802.11ax Wi-Fi 6 Modulation and Coding**

<b>MCS index</b>	<b>modulation</b>	<b>bits</b>	<b>coderate</b>	<b>bps/Hz</b>	<b>802.11 dBm</b>	<b>Sensitivity dBm</b>	<b>SNR dB</b>
0	BPSK	1	1/2	0.50	-82	-95	3
1	QPSK	2	1/2	1.00	-79	-92	6
2	QPSK	2	3/4	1.50	-77	-90	8
3	16-QAM	4	1/2	2.00	-74	-87	11
4	16-QAM	4	3/4	3.00	-70	-83	15
5	64-QAM	6	2/3	4.00	-64	-77	21


<b>MCS index</b>	<b>modulation</b>	<b>bits</b>	<b>coderate</b>	<b>bps/Hz</b>	<b>802.11 dBm</b>	<b>Sensitivity dBm</b>	<b>SNR dB</b>
6	64-QAM	6	3/4	4.50	-65	-78	20
7	64-QAM	6	5/6	5.00	-64	-77	21
8	256-QAM	8	3/4	6.00	-59	-72	26
9	246-QAM	8	5/6	6.67	-57	-70	28
10	1024-QAM	10	3/4	7.50	-54	-67	31
11	1024-QAM	10	5/6	8.33	-52	-65	33

Modulations with higher bits per symbol also require higher signal to noise ratio and thus higher input level. When modulation level is increased so that two addition bits are loaded onto each subcarrier, the signal to noise ratio penalty is 6 dB. All things being equal 1024-QAM which has 10 bits per symbol needs 6 dB higher signal to noise ratio than 256-QAM which has 8 bits per symbol.

Boltzmann constant is  $1.38 \times 10^{-23}$  Joules per degree Kelvin. The thermal noise power spectral density can be calculated by multiplying Boltzmann constant by the temperature. For temperatures encountered on earth, the thermal noise power spectral density when converted to dBm is about -174 dBm/Hz. For a 20 MHz channel width and 3 dB noise figure receiver the noise floor of the receiver is -98 dBm. With a receive signal level of -95 dBm and a noise floor of -98 dBm, the SNR is 3 dB. A 3 dB SNR is plenty for a BPSK and 1/2 LDPC code rate signal. With a 3 dB noise figure receiver and a demodulator that can work at 3 dB SNR for a BPSK 1/2 rate LDPC code the receiver sensitivity is -95 dBm. With four receivers and a maximum ratio combining gain of 6 dB, the sensitivity is -101 dBm per chain for MCS0 20 MHz channel width.

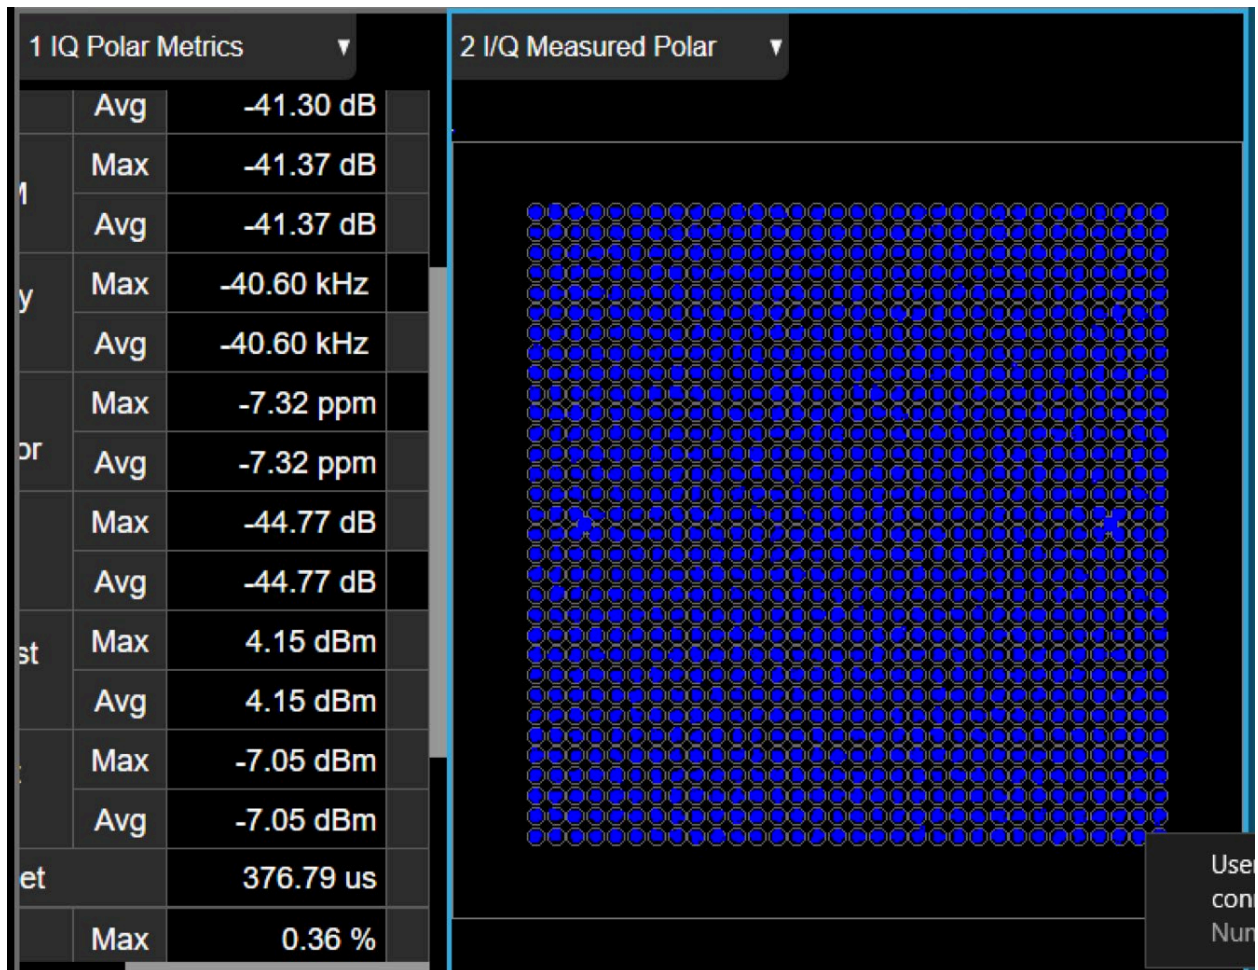
The IEEE dBm labeled column in Table 3 is the IEEE 802.11 receiver sensitivity threshold levels. The sensitivity dBm labeled column in Table 3 shows the levels for a 3 dB noise figure receiver that can demodulate MCS0 20 MHz at 3 dB SNR.

Figure 9 shows the constellation diagram of BPSK and the spectrum of the 20 MHz OFDM symbol. As seen in Table 3 MSC 0 uses BPSK modulation, binary phase shift keying. BPSK sends 1 bit of information, 0 has a phase of 180 degrees and 1 has a phase of 0 degrees as seen in the constellation measurement of Figure 9. The 20 MHz signal is divided into OFDM subcarriers. Each of the data subcarriers are BPSK modulated for MSC 0 carrying 1 bit for each data subcarrier. There are also left and right guard subcarriers, DC null subcarriers, and pilot subcarriers so not all of the FFT subcarriers carry data bits, but most do.




**Figure 9 – BPSK modulation has two constellation points sending 1 bit for each data subcarrier in the 20 MHz OFDM symbol.**

Figure 10 shows the measured constellation diagram of a 802.11ax 160 MHz MCS11 1024-QAM 5/6 rate LDPC signal. The constellation diagram has 1024 points. Each point of the constellation diagram is mapped to 10 bits. Thus 1024-QAM modulation carries 10 bits in each data subcarrier forming an OFDM symbol. Not all of these bits are information bits, the bits that modulate data subcarriers come from LDPC codewords, for MCS 11 the LDPC codeword rate is 5/6 so only 5 out of every 6 bits carry information. At the demodulator the final SNR must be 32 to 33 dB in order to demodulate MCS 11. With a 3 dB noise figure receiver, the receive level must be -56 dBm for a 160 MHz channel width MCS 11 signal. Since the noise floor of a 3 dB noise figure receiver is -98 dBm 20 MHz channel width, the noise floor for 160 MHz is 9 dB higher, -89 dBm. 33 dB above -89 dBm is -56 dBm. The 33 dB SNR needed for 1024-QAM 5/6 LDPC modulation is the sum total of all noise and interference into the receiver demodulator. In a well designed system, the noise will be dominated by the receivers noise figure and the fundamental thermal noise given the channel width. The transmitter EVM, error vector magnitude, is a contributing factor in the cumulative noise into the receiver. The transmitter EVM target should be about 10 dB better than the final required SNR where practical, this is a reasonable compromise between transmitter complexity and system impact.



**Figure 10 – MCS 11 uses 1024-QAM modulation carrying 10 bits per subcarrier with measured MER and constellation diagram.**

The measured packet error rate versus estimated RSSI measured for various MCS indices for a 20 MHz 802.11ac signal are shown in Figure 11. The measured levels in Figure 11 correspond closely to the calculated levels in Table 2 which were based upon first principles of noise figure and modulation.



**Figure 11 – Measured Receiver sensitivity shows lower MCS levels work at lower received signal levels.**

Table 4 shows a summary of measurements made in a residential home. The STA is a 2x2 160 MHz Wi-Fi 6 device with a maximum PHY rate of 2400 Mbps. The STA is a PCIe M.2 card in a Windows 10 computer. The computer was moved throughout the house to get MCS11 all the way down to MCS1. The RSSI received into the four antennas of the 4x4 Wi-Fi 6 AP were measured and reported in Table 4. At a receive level just above -50 dBm, the STA ran at full 2400 Mbps with a download data rate of 1.44 Gbps. The download data rate was measured between computer connected to a 2.5 Gbps Ethernet port on the cable modem wireless router and the notebook computer with the 2400 Mbps PHY Wi-Fi 6 card. At further distances and more obstructions such as walls and floors, the RSSI dropped as did the MCS and the channel width along with the data rate.

The MRC, maximal ratio combining, uses the formula shown in the equation below.

Equation 7

$$P_{mrc} = 10 \cdot \log_{10} \left( 10^{\frac{P_{rx1}}{10}} + 10^{\frac{P_{rx2}}{10}} + 10^{\frac{P_{rx3}}{10}} + 10^{\frac{P_{rx4}}{10}} \right)$$

In Table 4, the four receiver levels measured by the AP are reported. For a single information stream it is possible to vectorially add the signals from all four antennas. This is calculated and reported in the column labeled MRC in Table 4 using Equation 7. This is sometimes referred to as receiver beamforming since the receive and take the four received signals and adjust the phase of each so that the vectors add

constructively. This results in an equivalent receive level that is higher than that of each chain of the receiver individually.

**Table 3 - Measured levels, MCS, PHY rate, data rate moving throughout a residential home 2x2 160 MHz Wi-Fi 6 station**

RX1	RX2	RX3	RX4	MRC	Mode	MCS	Nss	bw	GI	PHY	data
dBm	dBm	dBm	dBm	dBm		index	streams	MHz	μs	Mbps	Gbps
-55	-49	-50	-47	-43.4	he	11	2	160	0.8	2401.9	1.44
-55	-56	-54	-56	-49.1	he	10	2	160	0.8	2161.76	1.34
-59	-61	-59	-59	-53.4	he	9	2	160	0.8	1921.54	1.23
-65	-58	-56	-62	-53.0	he	8	2	160	0.8	1729.41	1.18
-68	-62	-60	-63	-56.4	he	7	2	160	0.8	1441.17	0.945
-68	-64	-61	-68	-58.2	he	6	2	160	0.8	1297.05	0.978
-67	-72	-68	-66	-61.7	he	5	2	160	0.8	1152.94	0.645
-70	-70	-66	-69	-62.4	he	4	2	160	0.8	864.7	0.687
-72	-71	-71	-75	-66.0	he	5	2	80	0.8	576.47	0.39
-78	-72	-76	-76	-68.9	he	4	2	80	0.8	432.35	0.342
-80	-79	-74	-79	-71.3	he	3	2	80	0.8	288.23	0.22
-83	-79	-81	-82	-75.0	he	1	2	80	0.8	144.11	0.126

**Table 4 - Measured results in each room of a residential home with 2x2 160 MHz Wi-Fi 6 station.**

Location	RX1	RX2	RX3	RX4	MRC	mode	mcs	Nss	bw	GI	PHY	data	iperf
Units	dBm	dBm	dBm	dBm	dBm		index	streams	MHz	μs	Mbps	Mbps	Gbps
family room	-60	-60	-55	-55	-50.8	he	11	2	160	0.8	2401.5	1522.4	1.5
office	-51	-46	-51	-51	-43.1	he	11	2	160	0.8	2401.5	1544	1.47
kitchen	-58	-60	-58	-58	-52.4	he	11	2	160	0.8	2401.5	1501.5	1.42
library	-63	-53	-53	-58	-49.2	he	11	2	160	0.8	2401.5	1542.3	1.45
Annie's room	-50	-57	-53	-55	-47.0	he	11	2	160	0.8	2401.5	1515.5	1.5
Sam's room	-68	-63	-67	-73	-60.4	he	8	2	160	0.8	1733.7	1213.8	1.23
upstairs bath	-60	-56	-55	-58	-50.8	he	11	2	160	0.8	2401.5	1502.3	1.5
Katie's room	-70	-69	-72	-73	-64.7	he	6	2	160	0.8	1295.3	911.5	0.881
master bath	-65	-74	-69	-70	-62.4	he	7	2	160	0.8	1441	1060.1	1.02
master bed	-81	-80	-75	-80	-72.3	he	3	2	160	0.8	575.9	477.3	0.455
master bed	-80	-78	-74	-78	-70.9	he	4	2	160	0.8	858.5	514.7	0.531
master bed	-77	-76	-74	-77	-69.8	he	4	2	160	0.8	862.3	671	0.665
master bed	-84	-80	-79	-83	-75.0	he	4	1	160	0.8	394.7	264.7	0.23
dining room	-66	-62	-60	-67	-56.8	he	9	2	160	0.8	1920.3	1328.1	1.3
washer	-69	-69	-63	-69	-60.6	he	8	2	160	0.8	1727.9	1204.6	1.19
powder room	-71	-70	-64	-70	-61.7	he	7	2	160	0.8	1441	1037.3	1.03
garage	-75	-72	-67	-72	-64.5	he	7	2	160	0.8	1441	1060	1.02
basement	-69	-65	-64	-69	-60.1	he	9	2	160	0.8	1921.5	1301.5	1.2
front porch	-69	-66	-65	-64	-59.6	he	7	2	160	0.8	1441	1069.9	1.03
back porch	-69	-65	-69	-71	-61.9	he	7	2	160	0.8	1441	1069	1.03
mailbox	-74	-71	-67	-72	-64.2	he	6	2	160	0.8	1295	880.3	0.867
back yard	-75	-76	-77	-75	-69.7	he	4	2	160	0.8	864.1	687.4	0.652

Table 5 shows download data rates, RSSI levels, MCS, channel width and other parameters measured in every room of the house as well as a few indoor to outdoor locations. These results match reasonably well with the theoretical and test set measured results. The download throughput was measured between a computer connected to the 2.5 Gbps Ethernet port of the cable modem wireless gateway and the notebook computer with 2x2 160 MHz 2400 Mbps PHY station. Of note is that 16 rooms of the house had measured data rate above 1 Gbps out of a total of 22 locations. The overall coverage was good both inside and outside the home.

# Data, Pilot, Null Subcarriers

Each data subcarrier is modulated with bits from LDPC codewords. We now know the bits per data subcarrier for each MCS index. But how many data subcarriers are there? The FFT size is 2048 for a 160 MHz he mode symbol.

The FFT of a 160 MHz OFDM symbol converts the frequency domain subcarriers into a time domain waveform for input to a digital to analog converter. The FFT size for a 160 MHz channel width in he mode is 2048, thus there are 2048 subcarriers. The types of subcarriers are null, data, pilots.

Null subcarriers are placed at each end of the channel width spectrum and at the carrier frequency. The null subcarriers at the ends of the spectrum allow for a guard band in the frequency domain between channels. The null subcarriers at the carrier center frequency allow the receiver to work around DC offset since for a direct conversion receiver the carrier frequency will be translated to a baseband DC.

The 160 MHz subcarrier assignment consists of two adjacent 80 MHz channel width OFDM symbols. It is possible to demodulate the left side of the 160 MHz OFDM symbol and the right side of the OFDM symbol with an 80 MHz channel width demodulator. You can try this for yourself if you have a signal analyzer that can demodulate he mode signals. For channel 100 and 160 MHz channel width the center frequency is 5570 MHz. The center frequency is set to 5570 MHz and the demodulator channel width set to 160 MHz will demodulate the entire 160 MHz OFDM symbol. The center frequency can be set to 5530 and the channel width to 80 MHz to demodulate just the left half of the 160 MHz signal. Likewise, the center frequency can be set to 5610 and the channel width set to 80 MHz to demodulate just the upper half of the 160 MHz OFDM symbol. In fact it is even possible to create a 160 MHz channel width signals with two antennas, each antenna transmitting only half of the 160 MHz signal.

The 160 MHz 802.11ax signal is just two side by side 80 MHz channel width signals. Let's therefore take a closer look at the 80 MHz channel width signal. The FFT size for an 802.11ax he mode signal is 1024, half that of a 160 MHz signal. The lower guard band consists of 12 subcarriers while the upper guard band consists of 11 subcarriers. We do not after all wish to have the subcarriers of two different channels too close to each other in frequency. 5 DC subcarriers are set to zero in the middle of the spectrum. The 80 MHz channel width he signal has 16 pilots. That leaves 980 data subcarriers in an 80 MHz channel width he signal. Equation 8 shows the allocation of subcarriers in an 80 MHz channel width he signal.

*Equation 8*

$$1024 \text{ FFT size} = (12 \text{ lower guard band}) + (5 \text{ DC}) + (11 \text{ upper guard band}) + (16 \text{ pilots}) + (980 \text{ data})$$

*Equation 9*

$$R_b = \frac{N_d \cdot b \cdot r}{T_{symbol}} = \frac{980 \cdot 10 \cdot \frac{5}{6}}{13.6} = 600 \text{ Mbps}$$

With the number of data subcarriers, the modulation bits per symbol, the LDPC code rate, and the symbol period, the data rate of the symbol can be calculated. As shown in Equation 9

the data rate is 600 Mbps for MCS11 1204-QAM, 5/6 rate LDPC code for an 80 MHz he mode symbol with 800 ns guard interval. 600 Mbps for 80 MHz channel width and single spatial stream for Wi-Fi 6 is worth committing to memory. If you can remember 600 Mbps per spatial stream for 80 MHz channel then the relationship between channel width and number of chains can be easily understood and derived for


both AP's and STA's. An AP with 8x8 and 80 MHz has a 4800 Mbps PHY. Why? One spatial stream has 600 Mbps so 8 spatial streams have 4800. An AP with 4x4 and 160 MHz also has a PHY rate of 4800 Mbps. With twice the channel width, 160 MHz compared to 80 MHz, the 600 Mbps is multiplied by 2 giving 1200 Mbps, then 1200 Mbps is multiplied by 4 for 4 spatial streams to reveal the 4800 Mbps PHY rate. A 2x2 phone with 80 MHz channel width has a 1200 Mbps PHY, 2 spatial streams times 600 Mbps per spatial stream. And a 2x2 160 MHz computer has a 2400 Mbps PHY, 2 times 600 to account for 160 MHz channel width, and 2 times 1200 to account for two spatial streams, yielding 2400 Mbps PHY rate. Don't worry if you are confused about the use of spatial streams, that is covered in the next section.

A 160 MHz channel width has 2 times 980 equals 1960 data subcarriers. That's a lot of subcarriers for one station. The 12 lowest frequency subcarriers and the 11 highest frequency subcarriers are set to zero to provide a guard band between adjacent channels. Subcarriers at the center of the channel around the carrier local oscillator are set to zero to avoid a DC offset in the direct converted baseband signal. The 160 MHz he mode signal has 7 null subcarriers in the band center.

Pilot subcarriers are needed in order to estimate frequency offset and channel frequency response. Since the symbol rate is known to be 13.6  $\mu$ s, the phase of a pilot can be measured over two consecutive symbols and the frequency offset can be calculated by dividing the phase difference by the symbol period. The channel frequency response can be estimated by comparing two pilots at different subcarriers frequency.

## Spatial Streams

MIMO allows for multiple spatial streams to be sent in one OFDM symbol when multiple antennas and chains are used in both the transmitter and receiver. With an AP having four transmit chains and an STA having two receive chains the channel matrix has two rows and four columns. The elements of the channel matrix are the impulse response between a transmit antenna and a receive antenna.



**Figure 12 – Diagram of channel matrix with 4 transmitters and 2 receivers**



Figure 12 Shows a diagram of the paths between four transmitters and two receivers. The channel matrix transmit antenna and each receive antennas, 12 impulse response elements in all. The relationship between the input signals and the output signals illustrated in Figure 12 and mathematically described in Equation 10 is a matrix equation with measured output a 2 element vector Y and four element input vector X. If an inverse to the 2 by 4 channel impulse response matrix H can be calculated, then the input signals can be calculated by knowing the received signals and the channel matrix. When multi-path due to obstructions in the home allow the channel matrix to be inverted, then two spatial streams can be sent in one OFDM symbol. It turns out fortunately, that with 4 transmit antennas and 2 receive antennas it is quite easy to get a 2400 Mbps PHY rate at reasonable distances in an indoor residential environment. The four transmit antennas can send two spatial streams, one spatial stream into one antenna pair and another spatial stream into the other antenna pair. Additionally, the two antennas sending a single spatial stream and incorporate 2 element beamforming in order to increase the SNR at the receiver and improve the MCS level.

Equation 10

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} & h_{14} \\ h_{21} & h_{22} & h_{23} & h_{24} \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}$$

Equation 11

$$E_{\theta} = j\omega\mu \frac{2I_m}{\beta} \frac{e^{-j\beta r}}{4\pi r} \sin\theta \frac{\cos\left(\left(\frac{\pi}{2}\right)\cos\theta\right)}{\sin^2\theta}$$

The equation for the electric field of a half wave dipole in the far field is show in Equation 11. The elements of the channel matrix shown in Equation 10 can be determined from Equation 11 in the ideal case of half wave dipole antennas with a direct line of sight path between transmitter and receiver and no reflections from surrounding objects. The electric field is a complex directionally dependent vector with amplitude that falls off as the inverse of the distance between the two antennas, r in Equation 11. The wave length of a 5.5 GHz signal is 54 mm, the phase of the electric field cycles through 360 degrees every wavelength, the  $e^{j\beta r}$  in Equation 11. The electric field above the half wave dipole is zero when  $\theta = 0$  degrees in the floor directly above the AP antenna and the electric field is a maximum when the AP and STA are on the same floor  $\theta = 90$  degrees.

In a customers home, the channel is much more complicated with many reflected rays and attenuation from walls and floor. The path between scattering objects follows the formula of Equation 11. The elements of the channel matrix are the complex vector sum of all the direct and reflected rays between a transmit antenna and a receive antenna.

The condition for two spatial streams is met when the eigenvalues of the channel matrix multiplied by it's complex conjugate transpose has two nonzero values. This condition is not met unless there are multipath reflections. In free space there will only be one spatial stream. Fortunately, there is enough multipath in residential homes for two spatial streams and high MCS rates with a 4x4 AP and 2x2 STA.

Now we have everything needed to calculate the data rate of an OFDM symbol. The symbol period is 13.6  $\mu$ s. The number of data subcarriers in an 80 MHz channel width is 980 and twice that for 160 MHz

channel. The bit loading is determined by the MCS index. In Equation 12, the PHY rate is calculated to be 2400 Mbps for a 2x2 STA with 160 MHz channel width for MCS 11.

*Equation 12*

$$R_b = \frac{N_{ss} \cdot N_d \cdot b \cdot r}{T_{symbol}} = \frac{2 \cdot 2 \cdot 980 \cdot 10 \cdot \frac{5}{6}}{13.6} = 2400 \text{ Mbps}$$

## PHY Rate and Speed Test Relationship

Sometimes there is confusion between the data rate of an OFDM symbol and the download rate of a file from the Internet to the customers computer or the result of a speed test. They both after all, are called data rate and have the same units of Mbps. For example, Figure 4 shows an example where the data rate of the OFDM symbol is 2400 Mbps while the result of a speed test measures a download rate of 1200 Mbps. The PHY rate of the OFDM symbol is determined by the channel conditions while the 1200 Mbps measured on the speed test is determined by many factors include speed test server, Internet path, WAN, speed, computer settings, other spectrum users, etc.

To avoid confusion for these two related but not identical metrics, the data rate of a particular OFDM symbol is referred to as a PHY rate whereas the download speed of a file from the Internet to the customers computer is referred to as data throughput. Both are data rates, and both are measured in Mbps generally for the same signals, although bits/sec, MB/s, GB/s, kbps, Gbps are equally valid units of measure.

A PHY rate greater than the target data throughput is a necessary but not sufficient condition. If the PHY rate is less than the target data rate then it is not possible to get the download speed we are trying to deliver to our customers. For example, if the PHY rate is 17.2 Mbps for a 40 MHz channel width at MCS0 and one spatial stream, as will be explained later in the paper, the measured download throughput was 16 Mbps, it is simply impossible to deliver 1 Gbps download rate when the PHY rate is 17.2 Mbps. Not going to happen. The flip side of this line of reasoning is that many factors in addition to PHY rate determine the data throughput experienced by a customer. These factors include preambles, aggregation, spectrum sharing, scheduling, network processor loading, latency from Internet server to CMTS, DOCSIS feed to Wi-Fi. A PHY rate of 2400 Mbps does not mean that speed test results will always be 1200 Mbps. Because of this difference between PHY rate and data throughput, the PHY rates are often incorrectly called “theoretical”. Nothing could be further from the truth. A PHY rate of 2400 Mbps will happen in customers homes all the time if they have a 2x2 160 MHz Wi-Fi 6 network adapter in their computer. In fact, the download speed measured in static and radio silence channel conditions with tools designed to generate as much throughput as possible are much more “theoretical” than PHY rates. PHY rates quoted in the Wi-Fi 6 device specifications such as 2400 Mbps happen all the time whereas the throughput measurements made in radio silence and static channel conditions with tools to generate maximum throughput will never happen in customers homes. Customers do not have radio silence due to neighbors and other radio devices inside the home, people and pets and things will be moving in the home and outside the home creating dynamic channel conditions, and customers will be using applications that have a different download and upload demand than test tools.

# Resource Units and OFDMA

The FFT duration of a Wi-Fi 6 symbol is 12.8  $\mu$ s. For the standard guard interval of 800 ns the symbol period is 13.6  $\mu$ s. The ratio of the guard interval to the FFT duration is 1/16 so that guard interval is only 5.88% of the symbol period. The efficiency in this sense is 4 times greater for Wi-Fi 6 than preceding OFDM Wi-Fi versions. A 2x2 160 MHz Wi-Fi 6 STA has a peak PHY rate of 2400 Mbps. When not running a speed test or downloading a many GB file from a Gbps server the traffic demand is much less than required to fill up 160 MHz wide symbols. So the subcarriers of an OFDM symbol are divided up into smaller resource units allowing multiple STA's to create a single OFDM symbol. This is called OFDMA since multiple access to the resource is accomplished by allocating users resource units. Resource units, RU, are blocks of subcarriers making a full OFDM symbol.

OFDMA is a key part of Wi-Fi 6. OFDMA uses frequency division as a multiple access scheme. Rather than each STA taking turns in time using all the channel width of a symbol, multiple STA's can use different blocks within the channel width at the same time in a single OFDM symbol. Since the tones generated by each STA are formed from a synchronized FFT they are orthogonal. The tones from many STA's are called orthogonal because they do not interfere with each other. Orthogonal signals do not interfere with each other so they do not require the same guard band between resource blocks that non-orthogonal signals do. That is why OFDMA is much more efficient than simple FDM systems. An FDM system would be AM radio where each radio station is assign a block of spectrum and guard bands are needed between channels. Another example is DOCSIS 3.0 channel bonding. The QAM signals that are bonded are not orthogonal, while the symbol rate is 5.3 MHz the channel spacing is 6 MHz. Orthogonal tones in 802.11ax are spaced at 78.125 KHz equal to the symbol rate. As we've seen the guard interval needed to prevent inter-symbol interference reduces the spectral efficiency of OFDM signals.

For a 20 MHz channel width, there are 2 pilots in a 26-tone resource block, 4 pilots in a 52 and 106 tone resource block, 8 pilots in a 242-tone resource block, and 16 pilots in a 484 and 980 tone resource block. This is show in Table 6

**Table 5 - OFDMA tones and pilots for each resource unit, RU**


tones	pilots
26	2
52	4
106	4
242	8
484	16
996	16

In determining the allocation of subcarriers to resource units the 802.11ax task group started at the beginning; the original 802.11a OFDM symbol. The 802.11a signal is shown in Figure 13, the active tones are spread out over 20 MHz and the EVM of the 48 data subcarriers in red and 4 pilot subcarriers in white are shown. The FFT size of an 802.11a signal is 64, consisting of 52 active tones with 4 pilot tones.

52 tones with 4 pilots was selected as a resource unit in he mode. The 48 data tones for this RU size matches the data tones of an 802.11a signal. The occupied bandwidth of 48 tones spaced at 78.125 kHz is


4.1 MHz. Simulations of packing efficiency versus RU bandwidth during the 802.11ax task group work revealed that above 2.5 MHz the efficiency dropped. A smaller RU size was needed.

The 52 RU was divided into two. The smallest RU size was decided to be 26 tones with 2 pilots. The RU having 26 tones occupies 2.0 MHz of spectrum and has a maximum PHY rate of 29.4 Mbps with two spatial streams.




**Figure 13 – Back to the drawing board, 802.11ax task group formed resource units, RU, to match original 802.11a signal**

While it may seem that 802.11a signals are hopelessly old school in these times, in fact 802.11ax still makes use of 802.11a signals. Signalling messages are still sent with 802.11a protocol at 6, 12, 24 Mbps data rate in 20 MHz channels with 48 data subcarriers and 4 pilots for a total number of active subcarriers. As shown in figure 14, a block acknowledgement, Block ACK, after receiving successfully a 160 MHz channel width 802.11ax MCS11 1200 Mbps PHY data burst, the Block ACK does not just send one of these 802.11a 20 MHz channel width signals but eight of them. Each 20 MHz segment of the 160 MHz channel width contains an 802.11a block ACK.




**Figure 14 - 802.11a signal are still used in Wi-Fi 6, here for Block ACK**

To verify that each of the 20 MHz spectrum blocks are indeed 802.11a signals, the vector signal analyzer demodulator was set to a center frequency of 5520 MHz, just above the primary channel 100. The demodulator was set to 802.11a mode. The measured constellation is shown in figure 15. The modulation is 16-QAM for the block ACK with data rate 24 Mbps and 52 active tones. The duration in time of these block ACK's measured 28  $\mu$ s. The time duration of the block ACK is measured in figure 16. The block ACK contains two OFDM symbols and preamble. The symbol time for 802.11a is 4  $\mu$ s so the two symbols of the block ACK occupy 8  $\mu$ s. The 802.11a preamble occupies 20  $\mu$ s.



**Figure 15 - 802.11a constellation for 24 Mbps Block ACK 16-QAM**



**Figure 16 - Time duration of 24 Mbps 802.11a signal, 1/8 of the total ax Block ACK**

**Table 6 - OFDMA resource units RU, for 2x2 STA**

RU	2.0	4.1	8.3	18.9	37.8	76.6	MHz
MCS	26	52	106	242	484	996	tones
0	1.8	3.5	7.5	17.2	34.4	72.1	Mbps
1	3.5	7.1	15.0	34.4	68.8	144.1	Mbps
2	5.3	10.6	22.5	51.6	103.2	216.2	Mbps
3	7.1	14.1	30.0	68.8	137.6	288.2	Mbps
4	10.6	21.2	45.0	103.2	206.5	432.4	Mbps
5	14.1	28.2	60.0	137.6	275.3	576.5	Mbps
6	15.9	31.8	67.5	154.9	309.7	648.5	Mbps
7	17.6	35.3	75.0	172.1	344.1	720.6	Mbps
8	21.2	42.4	90.0	206.5	412.9	864.7	Mbps
9	23.5	47.1	100.0	229.4	458.8	960.8	Mbps
10	26.5	52.9	112.5	258.1	516.2	1080.9	Mbps
11	29.4	58.8	125.0	286.8	573.5	1201.0	Mbps

Table 7 shows the PHY rate for two spatial streams for each MCS index for each RU size. The RU size is expressed in MHz of occupied bandwidth in the first row and number of tones in the second row. The smallest PHY rate is RU 26 tones with MCS0 having a data rate of 1.8 Mbps. The largest PHY rate is 1201 Mbps for RU 996 tone occupying 76.6 MHz of bandwidth with MCS11. This should be familiar by now, the 1200 Mbps maximum PHY rate of a 2x2 80 MHz Wi-Fi 6 phone.

Frequency division multiplexing and time division multiplexing are duals of each other. Two stations running at 80 MHz channel width on the same channel and 50% duty cycle have equal throughput to two stations running at 40 MHz channel width with 100% duty cycle on separate channels. There is no inherent advantage in throughput of one method over the other.

With perfect implementation running at twice the channel width half the time equals running at half the channel width all the time. It is in the imperfections of the implementation that in certain cases one may have advantages over another.

OFDMA implementation in the mode is a good example. In some cases stations are better off running at full channel width and sharing the same channel either with time division or with MU-MIMO spatial division multiplexing based upon antenna beam forming.

While in other cases stations are better off each running in a separated block of subcarriers within the channel width at the same time.

OFDMA and MU-MIMO are complementary technologies. MU-MIMO works best in high SNR regions with high throughput demand. OFDMA works best in low SNR regions with low throughput demand. Consider a 2x2 160 MHz Wi-Fi 6 STA, the PHY rate is 2400 Mbps.

By multiplying 2400 Mbps PHY rate by the symbol time of 13.6 microseconds we see that each symbol needs to be loaded with 32,640 bits. Most applications just do not generate enough bits consistently fill 2400 Mbps data rate symbols. When traffic demand is low and PHY rate high, latency, memory overloading, and zero filled symbols negatively impact customer experience. With a 2400 Mbps PHY rate it is unlikely that enough stations each with low traffic demand could fill the pipe.

So while the 2400 Mbps PHY may be inefficient at low traffic demands, it is a moot point since it still works fine. But as the users get further away from the AP and more walls, floors and obstacles attenuate the signal, the PHY rate drops. At lower PHY rates stations can benefit from OFDMA. A station can only transmit at a finite level, many are limited to around +20 dBm of transmit power.

Transmit power of stations need to be restricted for regulatory reasons and for battery life. Without OFDMA, two stations may transmit at 80 MHz channel width taking turns in time. The +20 dBm transmit level of each station is spread out over the full 80 MHz. With OFDMA, the two stations can transmit each in an adjacent 40 MHz channel width at the same time. The +20 dBm transmit level of each station is now spread out over the 40 MHz channel. The power spectral density of each station transmission is increased by 3 dB.

The signal to noise ratio at the AP receiver is increased by 3 dB. With 4 stations running in uplink OFDMA mode the SNR increase is 6 dB. With 8 stations running in uplink OFDMA mode the SNR increase is 9 dB. A 9 dB increase in SNR can increase the MCS level, resulting in a throughput improvement.

## Speed Tests and File Downloads

Figure 17 shows a screen shot of a 2x2 160 MHz Wi-Fi 6 STA inside a notebook computer. The speed test result is 1129 Mbps download and 38.26 Mbps upload. The PHY rate is consistent 2400 Mbps, the mode, 160 MHz channel width, MCS11, 1024-QAM, 5/6 LDPC code rate, 2 spatial streams. An iperf3 download from the 2.5 Gbps Ethernet LAN port of the router to the STA measured a download throughput of 1.43 Gbps.




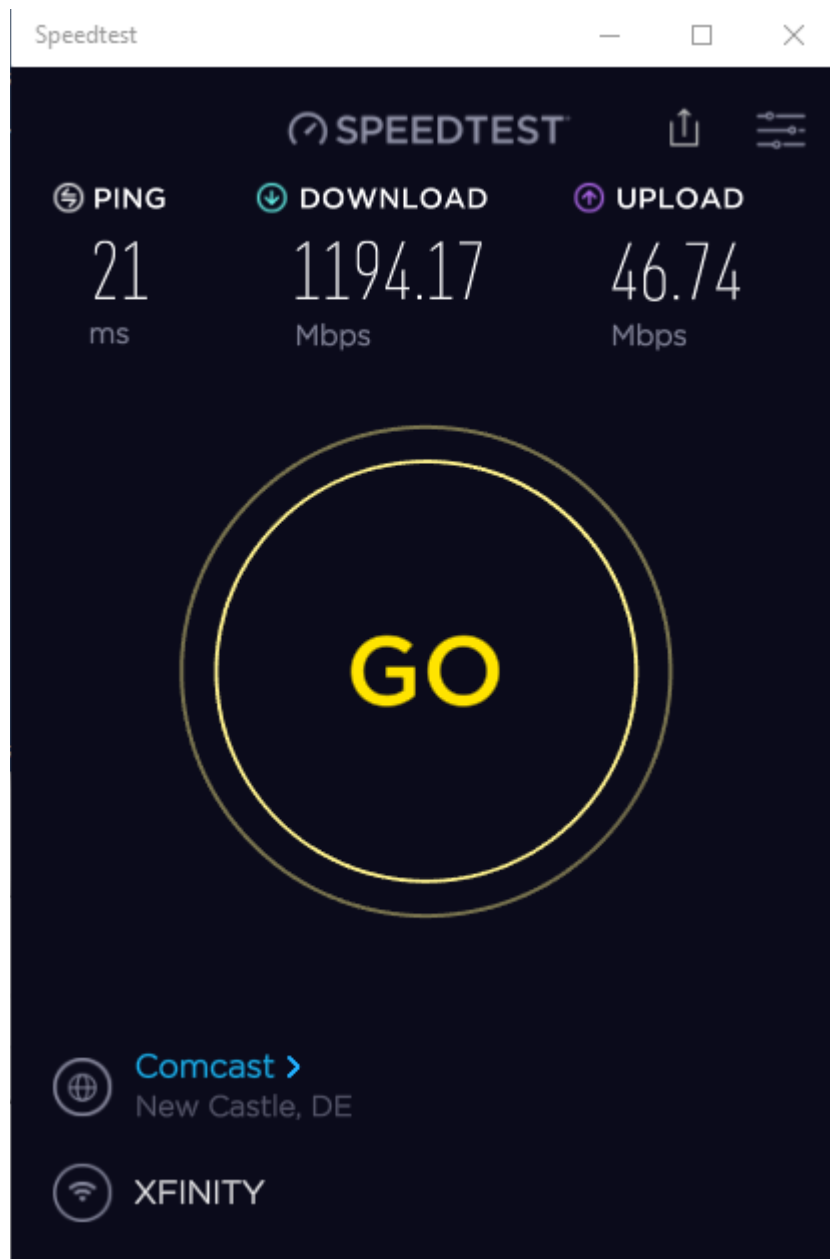



Figure 17 - Screen shot of notebook computer with 2400 Mbps Wi-Fi 6 2x2 160 MHz station



**Figure 18 - Speed test with 2400 Mbps PHY notebook computer 2x2 160 MHz Wi-Fi 6**

Figure 19 shows the screen shot of a speed test with a 2x2 80 MHz Wi-Fi 6 phone. The download speed was 838 Mbps and the upload speed was 39.9 Mbps. The PHY rate was 1200 Mbps, the mode, 80 MHz channel width, MCS11, 1024-QAM, 5/6 LDPC code rate, 2 spatial streams.



**Figure 19 - Speed test with phone having 1200 Mbps Wi-Fi 6 2x2 80 MHz station**

## Conclusion

Wi-Fi 6 based on the IEEE standard 802.11ax adds many critical features for service providers. Wi-Fi 6 is key to delivering consistent, reliable, stable Gbps service with whole home coverage. First, the speeds are higher, with 1200 Mbps symbol data rate for many phones and 2400 Mbps symbol data rate for many notebook computers. The paper has shown examples of the phone delivery speed test results of 800 Mbps and the notebook computer delivering speed test results of 1200 Mbps. Second, Wi-Fi 6 adds OFDMA as a new multiple access network. This allows many devices each with a small amount of traffic demand to all be served with a shared symbol, each device using only a small number of subcarriers of the OFDM symbol. Ensuring that the customer traffic demand is much less than the offered capacity is the key to consistent, reliable, stable service. Delivering 1200 Mbps speeds with a 2400 Mbps symbol rate allows for other users of the unlicensed spectrum such as neighbors as well as many other lower traffic demand devices in the customers home.

## Abbreviations

AP	access point
STA	station

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
MIMO	Multiple input multiple output
OFDMA	Orthogonal frequency division multiple access
FFT	Fast Fourier Transform

# **Next Player Video Service**

## **The Case For Bringing Playlists to TV**

A Technical Paper prepared for SCTE•ISBE by

**Arash Pindari**

CEO.

VionLabs AB.

Drottninggatan 25, 111 51 Stockholm, Sweden.

+46 8-410 163 16.

arash@vionlabs.com.

**Giles Wilson**

CTO.

VionLabs AB.

Drottninggatan 25, 111 51 Stockholm, Sweden.

+46 8-410 163 16.

patrick@vionlabs.com.

**Michael Eagles**

Director Strategy & Innovation.

Liberty Global.

Boeing Avenue 53, 1119 PE, Schiphol Rijk.

+31 20 7788339.

meagles@libertyglobal.com.

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	4
Conclusion .....	15
Abbreviations.....	17
Acknowledgements .....	18
Bibliography & References .....	18

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Video Viewing Is Growing But Broadcast Live Linear TV Is In Decline .....	4
Figure 2 - Past, Current and Future TV Viewing .....	5
Figure 3 - Increasing Competition For Viewer Time.....	6
Figure 4 - Personal Carousels Made into Channels.....	8
Figure 5 - Path To Future For Emotional Content Fingerprinting.....	9
Figure 6 - Machine Process Content To Produce Emotional Content Fingerprint.....	12
Figure 7 - Emotional Content Fingerprint A/B Test Shows Improved Performance.....	13
Figure 8 - Next Player Ad Revenue Opportunity.....	14

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Technology Enablers and Viewing Behavior Evolution .....	7
Table 2 - Next Player Existing vs New Components .....	9

# Introduction

Live TV was, and still is, a brilliant discovery tool, but it needs a makeover. Sooner rather than later.

Live TV scheduling is all about audience flow. The ability to engage the viewer and retain them within that live broadcast channel is the primary function of good broadcast scheduling and how live TV has been able to maintain its dominance over viewing minutes for so many years.

When pay-tv was young, the linear channel was the only way of finding new content, with its limited possibilities it still served as a great tool in guiding viewers to content and providing a quick way of sampling a large amount diverse content. The TV channel logo and brand helped guide to the direction of channels and provided comfort in what to expect behind a certain logo. You would expect to find nothing but boats behind the “Sailing Channel” logo, and inspirational mood programs behind the “Travel Channel” logo. A channel without a clear enough content profile would often become less and less relevant and lose viewership.

In an SVOD and non-linear world the brands of the broadcasters play a less important role in guiding between content and the brand properties of the logos have been given less relevant. Instead, we can see playlists, recommendations and series emerging as the guiding stars for this. In this paper we will show how a strong playlist capability (“Next Player”), that uses Machine Learning and A.I. to compare similarities in types of content, can provide a good proxy for that scheduling art form.

Whilst the possibility of watching any content asset at any given time represents a superior customer value to the defined time slots of the linear channels, it also represents an expectation of an increasingly active customer and places a lot more of the burden of discovery on the consumer. This is a burden many customers are neither capable, nor willing, to shoulder. The lack of guidance in the user experience is now creating a problem for consumers in navigating and finding new content. We intend to show that this burden can be eased through content similarity analytics for matching, combined with user-created profiles; shifting the audience flow story and related value creation opportunity towards the platform.

The effect of the discovery problem we describe has been an increased amount of the consumers’ time being spent on the process of discovery instead of watching content - searching for content or becoming paralyzed by the many choices. This in turn creates the perception of a significantly smaller content catalog than most players actually have available. By introducing the Next Player concept, we aim to reduce customer frustration and increase satisfaction as well as re-invigorating the traditional TV business model by modernizing the Linear TV viewing experience and creating new advertising inventory.

# Content


## THE BIG PICTURE

With live linear broadcast TV in decline, an explosion of choice is emerging in direct to consumer video. At the same time there are new enabling advances in personalization and content metadata emerging. In the context of these trends, there exists an opportunity to re-invent the lean-back live TV experience so viewers can spend more time watching TV and less time searching and browsing.

### Live Linear Broadcast TV Is In Decline

Over the last 5 years, from 2015 to 2019, viewing and engagement with traditional linear broadcast TV has been declining at an increasing rate. Customers are substituting linear TV viewing by viewing content outside of the platform. Integration of direct to consumer video applications is becoming relatively common place for TV platform operators and these integrations are cannibalizing traditional “tune-in” live linear TV viewing.

As shown in Figure 1, over an 8 years period the live linear TV consumption in the UK will have declined each year (almost 30% from 225 to 160 minutes per person per day), whilst overall video viewing will have increased 11% from 258 to 288 minutes per person per day. It’s clear that new direct to consumer video consumption applications are gaining traction and this has given rise to an explosion of choice for viewers.



**Figure 1 - Video Viewing Is Growing But Broadcast Live Linear TV Is In Decline<sup>i</sup>**




Several studies in early 2019 have revealed that the number of linear TV ads seen by 16 to 34-year-old viewers has fallen by a fifth compared with a year ago <sup>ii</sup>. Recently Bank of America Merrill Lynch warned that “seismic changes in the viewing habits of young people and sliding TV advertising revenues would hit broadcasting revenues and put pressure on broadcaster margins” <sup>iii</sup>. They also noted that “TV broadcasters were underestimating the declines in audiences and the threat from digital video services” <sup>iv</sup>. In mid-2019 OFCOM noted that, for the UK market, on-demand usage has increased “driven by increased use of subscription video-on-demand services such as Netflix and Amazon Prime Video”, noting no change in the proportion of people watching broadcasters free catch-up services (BBC iPlayer, ITV Hub, All4 and My5) and declines for some age groups <sup>v</sup>. The report further drew attention to the fact that “declines in the amount of time spent watching live TV are more pronounced among younger audiences” with 18-24s watching 18 minutes less live TV per day than the previous year <sup>vi</sup>.

It is clear that digital video services have experienced an “explosion of choice” across a range of new direct to consumer platforms, which is appealing to younger audiences and putting the traditional broadcast TV models under pressure.

### An Explosion of Choice

Historically the TV experience was based on “scheduling”, and the choices for the viewer were limited to zapping channels. With the evolution of PVR, Video On-Demand, and Replay TV that content could be viewed at different times of day at the discretion of the consumer. Now, with the addition of new viewing devices and direct to consumer app TV there is more choice than ever before.

Today, the viewer is faced with an array of devices and different user experiences for discovering content across many platforms, forcing them to spend more time in the discovery process. As shown in Figure 2, in the past the “what to watch” discovery process was relatively simple as it was largely “time” bound. Today, with the explosion of choice in direct to consumer video the discovery process more challenging and suited to younger audiences who are able to switch between apps and devices. This enables not just “time-shifting”, but also “place-shifting”, with viewing on potentially an array of different “devices” and multiple “TV apps”. Thus it creates a new, more fragmented, viewing landscape.



**Figure 2 - Past, Current and Future TV Viewing**

With emerging personalization capabilities, however, we may see a simplification of that fragmented viewing landscape in the future, where the TV experience follows the viewer; based on a machine-generated, personalized schedule or playlist. This will allow the viewer to spend less time browsing through endless box art tiles on an array of devices and more time enjoying the content.

What are the key market factors that may be accelerating the need for a personalized schedule or playlist? We see that despite viewers' limited time to discover TV content, there is increased competition in the video services market.

### Increasing Competition for Viewer Time

We expect that with the addition of new streaming video services the amount of time spent browsing on each service will be even less. With Disney+, Apple TV Plus, HBO Max, Viacom and NBC Universal all launching into an increasingly competitive market <sup>vii</sup> and with more direct to consumer video services and subscribers <sup>viii</sup>, we expect further pressure on the time a viewer has to spend searching for something to watch.

Netflix noted in a 2016 paper <sup>ix</sup> that the average subscriber spends 60 to 90 seconds scanning movies and TV shows on the platform before giving up; and in that time, the subscriber will roughly 10 to 20 titles—about three in detail—on one or two screens. If a viewer is splitting time between 3 different video services, there may only be 20 to 30 seconds of browse time per catalog, making first time right playlists and recommendation more critical than ever.



**Figure 3 - Increasing Competition For Viewer Time**

What enablers becoming available to address these challenges? We next explore what the technology enablers are behind the current explosion in choice, and how emerging personalization, artificial intelligence and machine learning may activate the next wave of market models and behaviors.

### Technology Enablers of New Market Models & Behaviors

In the 1990's digital multi-channel zapping was the predominant TV experience. In the 2000s, with the emergence of low-cost personal storage, it became possible to record TV content. This enabled a new viewing behavior - "Time shifting" – in which TV content could be viewed at any time convenient for the viewer. By 2010 the Internet and Wi-Fi technology were ubiquitous enough to enable TV viewing on an explosion of new devices such as phones, tablets and HDMI sticks. This allowed TV viewing to be "Place shifted", with viewing no longer constrained to just the home or the living room. In turn, this also resulted in new types of TV apps and a direct to consumer model and a new type of viewing behavior - "Snacking" on TV content fueled by the rise of short form internet content providers.

As we enter the 2020s generation of viewing enablers we see that they are Personalization, Compute and Artificial Intelligence (AI). Applied to content as enablers these can support a move to "Person shifted" TV viewing, in which personal TV playlists can follow the viewer from device to device and computed or content fingerprint can form the basis for adaptive personal channels with new and interesting content options for the viewer without ever leaving the viewing experience.

**Table 1 - Technology Enablers and Viewing Behavior Evolution**

Format	Type	Year	Enabler	Browse Behavior	Viewing Behavior
Live linear Multi-channel	Scheduled	1990s +	Digital	TV Zapping/ Tune-in	Tune-in, Lean back
PVR, Video On-Demand, Replay	Time shifted	2000s +	Storage/ Path Return	EPG/ Interactive/ PVR menu	Recorded, Lean forward
New Viewing Devices, OTT, App TV	Place shifted	2010s +	W-Fi/ Internet/ Mobile App Stores	TV Apps "Launch"	On The Go Snacking
Personal Playlist, Adaptive-channel	Person shifted	2020s +	Personalization, Compute, AI, "Emotional Content Fingerprinting"	Playlist/ Follow-me	"Next" Binge

**Next**

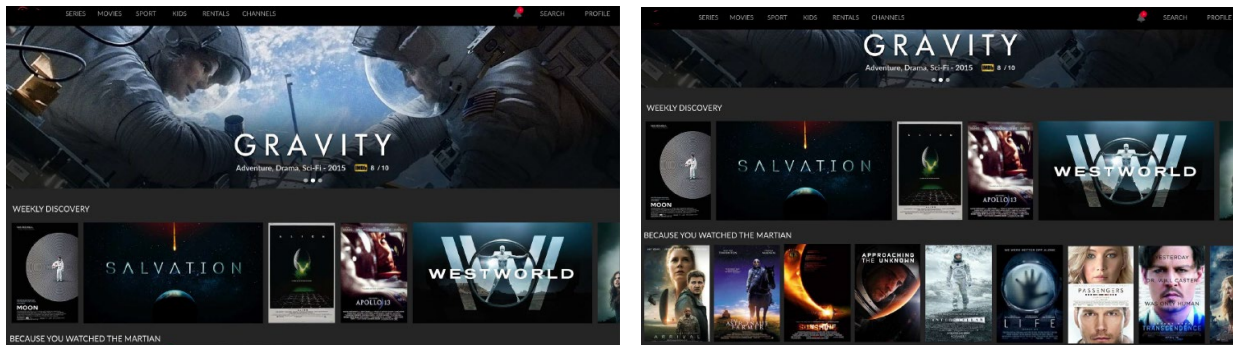
To support this hypothesis we propose a hypothetical "Next Player" that will leverage Personalization, content-focused AI and "Emotional Content Fingerprinting" to bring this new TV experience to the market and enable this next wave of viewing behavior.

## PROPOSED NEXT PLAYER

### What Is The Next Player ?

The proposed “Next Player” is video discovery service that can be presented to the viewer in the EPG, within the context of the viewers’ personalization profile, but which is built as a playlist of on-demand assets. Of course, the challenge in doing this is in determining what content should be included in the playlist, in what order, and at what time in order to make the channel both valuable and engaging for the consumer.

These are similar classes of problem to that of producing on-demand recommendations for a consumer in a carousel, but with the added complexities of needing to take the selected personalization profile and most appropriate ordering of content into account.



**Figure 4 - Personal Carousels Made into Channels**


### What Components Are Existing ?

In developing the proposed Next Player, we assume that in an MVDP there will be existing content metadata, on-demand platform, content library, recommendation system and personalization service underpinning existing video services and available to support the proposed Next Player.

### What's Components Are New ?

The key new element needed for the proposed Next Player is the addition of an ‘Emotional Content Fingerprinting’ service that can be combined with viewer personalization, to create personalized playlists per viewer.

For our approach we can summarize the existing and new components follows:



**Figure 5 - Path To Future For Emotional Content Fingerprinting**

1. Existing: Typical components that an MVDP will already have in place; such as existing Recommendation and Personalization service to support per viewer preferences.
2. New: The expected addition of a new “Emotional Content Fingerprinting” to produce high quality content similarity playlists.

**Table 2 - Next Player Existing vs New Components**

Existing	New
<ul style="list-style-type: none"> <li>- On-Demand platform</li> <li>- Uniform Content Identification Structure</li> <li>- Recommendation platform</li> <li>- Personalization service</li> <li>- Bookmarking service</li> <li>- Ad insertion services</li> </ul>	<ul style="list-style-type: none"> <li>- Emotional Content Fingerprinting for Next Generation Content Similarity.</li> <li>- Content Discovery In-line (i.e. Ability to skip to “Next”)</li> <li>- Personalized Playlist Discovery Via The EPG</li> </ul>

## **Existing Components**

### *On-Demand Assets*

The foundation for the assets used as input in the creation of the weekly/daily Discovery can be driven from both SVOD, TVOD, Replay, Catch-up and Linear TV viewing as long as a uniform content identification structure is in place.

### *Uniform Content Identification Structure*

The Next Player requires an existing uniform content identification structure, provided either by the MVDP or a 3<sup>rd</sup> party that can support a timely, consistent content identification between parties.

### *Recommendations*

A recommendation service is a key foundation. The proposed way to create a weekly discovery is to base it on a watch history of the last assets viewed. If a larger sample of viewership exists we recommend that it be parsed based on time of day in order to create a channel that is dynamic over the day and that shifts with the time based viewing patterns. Once a customer has reached the end of a series, a new similar series or a movie could be recommended.

### *Personalization*

A personalization service is needed to implement relevant recommendations. Channels are personal. Instead of the old way of serving one channel to all, the channels that we present are completely tailored to an individual or at the very least based on an existing personalization service for the viewer rather than being generic for the household. The requirement for amazing personalization is higher. If we fail to capture the viewer and do not show an in-depth understanding of the content viewed, the channels will become irrelevant and will fail.

### *Bookmarking*

A bookmarking service is a key foundation, as “continue watching” is a key user experience capability. One of the most common behaviors in OTT/On-demand viewing is resuming the last watched episode or movie and therefore bridging that experience into the main screen and lean back EPG behavior is a very natural step. This Channel-type can be created from a separate Movie title and would then function in a similar fashion as the “Because you watched...” channel creation, or would be based on a series viewership where the next available episode would automatically start playing.

### *VOD Ad Insertion*

In order to support advertising within the Next Player playlist, an VOD ad insertion platform able to support the ad decision and ad serving, is also a key foundation.

## **New Components**

### *Emotional Content Fingerprinting*

Determining content similarity is the key new capability required for the generation of the “Next Player” channel. Excellence in this step is critical to providing programming that will represent value to the user.

Traditional approaches to determining content similarity have been based on the use of metadata, including such items as genre, release year, popularity and production principals (actors, directors, producers etc.). Provided that the metadata is consistent and of high quality (which it is not always) these techniques can provide a reasonable level of similarity measure. However, we believe that such methods miss out on many of the characteristics of audio-visual content that truly establish similarity. As the old adage goes “a picture is worth a thousand words” – and trying to condense the essence of a multimedia asset into even a moderate number of metadata fields is an unsatisfactory approach.


Over recent years advances in cloud computing and the ability to leverage high performance GPUs have in turn led to the democratization of artificial intelligence and advanced data analytics. In particular there has been much progress in the fields of deep learning neural networks with respect to network topologies for varying applications and training techniques. We can leverage these advances to enable content fingerprinting and similarity determination.

## HOW IT WORKS


In order to determine how similar content assets are we compute numerical fingerprints for each piece of content using a variety of deep learning neural networks. These neural networks extract features from the video and/or audio components of the content over the complete timeline of the assets. Producers of content use a variety of visual and audio tools and techniques to tell stories, convey emotion and generally influence the consumers perception of the content. Examples are the emotional cues contained in soundtracks, color palette employed, stress levels induced by speed of cutting etc. By analyzing the content for these features and determining how they develop through the timeline of an asset we are able to generate fingerprints that truly capture the essence of each asset in dimensions beyond just the subject matter.

Examples of the feature timelines are show in Figure 4, below. These demonstrate the development of emotional dimensions “valence” and “arousal” through a movie. Valence is a measure of how positive or negative an emotion is, Arousal measures the amount of energy in the emotion from calm and relaxed to intense and energetic. Such information is much richer than the aggregated information contained in conventional metadata and the inclusion of the sequence data in the fingerprinting of assets leads to vastly superior results in determining content similarity.

Once the numerical fingerprints for the assets are obtained it is possible not only to determine which assets are similar, but also to determine how similar they are and hence relative similarity measures amongst sets of content. This information can then be used to generate final sets of content for playlists.



*No depth or timeline, single words to describe complex stories*



*More depth in the data as you can follow a stories evolvement*


**Figure 6 - Machine Process Content To Produce Emotional Content Fingerprint**

*A/B testing of the service to validate improvement*

Together with the one of Nordics biggest TVOD and EST streaming services SF Anytime VionLabs A/B tested purely metadata based result with results that included emotional fingerprints.

The outcome of the test that started 2018 November on Web browsers was on 50% of the viewer base (total 900,000 viewers) over 6 month time. We saw a clear positive effect from using emotional fingerprints where the relevance of recommendations were leading to 27% more content added to the watch list, directly leading to an overall increase of 11% in total number of transactions. We attribute this success directly to the increased relevance of the recommendations through the strong correlation we saw where someone who added a piece of content to the watch list were 7x more likely to make a purchase. To summarize, more relevant recommendations leads to increased engagement (content added to watch list) which leads to increased revenue (11% increase in total transactions)"





**Figure 7 - Emotional Content Fingerprint A/B Test Shows Improved Performance**

The engagement on SVOD was similar where the viewers with the emotional fingerprints added 3 times more content to their watch list and raising the amount of started streams with 25% monthly. Test also done over 6 months in 2018 on Web browsers on over 1,100,000 viewers.


## THE KEY OPPORTUNITIES

We explore the key advantages of the proposed Next Player including a). New services and advertising revenues, b). Delighting viewers and improved Net Promotor score; and c). As a result lower subscriber churn or cord cutting.

### New Services & Advertising Revenues

The Next Player concept activates a new advertising revenue possibility and creates currently untapped advertising inventory. The viewer acceptance of advertisements in a linear stream is higher than that during on-demand viewing. By creating linear channels we provide new inventory for ad-sales. Further, the Personalization of the channel has the possibility to further add to that value and enable a new level of targeted ad's.

We propose an estimated revenue model that is based on platform reach and share of total viewing attributed to the "Next Player". Assuming an eCPM of USD 25 based on Video on Demand ad inventory pricing <sup>x</sup> and a basis of one Ad per 30 minutes of viewing, we can proceed to model the revenue potential against the foundation penetration of the Next Player. We assume an upper limit of 80% of the video homes will have TV devices (i.e. Set-top) capable of running a "Next Player" application and an assume an engagement level upper limit that represents up to 10% of all the minutes spent viewing TV. In our scenario, the upper limit combination of 80% reach and 10% of viewing share results in a revenue potential of almost USD 14 per advanced video subscriber per year.



**Figure 8 - Next Player Ad Revenue Opportunity**

#### Improved NPS and Reduced Platform Churn

As competition increases for viewing between platforms it will become important to retain subscribers. Those platform able to benefit from increased engagement levels we could also reasonably expect to see lower churn and higher Net Promotor Score (NPS) for video platforms that deliver on the “Next Player” concept. Specifically, the Next Player viewing type would be expected to appeal to the younger TV viewing demographic - resulting in higher NPS and less churn for that group.

#### Deeper Viewing of the Content Library

We would also expect the advanced algorithms employed for content similarity to allow the Next Player to expose assets to consumers that would not normally be discovered through conventional means, allowing better exploitation of the entire content catalog. As a result, we expect the impact of personalized recommendations for a Next Player would also provide a boost to the perceived catalog size offered to the viewer vs the actual catalog size; where effective catalog size (ECS) is the spread of viewing across catalog items. Netflix research<sup>xi</sup> noted that recommendations provide an effective catalog size up to 4x larger when personalized vs popular only.

## THE KEY RISKS

There is a risk that personalization and compute-based metadata capabilities result in viewing bubbles if poorly implemented, for example a) scale and device reach is limited; b) the personalization implementation is limited or not used by the viewer; and c) where new content paths are not made available the viewer limiting new experiences and making the Next Player less engaging <sup>xii</sup>.

### On Demand Content

It is clear that for the Next Player concept to be successful it requires scale, so that it can be used by the largest possible base of viewers. Critical factors in achieving that scale include an existing VOD platform, existing Replay and Catch-up content library.

### Personalization Capabilities

To achieve the scale needed, a key risk is the lack of cross platform personalization capabilities. At a minimum it is envisaged that the same personalization service could be leveraged between device viewing types TV, Tablets, Phones, and PC's.

### Limited Paths To New Content

A key risk is that the Next Player implementation is unable to easily introduce new types of content into playlists, resulting in a “Rabbit-hole” type execution where the viewer is never exposed to new content. This is a critical consideration, as limiting the seeding of new content in the playlist can potentially limit the engagement and repeat viewing of Next Player playlists.

### Content Length Considerations

Platform in the music industry have demonstrated the successful use of personalized playlists to drive platform engagement, with Spotify's Weekly Discovery; and also music video with YouTube Music's “Your Mixtape Endless Personal Music”. Selecting a 3 minute song as the media for a personalized playlist, however, would most likely allow for a higher tolerance for inaccuracy than a 30 minutes TV show, or a 120 minute movie.

## Conclusion

The paper has outlined the concept of the “Next Player” application – a personalized linear channel of playlist, on-demand assets built on a combination of existing video platform capabilities and the additional of deep understanding of content similarity with the available content catalog.

The need for this is based on an understanding of a changing pattern of user behavior enabled by new services, applications and technologies. It is believed that these types of approach can be employed to re-energize the traditional TV business model and increase consumer engagement by providing tailored content that appeals to them on a personal level whilst removing the need for active decision making on their part.

Additionally it has been shown that there is the potential for significant revenue gain through the availability of new advertising spots within the “Next Player” service. This leverages the insight that viewers have a higher toleration of advertising within linear channels.

Further we outlined the potential risks and pitfalls; particularly in the areas of readiness and execution; including 1). on demand content needs to be available, 2) personalization capabilities need to be in place, 3). avoiding “Rabbit-hole” execution where viewers never experience new types of content; and 4) Content length considerations and the tolerance for accuracy.

In summary, with the explosion of video choice, we see a risk that viewers will become paralyzed by the multitude of options, with the viewers having a diminishing share of time to browse and discover new content within each video service.

Now is the right time to leverage advances in Machine Learning and AI to develop an appealing and engaging “lean-back” experience back to video consumption. With an improved content similarity capability, combined with platform personalization services, and the right execution, there is an opportunity to revitalize the linear TV experience; taking content assets that largely already exist on platforms today but bringing them out of the traditional on-demand catalog and into a new playlist-based viewing experience.

## Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
eCPM	Effective Cost Per Mille
ECS	Effective Catalog Size
Emotional Content Fingerprinting	AI-based analysis of content libraries for emotional similarity.
EPG	Electronic Program Guide
GPU	Graphics Processing Unit
JSON	JavaScript Object Notation
Linear TV	The Experience of Viewing Content In A Sequential Manner
Live Linear TV	Traditional Scheduled Live Broadcast TV
ML	Machine Learning
MVDP	Multichannel Video Distribution Platform
NPS	Net Promotor Score
Personalization	Typically User Created Profiles Representing an Individual
Playlists	A Manifest of Content For Playing
PSB	Public Service Broadcaster
SCTE	Society of Cable Telecommunications Engineers
SVOD	Subscription Video on Demand
TVOD	Transactional Video on Demand
VOD	Video On Demand

# Acknowledgements

We thank the Patrick Danckwardt at Netflix for his early thinking in this area and exceptional video product thinking. We are also grateful to Dan Stevenson at Ampere Analysis for outlook around UK BARB data.

## Bibliography & References

- 
- <sup>i</sup> Broadcasters' Audience Research Board (BARB) and Ampere Analysis forecasts. Publication date: Mar 2018
- <sup>ii</sup> "UK broadcasters face advertising tipping point, new study finds", Mark Sweney, The Guardian, Jan 31, 2019, Jan 31, <https://www.theguardian.com/media/2019/jan/31/uk-broadcasters-face-advertising-tipping-point-new-study-finds>;
- "Is the decline in young TV audiences accelerating?", Omar Oakes, Campaign, Mar 13, 2019, <https://www.campaignlive.co.uk/article/decline-young-tv-audiences-accelerating/1578795>
- <sup>iii</sup> "ITV sinks after analyst warning on deterioration in TV viewing", Kate Beioley, Jan 17, 2019, <https://www.ft.com/content/4c196140-1a53-11e9-9e64-d150b3105d21>
- <sup>iv</sup> "ITV sinks after analyst warning on deterioration in TV viewing", Kate Beioley, Jan 17, 2019, <https://www.ft.com/content/4c196140-1a53-11e9-9e64-d150b3105d21>
- <sup>v</sup> OFCOM Communications Market Report 2019, OFCOM, Publication Date: Jul 4, 2019, p. 4.
- <sup>vi</sup> OFCOM Communications Market Report 2019, OFCOM, Publication Date: Jul 4, 2019, p. 6.
- <sup>vii</sup> "Disney+ predicted to take on Netflix and Amazon", Alana Foster, IBC, Jun 3, 2019, <https://www.ibc.org/consume/disney-predicted-to-take-on-netflix-and-amazon/3901.article>
- <sup>viii</sup> "Global gross SVOD subscriptions will climb by 119 million in 2019: Digital TV Research", Simon Murray, MediaNews4U, Jun 3, 2019, <https://www.medianews4u.com/global-gross-svod-subscriptions-will-climb-by-119-million-in-2019-digital-tv-research/>
- <sup>ix</sup> "The Netflix Recommender System: Algorithms, Business Value, and Innovation", Carlos A. Gomez-Uribe and Neil Hunt, Netflix, Inc., ACM Transactions on Management Information Systems, Vol. 6, No. 4, Article 13, Publication date: December 2015.
- <sup>x</sup> "Big Jump Seen In Ad Inventory on Connected TV", BroadcastingCable, Jon Lafayette, Nov 6, 2017, <https://www.broadcastingcable.com/news/big-jump-seen-ad-inventory-connected-tv-170139>
- <sup>xi</sup> "The Netflix Recommender System: Algorithms, Business Value, and Innovation", Carlos A. Gomez-Uribe and Neil Hunt, Netflix, Inc., ACM Transactions on Management Information Systems, Vol. 6, No. 4, Article 13, Publication date: December 2015.
- <sup>xii</sup> "YouTube's Product Chief on Online Radicalization and Algorithmic Rabbit Holes", Neal Mohan, New York Times, Mar 29, 2019, <https://www.nytimes.com/2019/03/29/technology/youtube-online-extremism.html>

# **The Imperative of MSO Future Wireless**

## **Strategies, Services and Architectures**

A Technical Paper prepared for SCTE•ISBE by

**Drew Davis**

Executive Director  
Cox Communications  
6305 Peachtree Dunwoody Road, Atlanta, GA  
++1 (404) 269-0405  
Drew.Davis@cox.com

**Anish Kelkar**

Partner  
Bell Labs Consulting  
600 Mountain Avenue, Murray Hill, NJ  
+1 732 208 9692  
Anish.kelkar@bell-labs-consulting.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
1. Background.....	3
2. Architecture principles of the future MSO wireless network.....	5
3. Future Wireless Architecture – Converged and Virtualized with massive automation.....	6
4. Foundational Network Architecture- the starting point for the future .....	7
5. Full MVNO, Fixed Wireless and Private LTE architecture .....	7
6. Investment comparisons .....	9
Conclusion .....	10
Abbreviations.....	10

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Requirements of the future wireless services .....	5
Figure 2 - Future wireless architecture .....	6
Figure 3 - Foundational architecture for future wireless services .....	7
Figure 4 - Full MVNO service architecture.....	8
Figure 5 - Fixed Wireless service architecture.....	8
Figure 6 - Private LTE network architecture .....	9
Figure 7 - CAPEX investment comparison for foundational v/s full architecture.....	9



# Introduction

The traditional MSO business is facing unprecedented competitive threats resulting in customer churn, revenue decline and margin erosion. Telephone companies (Telcos) are deploying 5G services which they expect will provide a consumer experience comparable to the current cable experience. Over-the-top (OTT) players threaten to convert the MSO service to a connectivity play and skim the cream off the top by delivering high margin consumer and enterprise services. We already see this competition show up in the landscape in all segments such as:

## Consumer

- 5G Home broadband at 300Mbps
- NFL/Venues immersive experience
- Live TV streaming services

## Enterprises

- Fixed 5G for primary/secondary connection
- 5G enabled eHealth for hospitals
- 5G for education and training

## Public sector

- 5G First Responder –Public Safety
- Smart City applications

To reverse the revenue and margin declines, MSOs are planning new wireless services. These services improve competitive positioning of current consumer and enterprise products and add net new revenue through new product offerings.

However, wireless is a new business for MSOs, who often lack adequate resources to accelerate its deployment. Strategy teams are experimenting with many different business models and use cases, which result in product churn and indeterminate timelines of the new portfolio. With these challenges, it is important for MSOs to identify a low risk foundational architecture to learn, experiment and become technically and operationally ready to launch new wireless services aligned with market needs.

In this paper, we will discuss the future architecture which will support new wireless services. We will develop a foundational architecture common to most services, and which enables experimentation and readiness development in short order. We demonstrate that the foundational architecture needs 20-30% CAPEX investments of a service architecture and allows MSOs to prepare for a vast number of services. Adopting this approach, MSOs will be able to counter the lack of clarity of the future wireless imperative and reduce product development cycles through an investment protected architecture.

# Content

## 1. Background

MSOs are facing immense competitive threats to their traditional businesses. Telcos and new entrants are experimenting with 5G, LEO (Low Earth Orbit) satellites and HAPS (High Altitude Platform Station) to

expand their consumer broadband solutions to broader coverage areas, offer higher speeds and lower tariffs. Current and new OTT players are offering enterprise and consumer services that threaten to convert the MSO offerings into a low margin connectivity play. These threats, if not addressed will end up in lowering MSO revenues and margins. Hence, many MSOs are considering new wireless services to improve the competitive positioning of their current services and expand their product offerings. The main objective is to reduce churn, improve ARPU and add new revenue streams.

Samples of the wireless services that are being considered include

- Consumer
  - MVNO
  - Fixed Wireless
  - IoT-Home
- Enterprise
  - IoT -Verticals
  - Private LTE
  - 5G front/backhaul
  - vRAN hosting
- Public Sector
  - Smart City
  - Public safety
  - Fixed wireless

The challenge facing MSOs is uncertainty. MSOs will have to launch many products quickly knowing that changing market conditions will churn the product portfolio. This coupled with limited wireless expertise will result in unexpected surprises in commercial readiness of the new products. New partnership models may also pose threats to revenue.

The “Agile” way is to launch massive “distributed experimentation” of wireless services. Distributed experimentation unshackles cross-functional teams to rapidly ideate, innovate and test the business and technology viability of many different products. The emphasis is to encourage risk taking and allow for failures. If an idea fails to meet predetermined technical or business viability objectives, then learn from the failure and find some other idea to innovate.

However, distributed experimentation by many teams needs to have good governance so that the efforts are mutually beneficial and learnings from all projects benefit the final commercial product. For this strategy to work, the MSOs need to settle on two very important points on the roadmap

- 1) The end point– Future Wireless Network Architecture
- 2) The start point – Foundational Wireless Network Architecture

By fixing the start and end points of the network architecture, the teams working on wireless services will have a good platform to start experimentation and clear direction to execute their projects and test for success.

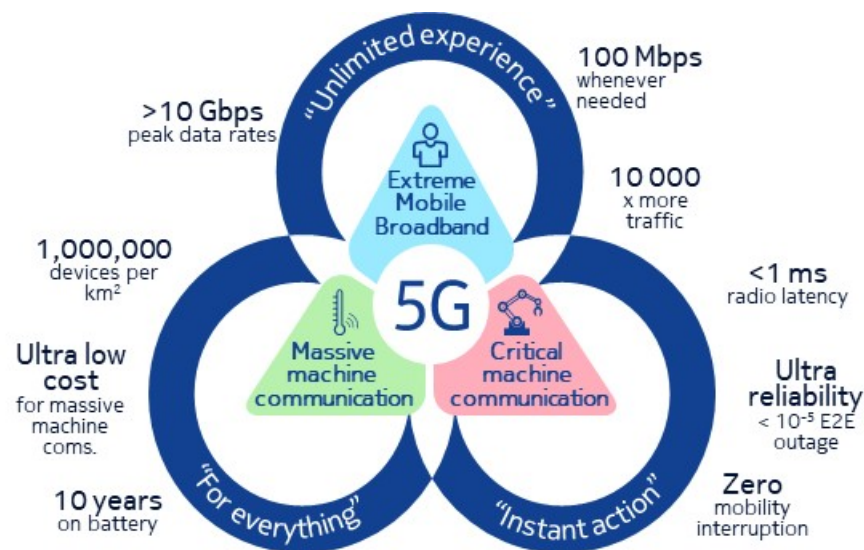
This paper demonstrates this concept by proposing the future MSO wireless architecture and the foundational architecture. Then taking a few sample services like Private LTE, Fixed wireless and Full MVNO, we demonstrate how the foundational architecture can grow to a service architecture. This paper also analyzes the CAPEX investments that would be needed for the foundational architecture and the extension of those to full service architecture.

## 2. Architecture principles of the future MSO wireless network

The new wireless use cases are typically categorized as below

- Enhanced Mobile Broadband (eMBB): requiring high data rates across a wide coverage area.
- Ultra-Reliable Low Latency Communications (URLLC): places a premium on latency and reliability performance SLAs and are interesting for mission critical communications, such as Industrial automation
- Massive Machine Type Communications (mMTC): support unprecedented number of connected devices, which may only send data sporadically, such as Internet of Things (IoT) use cases.

When we analyze the requirements from these use cases, we uncover a picture that is shown in Figure 1.



**Figure 1 - Requirements of the future wireless services**

As is obvious from Figure 1, the requirements of the future wireless services in the new world of 5G are diverse and will conflict with each other. This means that the traditional principles of network architecture need to be modified to support these services.

To make matters more complicated, there are considerations beyond just technology that also impact the network architecture.

The network architecture will need to be flexible to support experimentation. In addition, MSOs already have deployed networks which can serve as an asset to accelerate service deployment. If not managed properly, MSOs can end up with a complicated multi-layer network which exponentially increases operational costs.

In face of these challenges, we propose that MSOs build their new wireless network adhering to the following principles.


- Cost and risk reduction through **convergence** across access, transport, common NFVI and services

- Simplification and improved quality with service centric **automation**
- Scalability and low latency through edge cloud **virtualization**

These principles will reduce TCO, leverage existing assets and enable fast service deployment using network slicing. In the following sections, we will build the network architecture that adheres to these principles and meets the new wireless needs.

### 3. Future Wireless Architecture – Converged and Virtualized with massive automation

The proposed future wireless architecture, aligned with the concepts of convergence, virtualization and automation is shown in Figure 2.



**Figure 2 - Future wireless architecture**

The key points of this architecture are described below

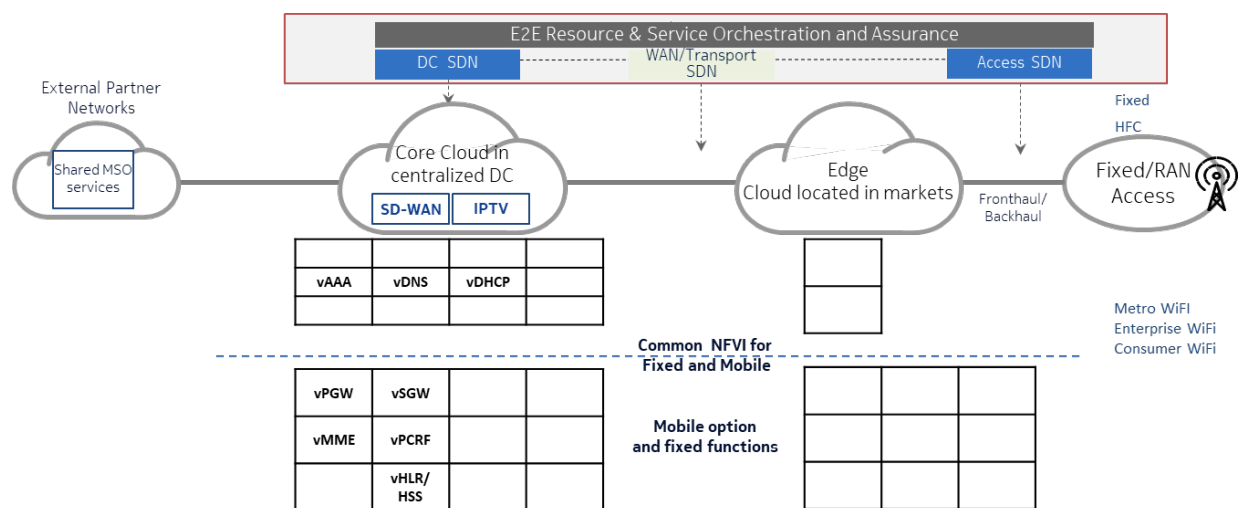
- 1) **Convergence:** The network functions that are above the dotted line in Figure 2 are used in both wireless and fixed networks and collapsing them reduces the TCO of the overall network.
- 2) **Distributed and Virtualized Edge Cloud:** MSOs have a unique advantage of assets like regional data centers and hubs that are deployed very close to where the services are consumed, which can be leveraged for the edge cloud. In addition to supporting the MSO needs, these edge cloud assets can also be monetized by offering them as a service to other operators (e.g.: BBU hoteling).
- 3) **Massive automation:** This architecture transitions from hardware centric physical assets to software defined networks which enables automations like Network slicing, engineering and deployment of services and zero touch close loop assurance.
- 4) **Interworking with partner networks:** This programmable network exposes control of functions via APIs to partner networks.
- 5) **Multi-wireless access:** It is expected that MSOs may need to deploy multiple different wireless technologies and networks including
  - a. WiFi: Enterprise, metro, hospitality, Consumer
  - b. RAN: 4G and future 5G Fixed wireless, Small cell offload, private LTE
  - c. IoT: LoRA, NB-IoT, LTE-M

It is important to note that the above architecture is the “end point” for the future network. It is not anticipated that MSOs will get to this end point in one step, rather we expect an evolution that will be driven by business needs.

## 4. Foundational Network Architecture- the starting point for the future

The MSOs need a minimal architecture that enables them to start the journey towards the future. The main considerations to select the network functions in this architecture is that they should be common across multiple services and support technical experimentation. The foundational architecture is not a complete service architecture.

This architecture is shown in Figure 3.



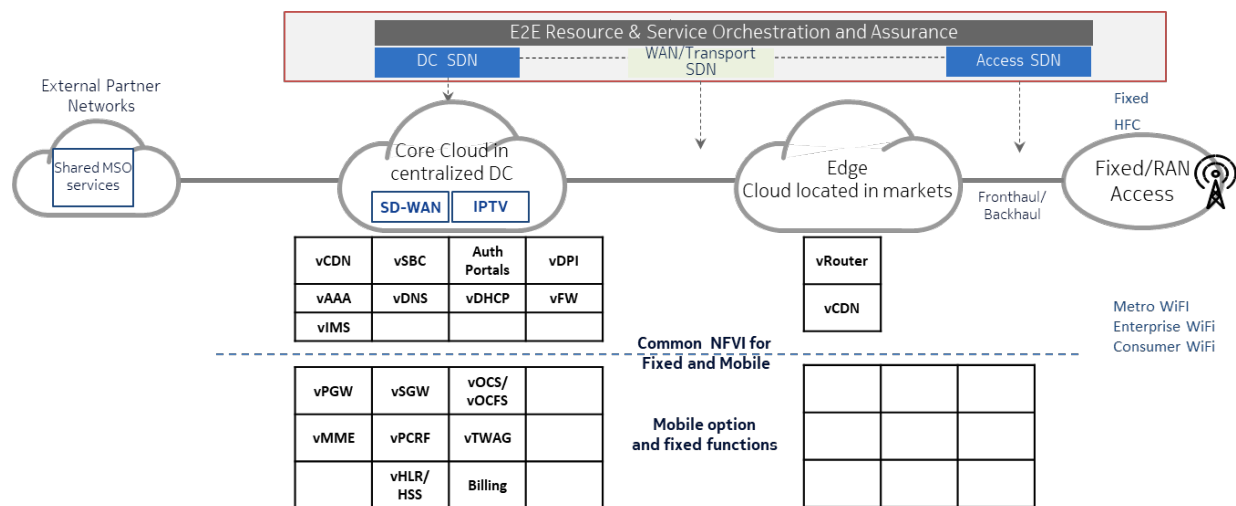
**Figure 3 - Foundational architecture for future wireless services**

The functions included in this foundational architecture include the Packet core for the wireless network (vMME, vPGW, vSGW, vPCRF), vHSS for control of the end user and vAAA for authentication function. These minimal functions show up in all service architectures and it is optimal to pull these together in the foundational network. The resulting platform can be used to start developing service architectures.

## 5. Full MVNO, Fixed Wireless and Private LTE architecture

In this section, we extend the foundational architecture to support Full MVNO, Fixed Wireless and Private LTE services.

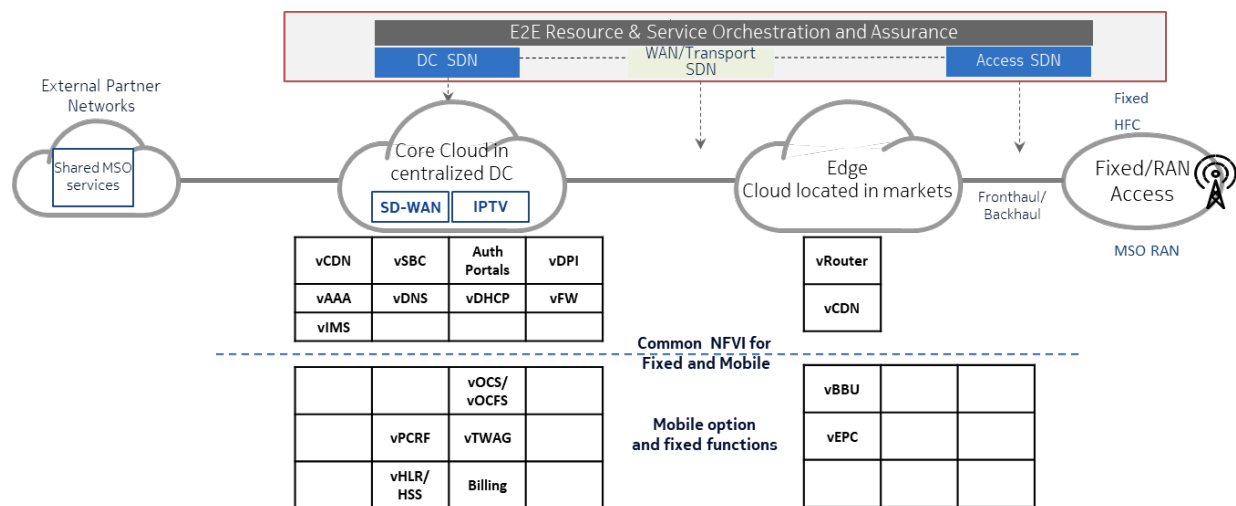
Figure 4 shows the service architecture that will support full MVNO services including core functionality to control user sessions and seamless offload.



**Figure 4 - Full MVNO service architecture**

We observe that the enhancements are on the Core, especially to support WiFi and Wireless interworking. This architecture uses the RAN from a partner MNO. In addition, the edge network has limited functionality, as there are no low latency services needed.

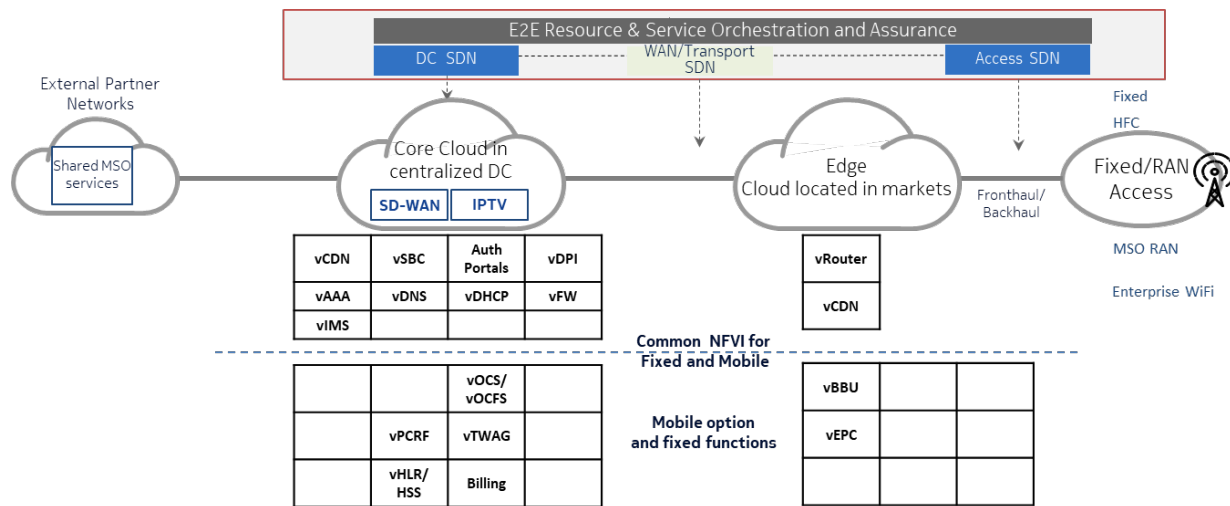
Figure 5 shows the service architecture for a fixed wireless service potentially over CBRS spectrum that can provide a Home Hotspot broadband solution. This architecture can also be used as a potential offload architecture for the MVNO service.



**Figure 5 - Fixed Wireless service architecture**

This architecture adds the MSO RAN components, which can be in the form of small cells, Fixed wireless CBRS or mmWave cell sites. Also, the vBBU and the vEPC network functions can be moved to the edge if there is a need for low latency services.

Private LTE network can be supported by an architecture shown in Figure 6.



**Figure 6 - Private LTE network architecture**


In this architecture, MSOs will need to be ready to deploy the edge cloud with vBBU and vEPC functions at the edge to support the services needed.

## 6. Investment comparisons

We observe that the foundational architecture is persistently needed in all service architectures. In some services a function like the vEPC may need to be relocated to the edge, which is simpler in the programmable virtualized architecture.

This proves the value of the foundational architecture to provide solutions to the challenges that we discussed earlier.

In addition, from an investment perspective, the foundational architecture serves as a low risk bet to get the MSO organization ready for business, technology and operations to support new wireless services.



**Figure 7 - CAPEX investment comparison for foundational v/s full architecture**

Figure 7 analyzes the CAPEX investments needed in the foundational architecture as compared to the service architectures for Full MVNO and Private LTE networks. We observe that an additional investment of approximately 20-30% will enable an accelerated launch of new wireless services.

## Conclusion

In this paper, we analyzed the imperative of a new approach towards wireless. Enabling rapid experimentation can be effectively utilized to satisfy both technical and business models. This approach is possible by deploying the foundational architecture before any of the service deployment decisions have been made. This low risk architecture needs only 20-30% of the investment and accelerates deployment of services such as Full MVNO, Private LTE and Fixed Wireless. This provides a cost-effective and yet flexible solution.

## 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
5G NSA	5G Non-Standalone
5G SA	5G Standalone
AMF	Access and Mobility Management Function
AUSF	Authentication Function
CBRS	Citizens Band Radio Service
CDN	Content Delivery Network
CPE	Customer Premises Equipment
IWK-SCES	Interworking Services Capabilities Exposure Server
LoRA	Long Range
LPWAN	Low Power Wide Area Network
MNO	Mobile Network Operator
MVNO	Mobile Virtual Network Operator
NRF	Network Repository Function
OCF	Online Charging Function
OFCF	Offline Charging Function
RPD	Remote PHY device
RMD	Remote MAC/PHY Device
SBC	Session Border Controller
SCS	Service Capabilities Server
SD-WAN	Software Defined Wide Area Network
SMF	Session Management Function
TWAG	Trusted Wireless Access Gateway
vAAA	Virtual Authentication, Authorization and Accounting



vBBU	Virtual Base band Unit
vEPC	Virtual Enhanced Packet Core
vHLR	Virtual Home Location Register
vHSS	Virtual Home Subscriber Server
vMME	Virtual Mobility Management Entity
vPCRF	Virtual Policy and Charging Rules Function
vPGW	Virtual Packet Gateway
vSGW	Virtual Serving Gateway

# **HFC Spectrum Expansion: Design and Component Impacts**

A Technical Paper prepared for SCTE•ISBE by

**Mark Vogel**

Director Network Architecture & Strategy

CommScope

6519 CommScope Rd

Catawba, NC 28609

828.241.6488

[mvogel@commscope.com](mailto:mvogel@commscope.com)

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
1. Options for Increasing Network Capacity.....	3
2. Component and Network Design Impacts.....	5
3. Upgrading Legacy Plants .....	8
4. Results .....	10
Conclusion .....	13
Abbreviations.....	14
Bibliography & References .....	14

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Node Splitting .....	4
Figure 2 - Data Capacity and CNR Requirements as a Function of Modulation Order.....	4
Figure 3 - Downstream Capacity as a Function of Modulation Order & Network Spectrum.....	5
Figure 4 - Trunk (PIII) and Drop Cable Attenuation Tables .....	6
Figure 5 - Gateway Architecture.....	7
Figure 6 - System Signal Level & Total Composite Power Profiles .....	8
Figure 7 - Model of Legacy 1002 MHz HFC Plant .....	9
Figure 8 - End-of-Line Signal Levels for 1.2 GHz Plant Extension .....	11
Figure 9 - End-of-Line Signal Levels for 1.8 GHz Plant Extension (1002 MHz Step-Down) .....	11
Figure 10 - End-of-Line Signal Levels for 1.8 GHz Plant Extension (Linear) .....	12
Figure 11 - End-of-Line Signal Levels for 3 GHz Plant Extension (800 MHz Step-Down) .....	12

# Introduction


As data demand continues to increase, cable operators are looking at ways to meet that demand by increasing the capacity of their Hybrid Fiber Coax (HFC) networks. Traditionally this has been done by pushing fiber deeper into the network, reducing amplifier cascades, reducing the number of customers in a service group and increasing modulation orders. However, the continuing development of the Converged Cable Access Platform (CCAP) and the introduction of DOCSIS® 3.1 has given operators additional methods to meet customer demands. At the top of this list is ability to support operation at extended frequencies up to 1.2 and even 1.8 GHz. There is also an effort currently underway at CableLabs® which is looking at the viability of increasing the upper frequency band to 3 GHz. Given that many of today's HFC networks have capacities ranging from 750 MHz up to 1002 MHz, how can an operator upgrade their networks without having to do a total rebuild or move to a Fiber to the Home (FTTH) network? This paper will discuss the impact of different network components and design options on the bandwidth expansion of an HFC network and present the results of a network model that applies these options to a 1002 MHz legacy plant design.

## 1. Options for Increasing Network Capacity

Cable operators have at least three approaches to increase the capacity of an HFC network. These can be implemented separately or together as they are not mutually exclusive but instead build upon each other. The first is probably the most familiar, and at its core, is modifying the network to reduce the number of subscribers associated with a radio frequency (RF) port on a CCAP platform. This can be achieved by implementing smaller node serving areas, splitting or segmenting a node, or just reducing the number of nodes on a CCAP port.

Figure 1 shows an example of node splitting where a 1000 households passed (HHP) node, that is fed by one CCAP port, is split into four 250 HHP nodes each served by a different CCAP port. This increases the data rate per HHP by a factor of four, which for a 1 GHz DOCSIS® 3.1 system operating at 256 QAM, provides ~6 Gb/s capacity to 250 HHP instead of 1000 HHP. Of course, this comes at the cost of additional CCAP ports, optical transmitters/receivers, nodes, and Dense Wavelength Division Multiplexing (DWDM) devices if spare fiber is not available.

Node splitting has limitations, as described in [Ulm/Maricevic, 2016], as service level data rates increase. This has traditionally been the approach to capacity expansion but going forward will not be as effective as the other options for increasing network capacity.



**Figure 1 - Node Splitting**

Increasing the modulation order is another method to increase capacity. DOCSIS ® 3.1 allows for a modulation order as high as 4096 QAM. Networks generally support 256 QAM which has a spectral efficiency of 8 bits/symbol. 4096 QAM modulation would raise this to 12 bits/symbol or a 50% increase in capacity per channel.

Of course, everything has its price and for modulation it is the received Carrier to Noise Ratio (CNR) at the cable modem as shown in the table in Figure 2 below (data rates are for a 258 to 1026 MHz downstream spectrum, and are the raw data rates which include framing/header overhead). As the modulation order increases so does the required CNR, which makes sense as there is a higher density of symbols in higher order modulations. This issue is somewhat mitigated as networks have evolved, because the number of amplifiers in cascade has been continually reduced which improves the resultant end of line CNR. One of the other benefits of DOCSIS® 3.1 is that it enables the use of multi-carrier Orthogonal Frequency Division Multiplexing (OFDM), where individual sub-carriers (within the channel) can each have a different modulation order. As a result the modulation of each sub-carrier can be tuned to accommodate the level of CNR at its operating frequency.

<i>Modulation</i>	<i>SG Capacity (Gb/s)</i>	<i>CNR (dB)</i>
256	6.1	27
512	6.8	30.5
1024	7.6	34
4096	9.1	41

**Figure 2 - Data Capacity and CNR Requirements as a Function of Modulation Order**

The third approach operators have for increasing capacity is bandwidth or spectrum expansion. By increasing the network operating spectrum from the current frequencies of 750, 860 or 1002 MHz to 1.2, 1.8 or 3.0, cable operators can add additional DOCSIS® channels and thereby increase the network data capacity. DOCSIS® 3.1 adds downstream capacity in blocks of 192 MHz, and so an increase in spectrum from 1 GHz to 1.2 GHz adds one 192 MHz block, to 1.8, adds another three blocks (four total), and to 3 GHz another six blocks (10 total).

The table in Figure 3 below shows the resulting network capacities again as a function of modulation order (assumes all DOCSIS® 3.1 channels). Utilizing modulation order and spectrum, an operator can move their 1 GHz network from a service group capacity of 6.1 GHz to 32 GHz. However, operating at

higher frequencies means greater attenuation in passive devices and therefore greater difficulty in meeting end of line signal level requirements. It will also mean that the network needs to be upgraded to accommodate the higher operating frequencies.

	<i>SG DS Capacity (Gb/s)</i>			
<i>Modulation</i>	<i>1 GHz</i>	<i>1.2 GHz</i>	<i>1.8 GHz</i>	<i>3 GHz</i>
256	6.1	7.6	12.2	21.3
512	6.8	8.6	13.7	23.9
1024	7.6	9.5	15.2	26.6
4096	9.1	11.4	18.2	31.9

**Figure 3 - Downstream Capacity as a Function of Modulation Order & Network Spectrum**

## 2. Component and Network Design Impacts

Of the three approaches to increasing capacity, expanding spectrum potentially has the greatest impact on the existing HFC networks since every element in the signal path must support the extended spectrum. It also will have the biggest impact on network capacity and support legacy systems while also utilizing the bit loading (i.e. ability to use different modulation on each sub-carrier) capabilities of DOCSIS® 3.1 to maximize capacity. This paper will therefore focus on that particular approach.

Inherently, operating at a higher frequency introduces additional attenuation in the signal path. A signal operating at 3 GHz will experience greater attenuation than one operating at 1 GHz due to the higher losses in the outside plant components. Hardline trunk and drop cable both have higher losses at higher frequencies as shown in the tables in Figure 4 below. The magnitude will be dependent on the cable size, but all will have a higher attenuation at higher frequencies.

For example, PIII 500 cable has a loss of 4.8 dB/100' of cable at 3GHz compared to a 2.5 dB loss at 1 GHz which is an increase of ~2.3 dB/100'. However, this works both ways. To overcome this additional loss, an operator can move to a larger cable size which will have a lower loss at higher frequencies compared to a smaller cable, and will also introduce a smaller negative tilt. Again, looking at the tables in Figure 4, an F50 cable has the same loss at 3 GHz as an F11 cable does at 1.2 GHz and an F6 cable at 550

MHz. In others words, by moving to an F11 cable (instead of an F6), an operator can gain 3.9 dB for a 100' drop @3 GHz, and ~7 dB by using an F50 drop.

Frequency (MHz)	Loss (dB/100')			Frequency (MHz)	Loss (dB/100')		
	PIII 500	PIII 625	PIII 875		F6	F11	F50
5	0.16	0.13	0.09	5	0.58	0.38	0.24
55	0.55	0.46	0.33	55	1.6	0.96	0.57
211	1.09	0.92	0.66	211	3.05	1.9	1.13
870	2.35	1.95	1.41	870	6.1	3.98	2.29
1002	2.54	2.11	1.53	1002	6.55	4.35	2.51
1218	2.84	2.36	1.72	1218	7.21	4.92	2.90
1800	3.55	2.96	2.18	1800	8.97	6.00	3.46
2000	3.77	3.15	2.32	2000	9.50	6.37	4.03
3000	4.79	4.02	3.00	3000	11.86	8.02	4.94
Dia Over Outer Cond (in)	0.500	0.625	0.875	Dia Over 1st Tape (in)	0.187	0.287	0.463

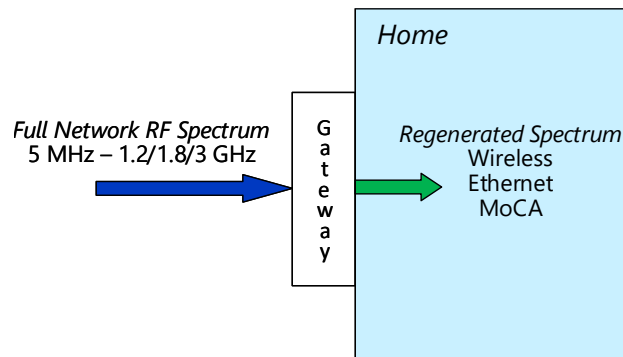
**Figure 4 - Trunk (PIII) and Drop Cable Attenuation Tables**

RF taps and splitters inside the home will also introduce a higher loss at higher frequencies. This could be 2-3 dB of additional loss at 3 GHz compared to operation at 1 GHz. Like cable, this can be overcome by using lower split levels in the case of a home splitter. Each reduction in split level, 8-way to 4-way, 4-way to 2-way, will reduce the loss by about 4 dB. For taps, the value of the tap can't be reduced since the signal level to the residence must be maintained, however, the number of taps in a tap run can be reduced. Each tap that is eliminated (actually moved from cable leg to another) not only eliminates the through attenuation of that tap, it also eliminates the cable attenuation for the cable that is connected to that tap.

It is not all bad news, however, with frequency extension. There are other changes to the network design that can be made to reduce the impact of additional plant loss. One is to allow a lower signal level at the end device. Typically, operators will set a device input level of 0 dBmV (analog signal level, digital signals -6 or more dBs below that). However, this level could be reduced at higher frequencies to accommodate the higher losses at those frequencies. Today's cable plants are much cleaner from an SNR/CNR perspective due to the reduction in amplifier cascades, so the impact of lowering the level on end of line SNR will not be so significant. Also since DOCSIS® 3.1 enables modulation on a per sub-carrier basis, worse case it means that the modulation order will need to be lowered at higher frequencies to a level that supports the reduced SNR at that frequency.

Another design option is to move to a Gateway architecture as shown in Figure 5. In this approach, the RF network is terminated at a single point somewhere near the entrance to the subscriber location. The device that performs this function is called a Gateway. The Gateway terminates all DOCSIS signals and sends that data throughout the residence via Wi-Fi, wired Ethernet, or over the existing coaxial network via MoCA® or some other in-home protocol. The benefit of the gateway architecture is that the loss

associated with the home coaxial network (cable + splitter) is eliminated, which can be 10+ dB of attenuation.




**Figure 5 - Gateway Architecture**

The final tool available is the improvement made in gain block technology which enables a higher RF output level from the optical node and sub-tending amplifiers. New high output GaN technology increases the Total Composite Power (TCP) of gain blocks relative to existing GaAs/GaN technology. As a result TCP levels can be raised as high as ~74 dBmV in a Node + 0 architecture, and to ~71 dBmV for Node + X architectures. This is reduced by about 3 dB (71./68 dBmV) with the circuit board implementation losses. For this paper, all levels will be referenced to “Virtual” analog signal levels, which will be 6 dB higher than that actual implemented digital levels (77/74 dBmV). Bottom line is that for Node + X architectures assumed here, this means the max TCP will be ~74 dBmV. As shown in Figure 6, this allows the Legacy system signal levels and tilt to be extended to 1.2 GHz. Unfortunately, these Legacy levels cannot be maintained above 1218 MHz as the TCP increases to ~85 dBmV @1.8 GHz, and ~107 dBmV @3 GHz. This is well beyond the capability of even the new GaN technology. To avoid this and to maintain legacy levels, it has been proposed that the gain is dropped a number of decibels and remains flat above a given legacy frequency (800 or 1002 MHz for example), as also shown in Figure 6. This compromise allows continued operation of the legacy systems and amp spacings (up to the selected frequency), while still providing maximum data capacity above that using DOCSIS® 3.1 bit loading capabilities. For 1.8 GHz, the Legacy level can be maintained to 1002 MHz, with a step down to 50



dBmV. At 3 GHz, to enable more power at higher frequencies, the Legacy frequency is reduced to 800 MHz, with a step down to 47 dBmV.



**Figure 6 - System Signal Level & Total Composite Power Profiles**

Another profile was considered @1.8 GHz that has a linear slope that is less than the Legacy levels, but which maintains the target TCP. Relative to the step profile, this profile places more power in the higher frequencies where loss/attenuation is the greatest. It does not maintain the Legacy levels at the lower frequencies, but better utilizes the excess power that is available when implementing a Gateway architecture (as will be discussed).

### 3. Upgrading Legacy Plants


So what is the impact to an existing plant when an operator moves to higher operating frequencies and utilizes one or more of these network design options? To determine the answer, a power budget model was built based on a 1002 MHz, legacy HFC plant design. The model does not look at trade offs that could optimize network capacity such as end-of-line signal level vs. supported modulation level. Instead, the model simply looks at the end-of-line signal level and the impact of various network implementations on maintaining that target level. It assumes that tap technology and the technology required to generate

higher node/amp signal levels will be available. Also, in the model, all signal level profiles are referenced to virtual analog signal levels (as shown in Figure 6), and actual digital signal levels would be 6 dB lower.

The model looks at a single tap run from the node or last amplifier with the following characteristics:

- Six 4 port taps in cascade w/ equal but maximized spacing
- 5/42 – 54/1002 MHz system
- Target level of 0 dBmV (analog) at subscriber device
- 1x4 split + 50' F6 cable inside the residence
- 100' F6 drop
- PIII 625 (QR 540) hardline cable
- All single carrier QAM channels

Figure 7 show the modelled HFC design and provides the resultant levels, spacing, and tap values. The design is an attempt to represent a "typical" HFC design, however actual designs will differ based on subscriber densities and the design criteria used by each operator. The goal is use a consistent set of criteria to see, in general, what provides the most benefit in reaching the goal of an extended spectrum network.



**Figure 7 - Model of Legacy 1002 MHz HFC Plant**

To do that, 8 scenarios were considered as shown in the bullet list below. All build off of Scenario #1 which upgrades the taps and node/amp modules to one supporting the particular extended frequency (1.2, 1.8 or 3 GHz) and which also implements the signal profiles for the various plant extension frequencies shown in Figure 6 (assumes a 208/258 MHz split). These are followed by several other single or cascaded scenarios that include elimination of the inside wire by utilizing a Gateway architecture,

increasing the drop cable size, increasing the trunk cable size of the first section of the tap run, and eliminating a tap in a tap run. These scenarios were run for all three targeted operating bandwidths.

- Scenario #1: Upgrade taps and node/amp modules to support a 1.2, 1.8 or 3 GHz operating bandwidth and the associated signal level profile (Figure 6)
  - Maintain SC QAM levels up to selected frequency
  - DOCSIS 3.1 above that frequency
- Scenario #2: Scenario #1 + eliminate inside wire and implement Gateway architecture
- Scenario #3: #2 + upgrade F6 to F11 drop
- Scenario #4: #2 + upgrade F6 to F50 drop
- Scenario #5: #2 + upgrade 1<sup>st</sup> hardline section to PIII 875 (QR 860)
- Scenario #6: #5 + upgrade F6 to F11 drop
- Scenario #7: #5 + upgrade F6 to F50 drop
- Scenario #8: #5 + eliminate last tap in tap run

For all scenarios, the end-of-line signal level @800/1002 MHz was also calculated to monitor the impact of network changes to the Legacy levels.

## 4. Results

The results for a 1.2 GHz extension are shown in Figure 8, a 1.8 GHz extension Figures 9A/B and a 3 GHz plant extension Figure 10. The tables in these figures show the tap values selected and the end of line device input levels for each of the scenarios at both the legacy frequency of 800 MHz (not shown for 1.2 GHz extension), 1002 MHz, and at each of the targeted extension bandwidths. A green highlight

means a given tap in a scenario met the targeted end-of-line level, a yellow highlight means the end-of-line level was below the target by up to 0.5 dB, and a red highlight means the signal level was below that.

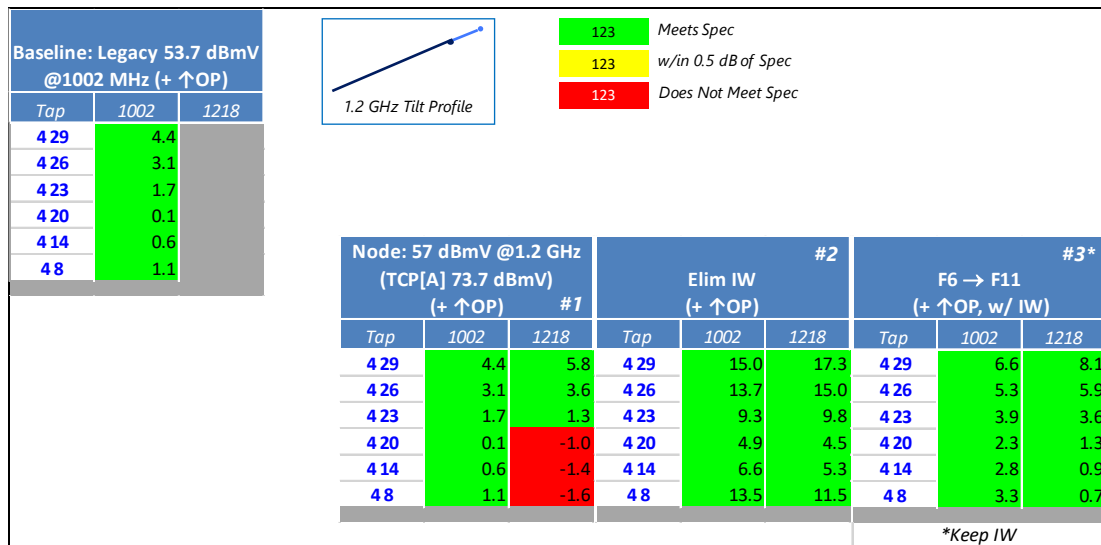


Figure 8 - End-of-Line Signal Levels for 1.2 GHz Plant Extension

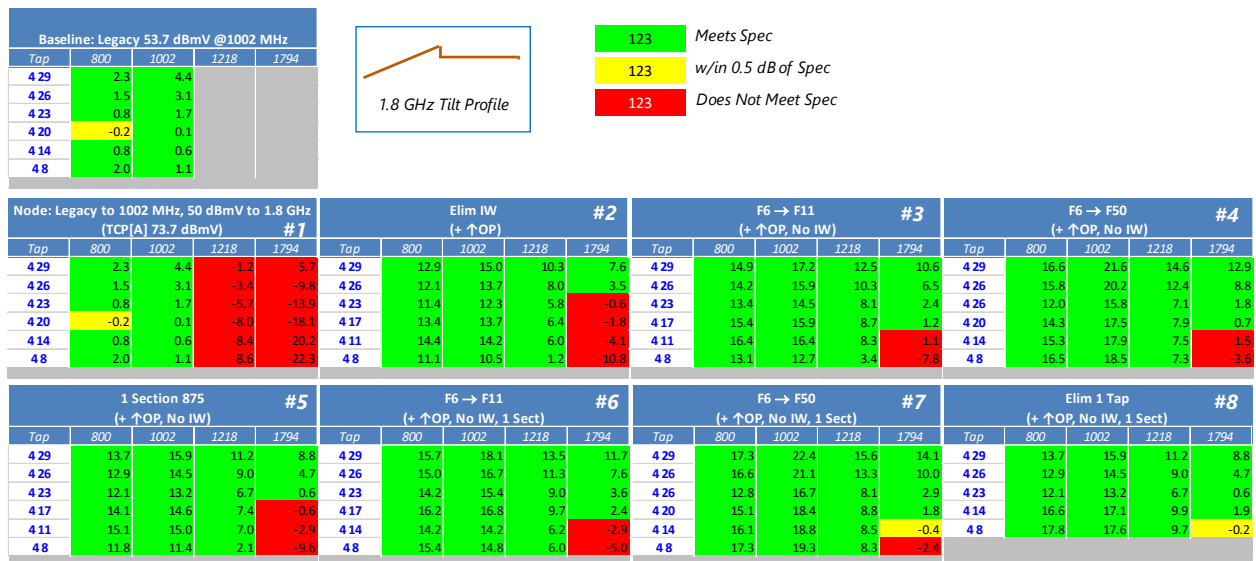


Figure 9 - End-of-Line Signal Levels for 1.8 GHz Plant Extension (1002 MHz Step-Down)



Figure 10 - End-of-Line Signal Levels for 1.8 GHz Plant Extension (Linear)

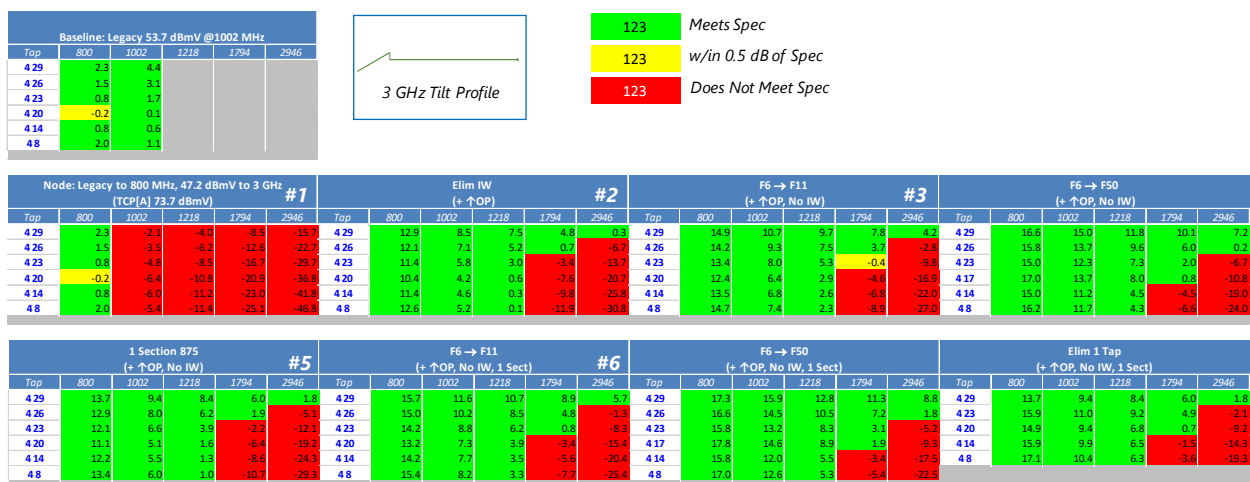


Figure 11 - End-of-Line Signal Levels for 3 GHz Plant Extension (800 MHz Step-Down)

As can be seen for the 1.2 GHz extension shown in Figure 8, scenario #1 (which changes taps and node/amp modules and raises the TCP to the target level) is mostly sufficient to achieve the desired levels at 1.2 GHz. The last taps in the cascade are slightly below the target (within ~1.5 dB), but this might be acceptable with the bit loading capabilities of DOCSIS® 3.1. All tap EOL signal levels are met with scenario #2 (Gateway), but that is over kill for this a 1.2 GHz extension, and a more reasonable approach to achieve the target levels is to implement scenario #3 (upgrade F6 drop to F11). These results show that in general, upgrading to 1.2 GHz is a relatively straight forward drop-in upgrade, as traditionally has been done for previous bandwidth extensions.

As expected, implementing scenario #2 in the 1.8 GHz and 3 GHz plant extensions is not sufficient to overcome the high losses at the higher operating frequency, but with the step down profiles, the Legacy levels are maintained. Moving to a gateway architecture and eliminating the home wiring (#3) enables

the 1.8 GHz upgrade to almost meet the EOL criteria (except for a few taps @ 1.8 GHz) and significantly improves the 3 GHz. Based on some of the CableLabs Extended Spectrum Effort initial feedback, many vendors and operators are assuming a Gateway architecture will be required for both a 1.8 and 3 GHz upgrade. One interesting aspect of this scenario is that it shows there will be an excess of power at the Legacy frequencies, particularly with the step down profile. For the 1.8 GHz scenario, @1002 MHz the max EOL Legacy signal level is 15 dBmV, where with the linear profile it is 9 dBmV.

Focusing on the 1.8 GHz plant extension, while scenario #3 doesn't quite meet the EOL target (w/in 11 dB for the step profile and 5 dB for the linear profile), it may be sufficient with the use of DOCSIS® 3.1 bit loading capabilities. Alternatively, upgrading the drop cable to F11 or F50 brings everything within the target EOL level for the linear profile, and within ~4 dB for the step profile. Implementing scenario #5 (upgrading the first section of trunk cable) has a slight impact, but scenarios 6, 7, or 8 are needed to have a significant impact. Again, the linear profile has better EOL performance than the step profile, and meets the EOL targets for scenarios #6-8, while the step down profile improves performance but only meets the EOL criteria with scenario #8. For the step profile, DOCSIS® 3.1 bit loading may enable sufficient performance for these scenarios.

Finally, looking at the 3 GHz extension, all the remaining upgrade scenarios (2-8) meet the EOL signal requirements in the legacy band, and in the 1.2 GHz band. In the 1.8 and 3 GHz band, however, only a couple of the taps meet the criteria. In the 1.8 GHz bands, EOL levels come within 4-12 dB of the target EOL criteria which may be sufficient with DOCSIS® 3.1 bit loading. However, in the 3 GHz band, the levels are only within 19-31 dB of the target EOL criteria which may be too low for DOCSIS® 3.1 cable modem receivers. Overall, the model shows that new technology or alternative architectures still need to be developed to fully utilize a 3 GHz frequency extension in today's HFC networks.

## Conclusion

Operators are looking for ways to expand the capacity of their HFC networks. Extended Spectrum DOCSIS or ESD, which is enabled by DOCSIS® 3.1, provides a platform for expanding capacity to 1.2 and 1.8 GHz and in the future 3 GHz. To determine how components and network designs impact the expansion of bandwidth in an HFC network, a power budget model was developed and various upgrade scenarios were run on a Legacy, 1002 MHz network. With the development of new high output GaN gain blocks, an upgrade to 1.2 GHz is achievable in much the same way the other drop-in bandwidth upgrades have been done in the past. A gain block improvement itself is not sufficient for bandwidth extensions to 1.8 or 3 GHz which minimally require an upgrade to a Gateway architecture. Additionally, upgrading drop cables from F6 to F11 or F50 enables the 1.8 GHz plant extension to meet end of line target levels, particularly with a linear tilt profile. Even though this signal profile lowers the signal below Legacy levels as compared to a step down profile, the lower power levels are offset by the lower attenuation of the upgraded network elements and design. Not surprisingly, an upgrade to 3 GHz will be the most difficult and even with a Gateway architecture, and drop, trunk, and tap upgrades, the EOL signal levels at 3 GHz fall short of the target by 20-30 dB. This indicates new technologies or other architectures will be required to fully utilize a 3 GHz network.

Overall, as the model shows, utilizing larger cables (particularly drop cable) will help facilitate the move to extended operating frequencies in an HFC network. Operators can start preparing for a future bandwidth expansion by using an F11 or larger drop cable instead of an F6 cable as they install new or maintain their existing drop plant.

## Abbreviations

CCAP	Converged Cable Access Platform
CNR	Carrier to Noise Ratio
CPE	Customer Premises Equipment
dBmV	Decibels with respect to a millivolt
DOCSIS	Data over Coax System Interface Specification
DWDM	Dense Wavelength Division Multiplexing
FTTH	Fiber to the Home
GaAs	Gallium Arsenide
GaN	Gallium Nitride
GHz	Gigahertz
HFC	Hybrid Fiber Coax
HHP	Households Passed
MoCA®	Multimedia over Coax Alliance
MSO	Multiple Systems Operator
OFDM	Orthogonal Frequency Division Multiplexing
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
SC	Single Carrier
SCTE	Society of Cable Telecommunications Engineers
SG	Service Group
SNR	Signal to Noise Ratio
TCP	Total Composite Power

## Bibliography & References

Arris (2016) John Ulm and Zoran Maricevic *Future Directions for Fiber Deep HFC Deployment*. Retrieved from <https://www.arris.com/globalassets/resources/white-papers/scte-future-directions-for-fiber-deep-hfc-deployments.pdf>

# Operating Legacy Cable Modems in an FDX Environment

A Technical Paper prepared for SCTE•ISBE by

**Shaul Shulman**

Senior Architect, Connected Home Division  
Intel Corporation  
Ha-Shfela 18, Petach Tikva, Israel  
972-3-9207943  
shaul.shulman@intel.com



# Table of Contents

Title	Page Number
Table of Contents.....	2
1. Introduction.....	4
2. Legacy Cable Modem Use Cases.....	5
2.1. Coexistence Use Case.....	5
2.2. FDX-L Use Case.....	5
2.3. “Soft Split” Use Case.....	6
3. Characterization of the Interference (Aggressor).....	6
3.1. Power Spectral Density (PSD) Shape.....	7
3.2. Power of non-self ACI.....	7
3.3. Time Domain Characteristics.....	9
4. Characterization of the Impact on the Cable Modem (Victim).....	9
4.1. Overview of Legacy CM Analog Front End (AFE) Architectures.....	9
4.1.1. Receiver types.....	9
4.1.2. Single Channel Receiver Architecture.....	10
4.1.3. Wide Band Receiver Architectures.....	11
4.1.4. Full Band Receiver Architectures.....	11
4.2. Adjacent Channel Leakage Interference (ALI).....	12
5. Mitigation Techniques.....	13
5.1. AGC Settings Adjustment.....	13
5.1.1. Scenario A: DOCSIS 3.0 Basic.....	14
5.1.2. Scenario B: DOCSIS 3.0 Stringent.....	14
5.1.3. Scenario C: SCTE-40 Based.....	14
5.1.4. Scenario D: DOCSIS 3.1 Basic.....	15
5.1.5. Scenario E: DOCSIS 3.1 with FDX.....	15
5.1.6. Scenarios Summary.....	16
5.1.7. Expected performance with added AGC headroom to contain ACI (the simplified model example).....	16
5.2. Filtering.....	18
5.3. Reducing US Tx power.....	18
5.4. Modulation and Coding Scheme (MCS) Downgrading.....	18
5.5. Partial FDX Band Deployment.....	19
6. Conclusion.....	19
7. Abbreviations.....	19
8. Bibliography & References.....	21

## List of Figures

Title	Page Number
Figure 1 - Frequency Plan on an FDX enabled Node (Source: [1]).....	5
Figure 2 - A Typical 2-port Tap Configuration (Source for Tap Schematic: Arris FFT*-*Q Series Datasheet).....	6
Figure 3 - Cable Modem Front End Band Separation.....	10
Figure 4 - Generalized and Simplified Diagram of a Single Channel Receiver.....	11
Figure 5 - Third Order Intermodulation Distortion.....	11
Figure 6 - Generalized and Simplified Diagram of a Wide Band Receiver.....	11
Figure 7 - Generalized and Simplified Diagram of Full Band Receiver.....	12

Figure 8 - Adjacent Channel Leakage Interference (ALI).....	12
Figure 9 - Simplified Model of Full Band Receiver .....	14
Figure 10 - Basic DOCSIS3.0 Input Signal Scenario.....	14
Figure 11 - Stringent DOCSIS3.0 Input Signal Scenario .....	14
Figure 12 - SCTE-40 based Input Signal Scenario .....	15
Figure 13 - DOCSIS 3.1 Input Signal Scenario.....	15
Figure 14 - DOCSIS 3.1 with FDX Interference Input Signal Scenario.....	15
Figure 15 - SNR Margin for 4KQAM for a CM Designed for Scenario D .....	17
Figure 16 - SNR Margin for a CM Designed for Scenario A .....	18

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Typical Drop Cable Attenuation (Source: CommScope 2275V WHRL RG6 Datasheet)	8
Table 2 - Different ACI to Received Signal Level Scenario .....	8
Table 3 - Receiver Types for CMs of Different DOCSIS Generations .....	9
Table 4 - Possible CM Benchmark Input Signal Scenarios .....	16
Table 5 - Reduction in FDX Interference Power due to Partial FDX Band Allocation.....	19

# 1. Introduction

The Full Duplex (FDX) capability of the DOCSIS® 4.0 specification allows significantly increasing the upstream capacity of the HFC network, without sacrificing the downstream capacity and without extending the upper band edge of the downstream band. This is achieved by means of echo cancellation on the remote PHY node, enabling it to transmit downstream signals and receive upstream signals simultaneously in the band that is assigned to downstream in a Legacy Frequency Division Duplex (FDD) partitioning scheme. The Legacy Cable Modems (CMs) that are connected to the same network will experience upstream transmissions from FDX capable modems in the frequency region intended only for downstream at the time when these CM were designed and manufactured.

This can potentially cause service disruption if not handled properly, thus raising several questions. Can Legacy CMs function properly at all in an FDX environment? Should all CMs be replaced when even a single FDX capable modem is connected to the node? What are the methods to mitigate the impact? Are there methods to guarantee robust operation of Legacy CMs, perhaps even at the expense of downgrading their throughput capabilities or other attributes? What can be done with and without firmware upgrade to Legacy CM? All these questions must be addressed and understood before starting FDX CM deployment. In this paper we suggest a framework to address these questions. We first describe the different use cases of Legacy CMs in an FDX plant. Then, we address the interference and its impact on the CM's receiver for different receiver architectures. Lastly, we suggest several methods that can be employed to mitigate, monitor, and control the impact on Legacy CMs.


The analysis proposed in this paper allows understanding the considerations behind Legacy CM capabilities with respect to dealing with FDX interference, with the purpose of helping MSOs to focus on the necessary parameters that need attention and to enable estimating the expected performance. Eventually, a careful lab characterization and certification is required for each specific Cable Modem or set-top box model to characterize its specific behavior under the FDX interference.

Note that while FDX CMs will create an environment of signals on the plant which on some parameters may go beyond what the Legacy CMs were designed for originally (primarily the upstream interference power), other signal conditions of an FDX-ready plant ( $N+0$  or  $N+1$ ) are by design better. Due to a small number or a lack of amplifiers on the node, the network SNR is expected to be much better than what these CM were designed to. Other impairments, such as nonlinear distortion products associated with analog optical to electrical conversion are not present either. These and other differences enable the necessary wiggle room for making adjustments that enable the Legacy CM to operate in an FDX plant.

The considerations presented in the paper also apply to a more general case where not all CMs in the plant adhere to the same frequency split, thus creating a situation where some CMs may transmit upstream in the frequency range designed to be for downstream for other CMs. Such scenarios occur when there is an “ultra-high split”, for example, when 10G CMs use all or part of the FDX band for a static upstream without actually employing FDX methods of operation.

## 2. Legacy Cable Modem Use Cases

An intended frequency plan for an FDX enabled node is depicted in Figure 1.



**Figure 1 - Frequency Plan on an FDX enabled Node (Source: [1])**

The use cases for deployment of Legacy CMs in an FDX enabled plant can be categorized into three types:

- Coexistence only
- Legacy DOCSIS 3.1 CMs that are able to receive in FDX region (namely, FDX-L)
- FDX capable modems used in a “soft split” FDD plant

We will address each type in more detail.

### 21. Coexistence Use Case

In this use case, DOCSIS 3.1, DOCSIS 3.0 or earlier CMs are expected to continue to operate, and perhaps share bonded channels with FDX enabled CM outside of the FDX frequency band. These CMs comply to test scenarios of a Legacy HFC plant.

### 22. FDX-L Use Case

The FDX-L CM is defined in DOCSIS 4.0 specifications as follows: A DOCSIS 3.1 CM with a software upgrade which can a) transmit in the 108 to 204 MHz Full Duplex upstream channels and receive in the 258 to 684 MHz Full Duplex downstream channels, in a high-split access network, or b) can receive in the 108 to 684 MHz Full Duplex downstream channels in a mid-split access network, with no access to Full Duplex upstream channels.

The FDX-L PHY receive performance requirements are not defined in the FDX specifications, but it is sensible to assume that these should not exceed the requirements for FDX CM.

## 23. “Soft Split” Use Case

One possible use case for the FDX technology is the elimination of the guard band between the upstream and downstream within the FDX band by means of echo cancellation. In such a scenario, the operator will be able to set the split frequency by setting the Resource Block Allocation (RBA) in a way that will effectively split between the upstream band and downstream band to any desired ratio with the granularity supported by the FDX frequency plan.

The signal conditions the Legacy CM has to deal with is identical to the coexistence and FDX-L use cases, with the exception of the properties of the upstream signals:

- The upstream band is smaller than 684 MHz.
- The upstream band is fixed to the lower part of the FDX band.

## 3. Characterization of the Interference (Aggressor)


The coexistence issue of the Legacy CMs in an FDX plant originates from the strong upstream signals the FDX CM transmits in a frequency range intended for downstream reception in all Legacy CMs.

Before proceeding to understand the impact of these transmissions on the CM receiver, let’s establish the interference characteristics.

- Spectral shaping
- Power
- Time domain characteristics
- Topology dependence

The term used to describe such interference in the industry jargon is “non-self Adjacent Channel Interference (ACI)” or “external ACI”, as opposed to self-ACI which originates from the transmission of the same CM.

The dominant cause of interference is the DOCSIS 4.0 CM sharing a same tap with the “victim” CM. The contribution of CMs on other taps is significantly lower and can be ignored for the purpose of this analysis.



**Figure 2 - A Typical 2-port Tap Configuration (Source for Tap Schematic: Arris FFT\*-  
\*Q Series Datasheet)**

### 31. Power Spectral Density (PSD) Shape

The RG-6 drop cable loss down tilt is 4 dB between the edges of FDX band for 100 ft, while the upward tilt of the “reference PSD” (the specification required tilt of the transmission) of the CM is 10 dB. Consequently, the PSD of the non-self ACI is expected to have a 6 dB up tilt over the FDX frequency band for 50 ft drop and 2 dB up tilt for 100 ft drop, assuming a flat response of the coupler isolation in this band.

### 32 Power of non-self ACI

The power of non-self ACI depends on the following factors:

- The transmit power of the aggressor CM
- The coupling power ratio from the aggressor to the victim

The power of the aggressor CM depends on its ranging, i.e. on the link budget from a particular CM to the RPHY, and the desired receive level of the upstream signal at the input to the receiver of the RPHY.

Ideally, the CMTS would like the CM to transmit the highest power so that its upstream receiver signal level is optimal and would like the CM to be able to receive the signal at the lowest power possible so that the downstream transmission power could be the lowest. The same is true from the CM perspective, only with opposite considerations. When analyzing the relative level of the non-self ACI power to the received signal power, it is possible to take the following approaches:

- Assume a typical network model and “sensible” CMTS behavior for Tx and Rx levels.
- Use the minimum and maximum levels of Tx and Rx from the DOCSIS 4.0 specification.
- Use the non-self ACI levels “as is” from the Downstream BER test requirements of the DOCSIS 4.0 specification.

The latter describes an aggressor level at the victim CM of 4 dBmV/6 MHz at 108 MHz and a value of 10 dBmV/6 MHz at 684 MHz, while the downstream receiver values for the desired and rest of the channels is 0 dBmV. This means a non-self ACI level of 4 to 10 dB over the FDX frequency range.

A more stringent assumption is to use the maximum Tx of the DOCSIS 4.0 specification and the minimal Rx levels of the DOCSIS 3.1 specification independently to compute the interference. Assuming the CM to CM coupling of 30 dB, that would mean a non-self ACI level of  $49 \text{ dBmV} - 30 \text{ dB} - (-6 \text{ dBmV}) = 25 \text{ dB}$  (!) at the top edge of the FDX band. However, a more consistent assumption would be to use the receive levels of the DOCSIS 4.0 specification for outside of the FDX band. The ACI level then becomes  $49 \text{ dBmV} - 30 \text{ dB} - (+1 \text{ dBmV}) = 18 \text{ dB}$ .

The CM to CM coupling depends on the tap’s port-to-port isolation and the drop cable loss, which depends on its quality and length. The port-to-port isolation is guaranteed by tap vendors to be better than 20 dB, and is usually at least several dBs better over the FDX frequency range. Measurements of typical taps show better than 30 dB isolation over most of the range. The RG-6 drop cable loss is about 6 dB for 100ft. This contributes 12 dB to the isolation for 100ft drop plants and 6 dB for 50ft.

**Table 1 - Typical Drop Cable Attenuation (Source: CommScope 2275V WHRL RG6 Datasheet)**

Frequency	Attenuation (dB/100 ft.)
100 MHz	2.01
200 MHz	2.86
400 MHz	4.23
700 MHz	5.96
900 MHz	6.96

So, while the datasheet-based CM to CM coupling on the same tap can be as bad as 26 dB for 50ft drop, the typical case is more than 40 dB for 100ft drop. Table 2 shows different ACI level scenarios.

**Table 2 - Different ACI to Received Signal Level Scenario**

Scenario	CM to CM Coupling [dB]	Interference Level (Top Edge of the FDX Band) Relative to +1dBmV/6MHz Input Level, Assuming Maximum CM Tx Power, [dB]	Interference Level (Top Edge of the FDX Band) Relative to +6 dBmV/6 MHz Input Level, Assuming Maximum CM Tx Power, [dB]
50ft drop, datasheet worst case scenario	26	22	17
100ft drop, datasheet worst case scenario	32	16	11
50ft drop, assuming 30 dB port to port isolation	>36	12	7
100ft drop, assuming 30 dB port to port isolation	>42	6	1
100ft drop, assuming 35 dB port to port isolation	>47	1	-4

Given all the considerations above, it is possible to conclude that for most of the cases the ACI interference will be close to the input signal level, and that assuming 10 dBc of ACI as worst case (as is assumed in the DOCSIS 4.0 BER test) is a reasonable working assumption.

### 33. Time Domain Characteristics

The time domain characteristics of the ACI interference follow the transmission pattern of the DOCSIS 4.0 upstream CM. Since there could be only one or a few dominant aggressors for a given Legacy CM (its tap “neighbors”), it is likely that most of the time the frequency/time resource of the interference will not be fully utilized. This greatly relieves the interference impact on the victim.

However, to guarantee robust service of the Legacy, it is recommended to assume an extreme case where the aggressor is utilizing the maximum of its power and frequency resource, and also behaves in a bursty fashion.

## 4. Characterization of the Impact on the Cable Modem (Victim)

### 4.1. Overview of Legacy CM Analog Front End (AFE) Architectures

The severity and mechanism of the impact on a Legacy CM depends on its analog front-end architecture. In the past 15 years, the architectures used in CMs have gone through significant changes, primarily to accommodate for the growing number of bonded channels. To better understand the potential coexistence issues of Legacy CMs with an FDX interference, it is beneficial to understand the architectures of different generations of CMs.

#### 4.1.1. Receiver types

All DOCSIS and set-top box receivers can be divided into the following three categories.

- Single channel receiver
- Wide band receiver
- Full band receiver

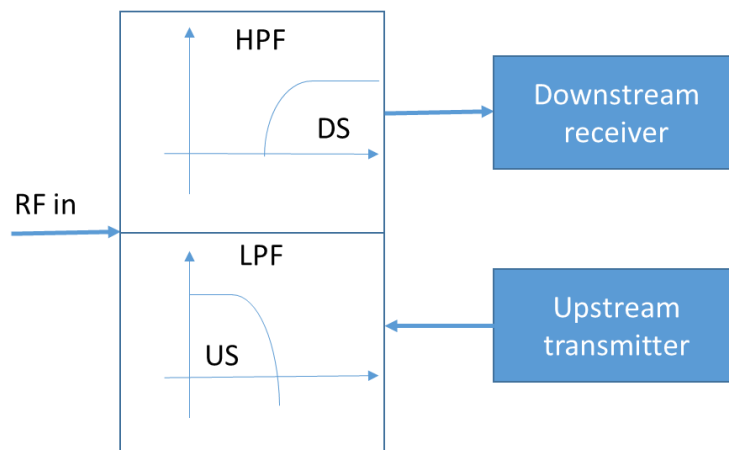
It is possible to identify the receiver type based on the DOCSIS standard support version and the number of supported downstream channels.

**Table 3 - Receiver Types for CMs of Different DOCSIS Generations**

DOCSIS Version	Number of Supported Downstream Channels	Receiver Type
DOCSIS 1.x – 2.0	1	Single channel receiver
DOCSIS 3.0	4-8	Wide band receiver
DOCSIS 3.0	16-32	Full band receiver
DOCSIS 3.1	2 OFDM, 32 SC-QAM	Full band receiver

All CMs use a diplex filter to separate upstream and downstream bands. The diplex filter consists of a low pass filter on the upstream transmission path, and a high pass filter on the downstream reception path. This ensures that the strong signal energy of the transmission in the upstream band and its harmonics in the downstream frequencies are filtered out by the combination of these two filters.

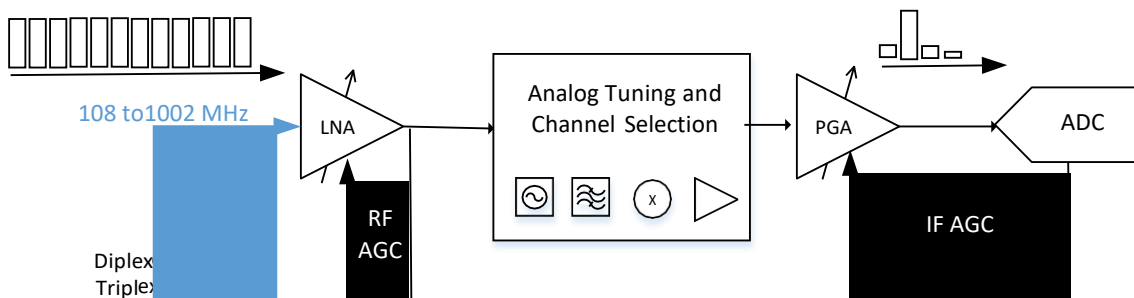





**Figure 3 - Cable Modem Front End Band Separation**

#### **4.1.2. Single Channel Receiver Architecture**

Single channel receivers are used for CMs compliant with the DOCSIS 2.0 specification or earlier. These receivers perform the tuning function and the channel selection function by analog means, using dual conversion, single conversion, low IF, or zero IF architectures. An oversimplified and generalized diagram of a single channel receiver is depicted in Figure 4. The RF gain is usually automatically controlled with no SW or FW intervention as the modem only “sees” the power in the single channel after it was “tuned” to IF or baseband and filtered out of the adjacent interference. The IF or baseband gain is controlled by the analog front end of the demodulator. The RF Automatic Gain Control (AGC) controls the gain of the low noise amplification stage of the receiver such that an optimal trade-off is made between the non-linear distortion and the thermal noise (noise figure). The IF AGC controls the signal level to optimize the input to the ADC. The total gain of the receiver, through both the RF and IF AGC, must present an almost continuous signal level to the demodulator to avoid performance degradation and packet loss. The AGC will not (and should not) be able to track significant and sudden changes in the total input power of the receiver. Such sudden changes in the total RF power at the input to the receiver, as in the case of FDX as an aggressor, may take the LNA out of its optimal point of trade-off between noise and interference, and cause degradation to the signal going out of the tuner to the demodulator. The level of impact and the frequencies affected may be somewhat unpredictable and dependent on the specific architecture of the tuner. The non-linear products that are created in the amplifier may fall on the frequency of the wanted channel. In this case, none of the filtering circuits further down the chain will be able to remove such an interference. The third order distortion, for example, produces not only harmonics of the signal at each threefold of the interferer frequency, but also intermodulation products of two or more signals. For example, intermodulation products of any pair of signals would fall on frequencies that are equal to the frequency of the interference plus and minus the gap between the interfering frequencies, as illustrated on Figure 5. The AGC at its converged state assuming a static downstream plant will damp the incoming interferer signal power to the level where those harmonics produce just as much distortion as designed by the CM vendor to pass the benchmark signal conditions during certification, testing, and typical plant operation. However, in case of a sudden rise of those interference signals beyond that level while the AGC is not yet able to converge, the level of interference will grow significantly. For third order distortions for example, the level of the distortion product grows by 3 dB for each 1dB of additional interference power.



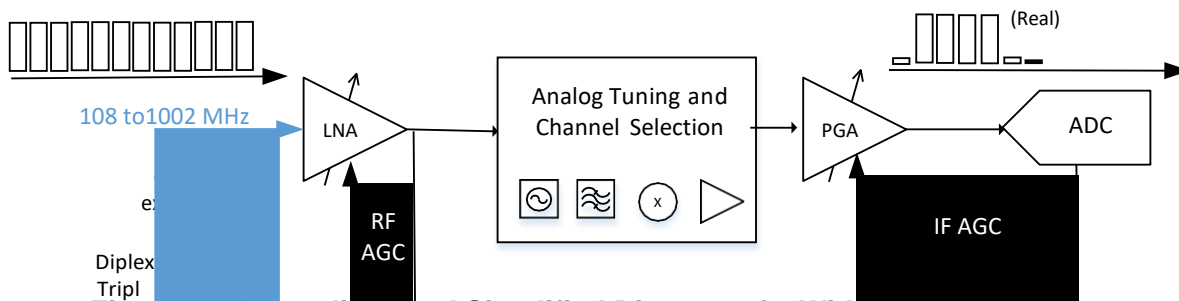
**Figure 4 - Generalized and Simplified Diagram of a Single Channel Receiver**



**Figure 5 - Third Order Intermodulation Distortion**

#### 4.1.3. Wide Band Receiver Architectures

Wide band receivers are used for CMs compliant with DOCSIS 3.0 specifications and are employed in CMs that support either 4 or 8 bonded downstream channels. The DOCSIS 3.0 specification requires the CM to be able to receive at least 4 channels simultaneously, with a total band “captured” of at least 64 MHz. In practice, some of the wide band receivers on the market are for dual 32 MHz (e.g. [2]), 100 MHz (e.g. [8]), and dual 96 MHz band widths (e.g. [9]). The architecture of the wide band receiver is similar in concept to the one of the single channel receiver, and the considerations mentioned in 4.1.2 apply here as well. The difference is in the width of the band passed to the demodulator. Also, these receivers support the 1002 MHz band edge. An oversimplified and generalized diagram of a wide band receiver is depicted in Figure 6.

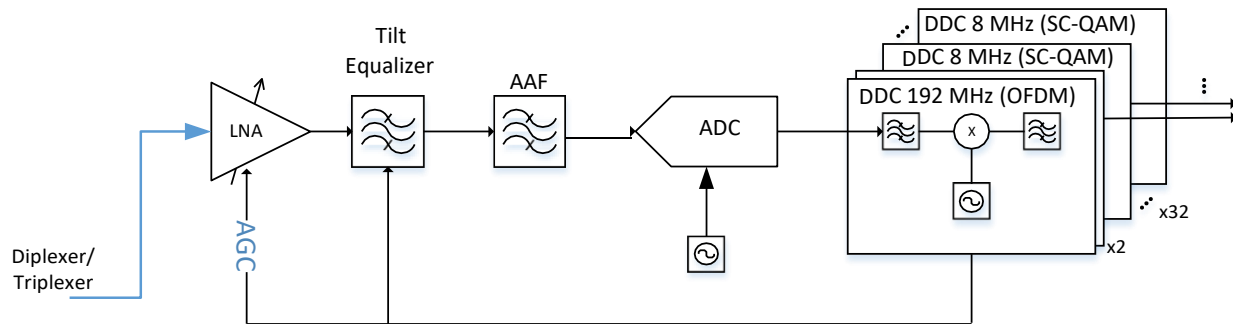


**Figure 6 - Generalized and Simplified Diagram of a Wide Band Receiver**

#### 4.1.4. Full Band Receiver Architectures

Full band receivers were introduced for a second generation of DOCSIS 3.0 CMs which support a number of bonded downstream channels of 16 or higher. The major difference of this tuner architecture is that all signal selection and down conversion processing is done in a digital domain. Such a receiver requires a high bandwidth and a high dynamic range ADC. There is no analog IF AGC but there is an RF AGC, which can now be controlled by the modem (using digital power

meter). The combined RF and digital AGC must ensure there is almost continuous signal channel power presented to the demodulator. It must adapt to changing signal conditions in the plant. The biggest challenge in designing such receivers is to ensure that the dynamic range of the ADC is high enough to enable sufficient SNR when the desired channel is surrounded by strong and numerous adjacent channels. The ADCs employed in these receivers stretch the bounds of feasibility for an integrated IC, and the performance of those is state of the art performance compared to publications (see [4] and the included references). To optimize the ADC dynamic range, the purpose of AGC scheme of the receiver is to bring the signal power to such level that would maximize the range of the ADC, while reducing clipping events to a level low enough statistically to not cause degradation. The downside of such headroom optimization is that any sudden total power change that the AGC is not tracking will cause ADC clipping. One mitigation to this problem is to add additional headroom, but that comes on the expense of the SNR of the ADC. This trade-off is further analyzed in section 5.1.




**Figure 7 - Generalized and Simplified Diagram of Full Band Receiver**

## 42 Adjacent Channel Leakage Interference (ALI)

Apart from the effects that are caused by the main lobe energy of the FDX signals, another contributor to interference is the leakage of the out-of-band emissions of the FDX signal to its neighbor. DOCSIS specifications guarantee a minimum 44 dB<sub>r</sub> rejection ratio between the Power Spectral Density (PSD) of the upstream signal and its out-of-band emissions. This is similar to the requirements from the downstream signal. DOCSIS 3.0 and earlier CMs are required to handle adjacent signals of at least 10 dB stronger and under stringent CNR conditions of 30 and 33 dB ([1]) which are not expected in N+0 plant. Therefore, it is likely that there should be no significant ALI impact on SC-QAM 256.

The case is different for FDX-L CMs. The DOCSIS 3.1 CMs are designed to deal with +3 dB adjacent. Any ACI level beyond that would mean some possible degradation to the signal reception capabilities of the FDX-L CM. This is generally true for all FDX CMs, not just for Legacy CMs.



**Figure 8 - Adjacent Channel Leakage Interference (ALI)**

## 5. Mitigation Techniques

It is possible to mitigate the impact of the interference by utilizing one or a combination of the following methods.

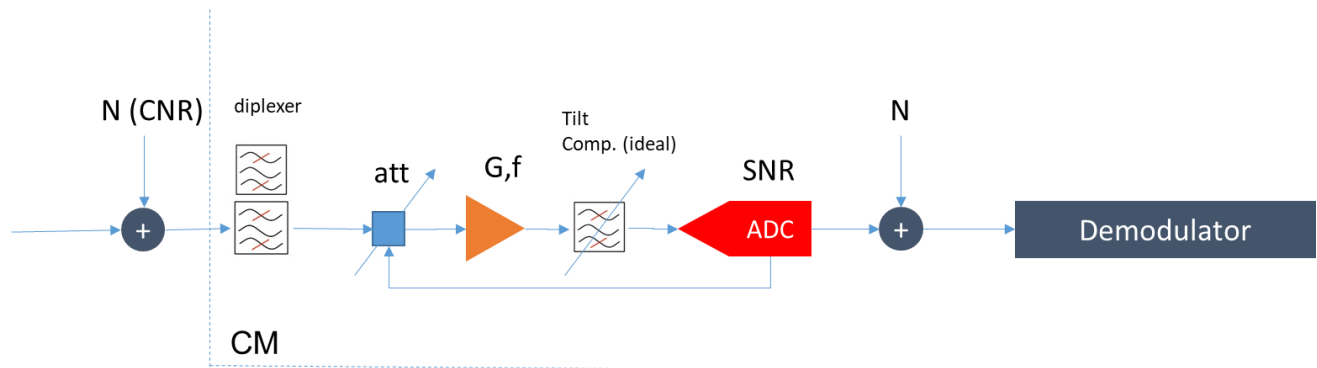
- Tap port to port isolation improvement
- External filtering
- CM AGC settings adjustment through firmware upgrade
- FDX upstream TX power control
- Modulation and Coding Scheme (MCS) downgrading
- Partial FDX band utilization

### 5.1. AGC Settings Adjustment

The essence of this method is to change the AGC setting (through CM firmware upgrade) so that sufficient dynamic range headroom is present to contain any possible sudden energy change which could overflow the full scale range of the ADC and cause ADC clipping. Such clipping, as described previously, could have a disruptive effect on CM performance. However, adding additional headroom is equivalent to raising the ADC noise floor, and will cause degradation to the SNR seen by the demodulator. A CM that was originally designed to meet high dynamic range conditions would be able to maintain the required SNR despite the additional headroom. The following is an intentionally oversimplified analysis to illustrate this point and draw equivalence between the known use cases for which CMs are designed and the usage of increased headroom for the purpose of containing the FDX non-self ACI bursts.

Consider the following simplified model of a full band receiver depicted in Figure 9. It consists of a constant gain LNA and a linear scale (in dB) attenuator, an ADC, and a demodulator. We fix the parameters so that the SNR at the slicer meets the requirements for a given signal condition scenario. We then vary the external parameters, namely the input CNR and the BO, to observe which other scenario such a CM would be able to withstand and what would be the SNR degradation, if any. The AGC is assumed to keep the total power constant at the level that maximizes the dynamic range of the ADC. The fixed parameters include:

- LNA noise figure at max gain
- PAPR headroom
- ADC SNR
- ADC jitter
- Demodulator “implementation loss”
- Non-linear distortion




**Figure 9 - Simplified Model of Full Band Receiver**

Let's consider several input signal scenarios.

#### **5.1.1. Scenario A: DOCSIS 3.0 Basic**


In this scenario, the plant is full of signals with two adjacent channels of +10 dBc.



**Figure 10 - Basic DOCSIS3.0 Input Signal Scenario**

#### **5.1.2. Scenario B: DOCSIS 3.0 Stringent**


In this scenario, the entire plant is assumed to be +10 dBc.



**Figure 11 - Stringent DOCSIS3.0 Input Signal Scenario**

#### **5.1.3. Scenario C: SCTE-40 Based**


This is a scenario that contains tilted spectrum with analog channels in the first half of the band and digital channels in the higher frequencies, two adjacents of 15 dB.



**Figure 12 - SCTE-40 based Input Signal Scenario**

#### **5.1.4. Scenario D: DOCSIS 3.1 Basic**


This scenario is assumed for BER performance tests of DOCSIS 3.1 CM



**Figure 13 - DOCSIS 3.1 Input Signal Scenario**

#### **5.1.5. Scenario E: DOCSIS 3.1 with FDX**

This is the assumed scenario in case of FDX transmission interference. The received signal is located anywhere in the yellow band.



**Figure 14 - DOCSIS 3.1 with FDX Interference Input Signal Scenario**

### 5.1.6. Scenarios Summary

Table 4 summarizes the scenarios together with the back-off of the desired signal (defined as the ratio between the 6 MHz band power of the desired channel and the total signal power “seen” by the ADC) and the external CNR.


**Table 4 - Possible CM Benchmark Input Signal Scenarios**

Scenario	Name	MCS	Back Off/ 6MHz	CNR	“Desired” Channel Power
A	DOCSIS 3.0 Basic	256QAM- J83B	22	30	-6
B	DOCSIS 3.0 Stringent	256QAM- J83B	32	30	-6
C	SCTE-40 Based	256QAM- J83B	33	33	-12
D	DOCSIS 3.1 Basic	4K- 8/9LDPC	26	41	-6
E	DOCSIS 3.1 with FDX	4K- 8/9LDPC	30	41	-6

### 5.1.7. Expected performance with added AGC headroom to contain ACI (the simplified model example)

Figure 15 shows the result of the simplified model in Figure 9 assuming the CM was designed for the basic DOCSIS3.1 signal conditions. It shows that for CNR of 38 and higher, and back-off of 30 dB,


the SNR is good enough for 2KQAM and about 2 dB short for 4KQAM.



**Figure 15 - SNR Margin for 4KQAM for a CM Designed for Scenario D**

For the SC-QAM modulations, it is clear that a CM designed for scenario B and C has the additional dynamic range required, as the back-off of those scenarios exceed the assumed FDX ACI. A modem designed to only satisfy the basic scenario A, does not have this additional dynamic range, as shown in Figure 16. However, it has enough dynamic range for 64 QAM.





**Figure 16 - SNR Margin for a CM Designed for Scenario A**

## 52 Filtering

A fool proof method to prevent any disturbance of a Legacy CM reception would be placing a physical filter between the tap port and the CM. This would be a band-stop type of filter suppressing the FDX band. Placing it in the premises closest to the CM rather than on the tap would be better in terms of return loss in this frequency region, due to the additional attenuation of the drop cable.

Alternatively, adding an attenuator produces the same effect. In an FDX plant, there should be enough margin for SC-QAM reception to allow reduction in the incoming signal power. The deployment of the filter can be limited to the tap that has the FDX CM on it, or even just the port that shares the final splitter stage with the FDX CM.

## 53. Reducing US Tx power

Another straightforward method of reducing the level of interference is reducing the Tx power of an FDX CM. It is possible to use “smart” and selective reduction of power only for those FDX CMs that actually interfere with Legacy CMs. Allocating these CM is possible by utilizing Proactive Network Maintenance (PNM) tools available to Legacy CMs as required by the DOCSIS 3.1 specification, and are common among DOCSIS 3.0 CMs as well.

## 54. Modulation and Coding Scheme (MCS) Downgrading

Together with reducing the Tx power mentioned in the previous section, MCS downgrading allows an agile compromise between the upstream and downstream throughput of the aggressor and victim CMs, respectively. DOCSIS 3.0 and earlier modulations allow dropping to 64 QAM from 256 QAM, while DOCSIS 3.1 allows practically continuous throughput reduction based on available SNRs. This is achieved using the bit loading as described in [7].

## 5.5. Partial FDX Band Deployment

While this is not an intentional mitigation method, it is worth mentioning that allocations smaller than the full FDX band reduce the level of potential interference. This is because according to the DOCSIS 4.0 specification the maximum transmit power is defined as a power spectral density (defined in the DOCSIS 4.0 specification as “reference PSD”), so that the total maximum power is achieved when the full FDX spectrum is modulated. Additionally, the reference PSD is up tilted so that higher frequency subbands have a higher contribution to the total TX power. The reduction in total power for different FDX spectrum allocation is depicted in Table 5.

**Table 5 - Reduction in FDX Interference Power due to Partial FDX Band Allocation**

<b>FDX Spectrum Allocation Width</b>	<b>Maximum Interference Power within Partially Allocated Band versus Full FDX Band Power (Assuming 6 dB up tilt) [dB]</b>
96 MHz	-10.7
192 MHz	-7.1
288 MHz	-4.8
384 MHz	-3
576 MHz	0

## 6. Conclusion

Backward compatibility and coexistence with earlier generations has historically been one of the strong points of DOCSIS technology. This allowed operators maximizing the return on investment by enabling longer life for deployed CPE, easier migration to new technologies, and avoiding disruption of existing services throughout the upgrades. What is different about FDX is that it does not provide such guarantee of coexistence by design due to the new frequency plan being incompatible with the old one. Therefore, it is important to carefully examine the potential disruption that FDX introduction may exhibit. In this paper, we characterized the interference and its potential impacts on different types of CMs, and proposed mitigation techniques. We believe that given the proper attention and careful examination of CM capabilities, together with the application of some of the methods mentioned in this paper, the coexistence issues of FDX and Legacy CMs could be greatly minimized if not eliminated altogether.

## 7. Abbreviations

ACI	Adjacent Channel Interference
ALI	Adjacent Leakage Interference
AWGN	Additive White Gaussian Noise
BER	1) Bit Error Ratio; 2) Bit Error Rate
BW	Bandwidth
CCI	Co-channel Interference
CM	Cable Modem
CMTS	Cable Modem Termination System

CNR	Carrier To Noise Ratio
CSO	Composite Second Order
CTB	Composite Triple Beat
dB	Decibel
dBc	Decibel Carrier
dBmV	Decibel Millivolt
dBr	Decibel Reference
DOCSIS	Data-Over-Cable Service Interface Specifications
DRFI	DOCSIS Downstream Radio Frequency Interface Specification
EC	Echo Cancellation or Echo Canceller
FDD	Frequency Division Duplexing
FDX	Full Duplex or Full Duplex DOCSIS
FEC	Forward Error Correction
GHz	Gigahertz
HFC	Hybrid Fiber/Coax
Hz	Hertz
IG	Interference Group
OFDM	Orthogonal Frequency Division Multiplexing
PAPR	Peak-To-Average Power Ratio
PHY	Physical Layer
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SCTE	Society of Cable Telecommunications Engineers
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

## 8. Bibliography & References

- [1] Data-Over-Cable Service Interface Specifications, DOCSIS® 4.0, Physical Layer Specification, CM-SP-PHYv4.0-D01-190628
- [2] Gatta, et. al. “An Embedded 65 nm CMOS Baseband IQ 48 MHz–1 GHz Dual Tuner for DOCSIS 3.0”, *IEEE Communications Magazine*, April 201
- [3] Richard S. Prodan, “Full Duplex DOCSIS PHY Layer Design and Analysis for the Fiber Deep Architecture”, SCTE Expo 2017 Technical Forum.
- [4] Olivier Jamin. “Broadband direct RF digitization receivers. Networking and Internet Architecture” [cs.NI]. Télécom ParisTech, 2013. English.
- [5] Murmann B., "ADC Performance Survey 1997-2019," [Online]. Available: <http://www.stanford.edu/~murmman/adcsurvey.html>
- [6] FFT\*-\*Q Series Datasheet. CommScope.
- [7] Arambepola, et. al. “Configurable constellation mapping to control spectral efficiency versus signal-to-noise ratio” US10158451B2
- [8] CSR MT2170 datasheet, avialble on-line.
- [9] MaxLinear MxL261 datasheet, <https://www.maxlinear.com/product/connected-home/cable-broadband/cable-modem-front-ends/mxl261>

# Ultra Low-Cost Injection-locked FP Laser Source for Coherent Access Networks

A Technical Paper prepared for SCTE•ISBE by

**Zhensheng(Steve) Jia, Ph.D.**

Distinguished Technologist

CableLabs

858 Coal Creek Circle, Louisville CO, 80027

303.661.3364

s.jia@cablelabs.com

**L. Alberto Campos, Ph.D.**

Fellow

CableLabs

858 Coal Creek Circle, Louisville CO, 80027

303.661.3377

a.campos@cablelabs.com

**Mu Xu, Ph.D., CableLabs**

**Haipeng Zhang, Ph.D., CableLabs**

**Junwen Zhang, Ph.D., CableLabs**

**Chris Stengrim, CableLabs**

**Curtis Knittle, Ph.D., CableLabs**

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	4
1. Coherent Optics for Access Networks .....	4
2. Motivation of Injection-locked FP Source for Coherent Optics.....	5
3. Operation Principle of FP Laser Injection Locking.....	6
4. Design of Injection-locking Scheme.....	7
5. Experimental Verifications.....	14
6. Applications and Cost Analysis.....	17
Conclusions.....	18
Abbreviations.....	18
Bibliography & References .....	19

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Coherent Optics for Cable Access Applications.....	4
Figure 2 – FP laser cavity modes.....	5
Figure 3 – Optical injection locking setup: (a) transmission style; (b) reflection style .....	7
Figure 4 – Multi-beam interference model for an Fabry-Perot laser cavity.....	8
Figure 5 – Gain spectra of the Fabry-Perot cavity versus (a) reflectivity of the front facet, and (b) cavity length.....	9
Figure 6 – SMSR versus injection power at different wavelengths.....	10
Figure 7 – FP laser I-V curve (a) and output power vs. bias (b) .....	10
Figure 8 – Optical spectrum of the FP laser: (a) free running; (b) injection locked.....	11
Figure 9 – Injection locking process study: (a) locking map under different injection ratio and frequency detuning; (b) optical spectrums of the slave FP laser under different locking conditions .....	13
Figure 10 – Side mode suppression ratio under various frequency detuning.....	13
Figure 11 – Delayed self-heterodyne laser linewidth measurement setup.....	14
Figure 12 – Measured ECL and injection-locked FP-LD linewidth.....	15
Figure 13 – Measured low-cost DFB laser linewidth .....	15
Figure 14 – System diagram for BER measurement.....	16
Figure 15 – BER performance with and without applying COIL for (a) 32-Gbaud DP-QPSK, (b) 60.375-Gbaud DP-16QAM, and (c) 40-Gbaud DP-64QAM .....	16
Figure 16 – Application I: Multiple fibers to different destinations.....	17
Figure 17 – Application II: Comb source for DWDM systems over single fiber .....	18

# Introduction

The market for coherent optical links to reach between 10 km and 120 km is emerging in many application scenarios, such as router-to-router and point-to-point data center interconnect, mobile xhaul and cable aggregation applications. These market opportunities have catalyzed huge investment and development of new power and footprint optimized pluggable products in optical industry. In the cable environment, access networks have been undergoing significant technology and architecture changes driven by the ever-increasing residential data service tiers and an increasing number of services types being supported, such as business services and cellular connectivity. Digital fiber technologies and distributed access architecture (DAA) for fiber deep strategies offer an infrastructure foundation for cable operators to deliver the best service quality to the end users in the years ahead.

CableLabs® has recognized the need of coherent optics in the access network and has been working on point-to-point 100G and 200G coherent optics specifications. On June 29th, 2018, CableLabs publicly unveiled for the first time two new specifications: P2P Coherent Optics Architecture Specification and P2P Coherent Optics Physical Layer v1.0 Specification [1]. These two new specifications are the result of a focused effort by CableLabs, its members, and the manufacturer partners to develop Coherent Optics technology for the access network and to bring coherent optical technology to market quickly. On March 11, 2019, CableLabs announced another addition to its family of Point-to-Point Coherent Optics specifications: The P2P Coherent Optics Physical Layer v2.0 Specification [2]. This new specification defines interoperable P2P coherent optics links running at 200 Gbps (200G) on a single wavelength.

Low-cost coherent transceiver design is of great interest in bringing coherent optics to access networks because the existing commercially available coherent opto-electronic subsystems are associated with a high degree of complexity and cost for long haul and metro applications. Among different components of the coherent transceiver, active continuous-wave (CW) laser is of critical importance performing both transmitter source and receiver local oscillator (LO) for such low-cost coherent systems. The laser used in current coherent system is an external cavity laser, or ECL, which generates a narrower linewidth for coherent system needs. The typical linewidth of ECLs is ~50-500 kHz level. It uses a reflector that creates the cavity outside of the gain chip and allows the cavity to be longer than if it was confined to the gain chip alone. However, it is high cost and complicated, which is not attractive in access applications.

Fabry-Perot (FP) lasers, or FP laser diodes (FP-LD), on the other hand, is simple and ultra-low cost. Cable operators have deployed inexpensive FP lasers in the HFC upstream optical links for years. Unfortunately, FP lasers are not applicable for coherent optics in its current use. In this paper, we propose low-cost injection-locked FP laser source for coherent access applications. Through injection locking, the child laser closely adopts the optical frequency and linewidth characteristics of the parent laser which can be shared between multiple child lasers. As a result, the cost of the coherent optical transceiver can be significantly reduced. In this paper the injection locking technique is investigated by applying the external parent laser to multiple child FP-LDs. The static and dynamic characteristics are studied in detail, which include injection locking and detuning condition, linewidth reduction feature, output power maximization, and optimized design for coherent system. In addition, this paper discusses applications of injection locked FP-LDs and optical frequency combs in coherent access system and presents an experimental comparison of transmission performance with and without injection locking.

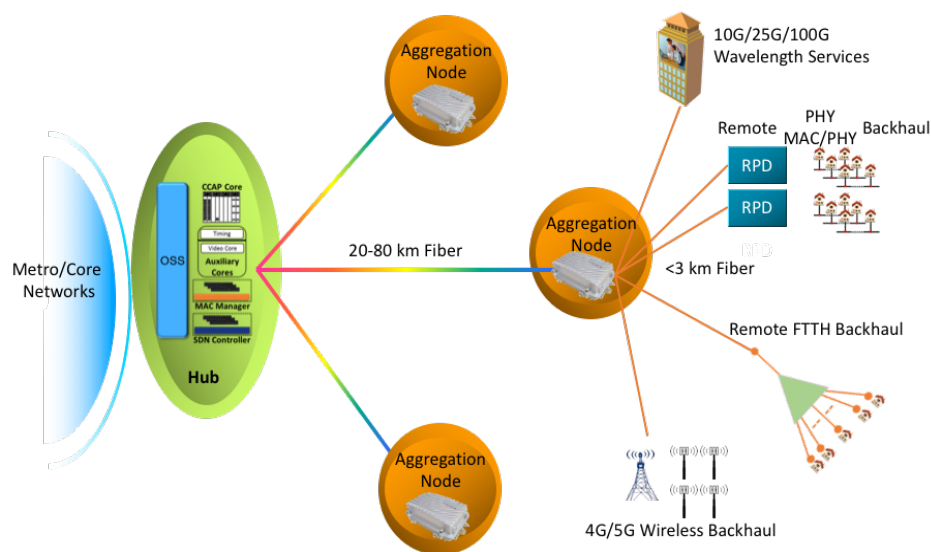
# Content

## 1. Coherent Optics for Access Networks

Leveraging further development of CMOS processing, reduction in design complexity, and cost reduction of opto-electro components, coherent solutions are moving from long haul and metro to access networks. Access networks have been steadily evolving in their capabilities, architectures, and the type of services they carry. Cable has been deploying fiber deeper and migrating towards DAAs that bring fiber much closer to subscribers. These networks have a transition towards much greater capacity per subscriber, more fiber in closer proximity to subscribers, and the simultaneous delivery of all types of services. Coherent optics is a very suitable technology to address the long-term evolution of access networks.

Coherent optics for access networks enable superior receiver sensitivity that allows extended power budget and high frequency selectivity for closely spaced dense/ultra-dense WDM channels without the need of narrow-band optical filters. Moreover, the multi-dimensional signal recovered by coherent detection provides additional benefits to compensate linear transmission impairments such as chromatic dispersion (CD) and polarization mode dispersion (PMD), and efficiently utilizes the spectral resource, benefiting future network upgrades through the use of multi-level advanced modulation formats.

Coherent optics technology can be leveraged in cable following two general approaches. First is when used as a means of multi-link aggregation, and the second is through direct end-to-end connectivity to the desired end-point as shown in Figure 1. Following capacity growth trends, it is obvious that initially the aggregation use cases are going to outnumber the direct end-to-end connectivity use cases. The aggregation use case supports any DAA, including Remote PHY, Remote MAC-PHY, and Remote optical line terminal (OLT) architectures and cellular aggregations.



**Figure 1 – Coherent Optics for Cable Access Applications**

Aggregation would likely take place at the optical terminal where fiber has already been deployed and from where multiple new deeper fiber links are being extended. This is the location of the original HFC node, which is becoming this evolved optical node. In this aggregation use case, a device host called the Optical Distribution Center (ODC) or Aggregation Node terminates the downstream P2P coherent optic link that originated at the Headend or Hub, and outputs multiple optical or electrical Ethernet interfaces




operating at lower data rates to connect devices that are either colocated with the ODC and/or exist deeper in the network. This aggregation or disaggregation function can be done by a router, an Ethernet switch, or a Muxponder, depending on the DOCSIS®/PON/business traffic demand, cost, scalability/flexibility/reliability, and other operational considerations. The distance between the Hub and fiber nodes or future Aggregation Nodes ranges from 20 to 80 km, and the distance from the Aggregation Node to each end point is less than 3 km. Today the hubs typically support tens of nodes support multiple (~60) Aggregation Nodes for different services [3].

## 2. Motivation of Injection-locked FP Source for Coherent Optics

Cost reduction is the key to successfully bring coherent optics into high-volume access environments. Therefore, simplified coherent transceivers have been intensively researched in recent years. Laser source is one of the high-cost items in coherent transceiver structure. The laser used in current coherent system is an external cavity laser, or ECL, which generates a narrower linewidth for coherent system needs. The typical linewidth of ECLs is ~50-500kHz in range. It has a reflector that creates the cavity outside of the gain chip and allows the cavity to be longer than if it was confined to the gain chip. Having the external cavity in addition to the gain medium semiconductor structure imposes conditions that leads to single frequency and very fine linewidth emission. However, ECL implementation is high cost and complicated for access applications.

FP-LD, on the other hand, is simple and ultra-low-cost light sources for low data rate short distance optical communication. Cable operators have deployed inexpensive FP lasers in the HFC upstream optical links before. Figure 2 shows the basic operation principle of a multimode FP Laser, which has a basic structure of resonant cavity in a gain media, typically III-V group semiconductor. The resonance cavity is formed by having a front reflection mirror surface, and a rear reflection mirror surface. For optimized output power efficiency, the rear reflection mirror usually has a higher reflectivity than the front mirror surface. A stable resonant mode is formed by meeting the following two conditions. First, an integer number of wavelengths is equal to the round-trip optical path of the laser cavity. Second, the resonant mode is within the gain media spectral range. Mathematically, the laser modal wavelength  $\lambda$  relates to the laser resonance cavity length  $L$  by the equation  $N \cdot \lambda = 2 \cdot L$ .  $N$  is a positive integer larger than one. Despite their low cost, FP lasers are not applicable for coherent system in their current use due to limited performance in terms of multiple modes and linewidth of each mode.



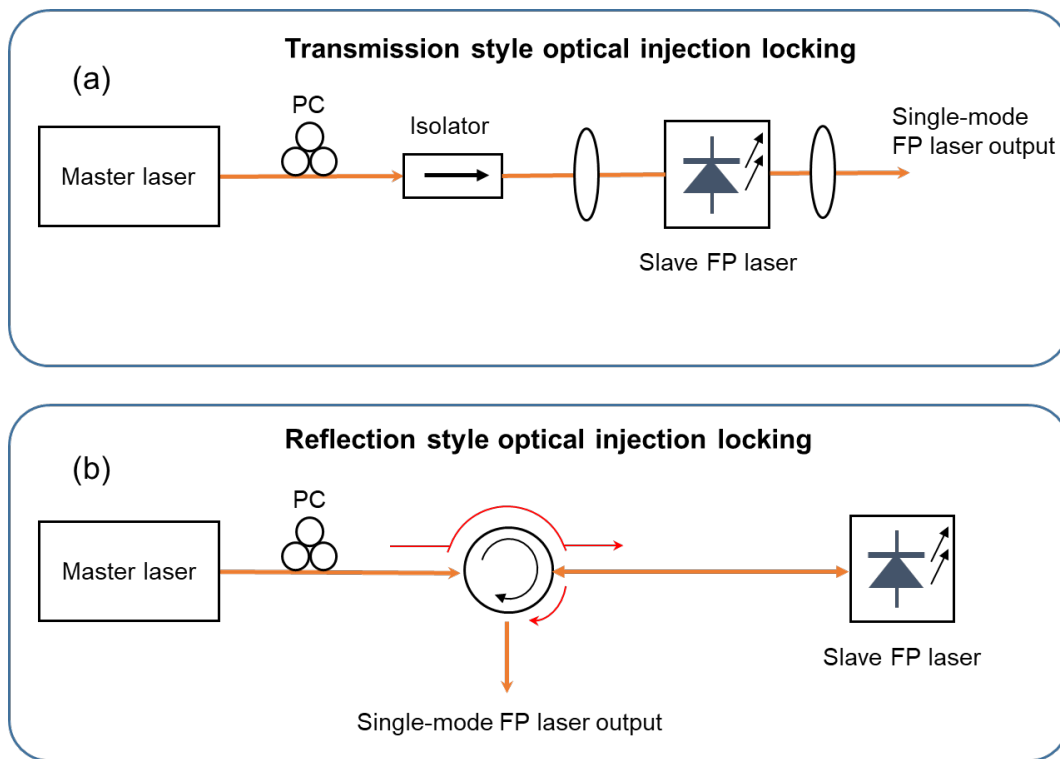
**Figure 2 – FP laser cavity modes**

In traditional operation of an FP laser, the wider linewidth that is typical for a FP laser structure would result in excessive phase noise and the inability to distinguish one constellation point from the other. In addition, FP-LD generation of multiple modes above the gain threshold and transmission of all these simultaneously modulated modes leads to excessive interference with no practical way to distinguish information.

The question we asked ourselves is if there is a way to predominantly excite and promote the generation and emission of a single mode of very narrow linewidth that can meet the signal criteria of coherent systems. Traditionally, this has been achieved by changes in the laser structure which has been the case in the more complex laser structures such as distributed feedback lasers and ECLs. This paper describes how a simple FP laser structure can be excited with an external optical signal to generate a high spectral purity source that suitable for use in coherent systems. A specific optical signal processing called optical injection locking can be used in generating this high spectral purity signal with an FP laser diode and enable low-cost FP lasers based coherent systems.

### **3. Operation Principle of FP Laser Injection Locking**

Optical injection locking (OIL) refers to the phenomenon under which the laser subjected to the external injection (the so-called slave laser or child laser) is phase and frequency locked to the external signal originating from a free running laser (the so-called master laser or parent laser). Since a semiconductor laser can be locked to frequency and phase of an externally injected optical signal through the injection locking process, a low-cost multi-mode FP laser diode can be turned into single-mode operation by injecting a high-quality single-mode signal into its cavity. The injection locked laser systems can improve a host of fundamental limitations in simple FP lasers: single mode operation and side-mode suppression, enhanced modulation bandwidth, suppressed nonlinear distortion, reduced intensity noise and chirp, etc. [4-6]. Schematic setup of an optical injection locking system is shown in Figure 3 ((a): transmission style; (b): reflection style). In the transmission style system, the light from the master laser is injected into a slave FP laser. It's polarization is adjusted by a polarization controller (PC) to match the slave laser cavity mode. An optical isolator is placed between the master and the slave laser to prevent back reflection. The injected master light enters the slave FP laser cavity through one of its facets and the output light exits through the other facet. In the reflection style system, the injected master light and the output light utilize the same facet to enter or exit the FP laser cavity, a laser design feature commonly found in commercial products. The master laser injects light into the FP laser via a three-port optical circulator, and the output from the FP laser exits the same optical circulator. When wavelength of the light from the master laser is set within a certain frequency detuning range with respect to the slave laser, the slave laser's wavelength is pulled towards the master's, until the dynamics of the laser settle with its frequency and phase both locked to the master laser.




**Figure 3 – Optical injection locking setup: (a) transmission style; (b) reflection style**

## 4. Design of Injection-locking Scheme

The peak frequency is the master laser frequency which the FP laser under test locks. One important characteristic of injection locking of FP laser is that the locking frequency peak is much higher than the side modes of the FP laser. This parameter is called side mode suppression ratio (SMSR). As a rule of thumb, SMSR of 30 dB or higher is regarded as good optical injection locking. Higher SMSR will provide better transmission distance with reduced phase and amplitude noise. Practical SMSR needed depends on specific link budget requirement.

The Fabry-Perot cavity can be simplified as a mirror-gain medium-mirror (MGM) structure as shown in Figure 4, where the gain medium in the middle is isolated by two reflective mirrors from the air. These mirrors could be made of thin films or coatings, which are characterized by certain coefficients of transmissivity and reflectivity. When the light is projected onto the surface of the mirror, part of the light is reflected, and another part of the light penetrates through the mirror and injects into the gain medium. The injected light passing through the front mirror will be bounced multiple times inside the cavity. After each round trip in the cavity, there will be a small part of light leaked out through the front mirror. At a certain wavelength when all these lights can constructively add together, the output power will reach a maximum value. Such a wavelength is referred as an intrinsic longitudinal mode in FP-LD, where the light waves can form a stable standing wave inside the cavity. The wavelengths that fulfill the constructive interference condition will exhibit multiple peaks at the power transmission curve, which is the reason for the multi-mode nature in an FP-LD. In the following section, the multi-mode property of FP-LD and the related factors affecting the design of FP-LD are analyzed in detail.



**Figure 4 – Multi-beam interference model for an Fabry-Perot laser cavity**

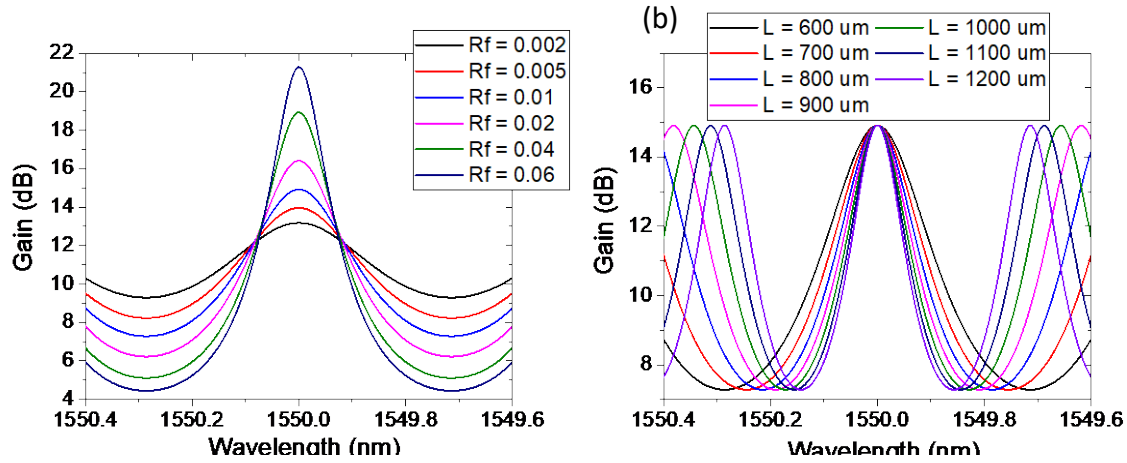
As shown in Figure 4, the transmission and reflection coefficients of the outer and inner layers for the front mirror are denoted as  $t_1$ ,  $t'_1$ ,  $r_1$ , and  $r'_1$  respectively. Similarly, the transmission and reflection coefficients of the inner and outer layers for the back mirror are  $t_2$ ,  $t'_2$ ,  $r_2$ , and  $r'_2$  respectively. Also, the gain and phase shift of one-round trip inside the cavity are provided as  $G$  and  $\delta$ , respectively. When electrical field  $A$  is given as the input, the output lights after the initial three round-trips inside the cavity can be given as  $r'_1 A$ ,  $t_1 r_2 t'_1 G^2 e^{j\delta t} A$ , and  $t_1 r_2 (r_1 r_2) t'_1 G^4 e^{j2\delta t} A$ . Adding these terms together, the output electric field of the light is

$$\begin{aligned} E_{out} &= r'_1 A + t_1 r_2 t'_1 G^2 e^{j\delta t} A + t_1 r_2 (r'_1 r_2) t'_1 G^4 e^{j2\delta t} A + t_1 r_2 (r'_1 r_2)^2 t'_1 G^6 e^{j3\delta t} A + \dots \\ &= r'_1 A + \frac{t_1 r_2 t'_1 G^2 e^{j\delta t} A}{1 - r_1 r_2 G^2 e^{j\delta t}} = A \left( \frac{r_1 - r_2 G^2 e^{j\delta t}}{1 - r_1 r_2 G^2 e^{j\delta t}} \right). \end{aligned} \quad (1)$$

Based on Stokes relations given as  $r = -r'$  and  $tt' + r^2 = 1$ , then the output power of the light can be calculated as

$$P_{out} = E_{out} \cdot E_{out}^* = A \cdot A^* \left| \frac{r_1 - r_2 G^2 e^{j\delta t}}{1 - r_1 r_2 G^2 e^{j\delta t}} \right|^2 = P_{in} \frac{(r_1 - r_2 G)^2 + 4r_1 r_2 G^2 \left( \sin\left(\frac{2\pi l \Delta \nu}{u}\right) \right)^2}{(1 - r_1 r_2 G)^2 + 4r_1 r_2 G^2 \left( \sin\left(\frac{2\pi l \Delta \nu}{u}\right) \right)^2}, \quad (2)$$

where  $\Delta\nu$  is the frequency deviation from the central resonance frequency, which is calculated as  $\Delta\nu = \nu - \nu_0$ ,  $l$  is the cavity length, and  $u$  is speed of light inside the gain medium, which is calculated as  $u = c/n$ . Based on Equation (2), given the reflectivity of the back facet of the cavity as  $R_b = |r_2|^2 = 0.99$ , the round-trip gain as  $G = 2$ , refractive index as  $n = 3.5$ , which is a typical value for InP, and the cavity length as  $l = 600\mu\text{m}$ , the gain spectra of the Fabry-Perot cavity as functions of reflectivity at the front facet,  $R_f = |r_1|^2$ , are plotted in Figure 5 (a).




**Figure 5 – Gain spectra of the Fabry-Perot cavity versus (a) reflectivity of the front facet, and (b) cavity length**

It can be observed that the gain is increased, and the bandwidth is also narrowed with a larger front-facet reflectivity. However, it is worth noting that in the real situation, a gain increase will be limited by the gain saturation effects of the gain medium. Moreover, a large front-facet reflectivity results in less optical power to be injected inside the Fabry-Perot cavity, thus reducing the injection-locking efficiency. The narrowing of the gain bandwidth also results in higher sensitivity of the FP laser towards frequency offset and drifts. Figure 5 (b) plots the gain spectra of the FP laser under different cavity lengths, where the reflectivity at the front and back facets are  $R_f = 0.01$  and  $R_b = 0.99$  respectively, round-trip gain here is set as  $G = 2$ , and refractive index is  $n = 3.5$ . It can be observed that the free spectral range (FSR) between the adjacent oscillating wavelengths will be reduced with a longer cavity length. The increased density of the oscillating wavelengths means it becomes easier for the wavelength of the injected light to be matched with one mode for injection locking. However, with a fixed size of the gain area, longer cavity length will increase propagation loss thus reducing the output power of the FP-LD.


Except from the impacts of frequency detuning towards the injection ratio and SMSR, another important aspect is the required power to achieve injection locking given a certain SMSR. The sample FP laser exhibits one original longitudinal wavelength at around 1563.09 nm in our design. During the injection-locking process, small wavelength detunings are introduced for the master laser and, with reducing the wavelength detuning, four wavelengths are tested, including 1563.03 nm, 1563.05 nm, 1563.07 nm, and 1563.10 nm. The measured SMSR versus injected power under different wavelength tunings is shown in Figure 6. An interesting phenomenon is that although the maxima SMSR can be obtained with a smaller wavelength detuning, the operational injected power range is actually expanded with a lightly increased wavelength detuning. The reason behind this may be that under the injection-locking operation, the laser

cavity will be slightly heated up causing red-shifts of the gain peak for an intrinsic longitudinal mode. There is also a trade-off between the optimal SMSR and operation power range for injection locking.



**Figure 6 – SMSR versus injection power at different wavelengths**


In our lab setup, the reflection-style coherent optical injection locking is used (Figure 3 (b)). Due to the requirement for a narrow frequency linewidth light source in a coherent system, a tunable ECL is selected as the master laser. The optical output power of the ECL is set to be 10 dBm. A FP-LD with 7-pin butterfly package,  $\sim 350\mu\text{m}$  cavity length, and a maximum output power up to +14 dBm is utilized as the slave laser. The ECL, with its wavelength tuned to match one of the longitudinal modes of the FP laser, injects light into the FP laser. Meanwhile, a DC bias is applied to the FP laser. The measured I-V curve (a) and output power at various bias voltage (b) of the FP laser are shown in Figure 7.



**Figure 7 – FP laser I-V curve (a) and output power vs. bias (b)**

By increasing the bias current gradually, one of the longitudinal modes of the FP laser is injection locked to the master ECL, while other modes are suppressed. As a result, single-mode operation of the FP laser is achieved at the selected wavelength. With its narrow spectral linewidth inherited from the ECL, this injection locked FP laser can be used as a coherent light source or local oscillator in the coherent system. Figure 8 shows the measured optical spectrum of the FP laser without external injection (a), and with ECL injection (b). The FP laser output power is set to be 5 dBm (1.11V bias, 20mA current). The center

wavelength of the ECL is set to be 1560.09 nm, which overlaps with one of the center cavity modes of the FP laser.



**Figure 8 – Optical spectrum of the FP laser: (a) free running; (b) injection locked**

The injection locking process, however, does not require a perfect overlap between the ECL and one of the FP laser modes. Under various injection conditions, when there's a frequency detuning between the ECL and the FP-LD modes, injection locking can still be achieved. It is important to understand the characteristics of the OIL laser frequency detuning, as it strongly affects the stability and reliability of the system. Theoretically, in the stable locking regime, operation of the OIL system can be described mathematically using a set of three differential equations [4-5]. The complex field of an injection-locked laser is very similar to that of a free running laser, with the addition of an injection term:

$$\frac{dE(t)}{dt} = \frac{1}{2}g\Delta N(1 + j\alpha)E(t) + kS_{inj} - j\Delta fE(t) \quad (3)$$

the slave laser's complex field  $E(t)$  can be split into three differential equations that describe its photon number  $S(t)$ , phase  $\phi(t)$ , and carrier number  $N(t)$ . When under steady state operation, since there's no time variation in the slave laser's photon number, phase, and carrier number, the time derivative parts of the equations are equal to zero. Follow the solutions in [5], the steady state injection system can be described by the steady state phase  $\phi_0$ , steady state carrier number  $\Delta N_0$ , steady state photon number  $S_0$ , and free-running photon number  $S_{fr}$ :

$$\phi_0 = \sin^{-1} \left\{ -\frac{\Delta\omega_{inj}}{k\sqrt{1+\alpha^2}} \sqrt{\frac{S_0}{S_{inj}}} \right\} - \tan^{-1}\alpha \quad (4)$$

$$\Delta N_0 = -\frac{2k}{g} \sqrt{\frac{S_{inj}}{S_0}} \cos\phi_0 \quad (5)$$

$$S_0 = \frac{S_{fr} - (\gamma_N/\gamma_P)\Delta N_0}{1 + (g\Delta N_0/\gamma_P)} \quad (6)$$

$$S_{fr} = \frac{J - \gamma_N N_{th}}{\gamma_P} \quad (7)$$

where  $g$ ,  $k$ ,  $N_{tr}$ ,  $\alpha$ ,  $J$ ,  $\gamma_N$ , and  $\gamma_P$  are the slave laser's linear gain coefficient, coupling rate, transparency carrier number, linewidth enhancement factor, current, carrier recombination rate, and photon decay rate, respectively. Under steady state, the phase value  $\phi_0$  falls into the range between  $-\frac{\pi}{2}$  and  $\cot^{-1}\alpha$  [4]. By choosing a  $\phi_0$  in the locking range, and rearrange (2) to determine the detuning frequency  $\Delta f =$

$\Delta\omega_{inj}/2\pi = (\omega_{ML} - \omega_{fr})/2\pi$ , which is the frequency difference between the master laser frequency  $\omega_{ML}$  and the slave laser free-running frequency  $\omega_{fr}$ . Also defining the field enhancement factor as  $R_{FE} = \frac{A_0}{A_{fr}}$ , which is the ratio between steady-state field magnitude and free-running field magnitude.

$$\eta_0 = \frac{c}{2n_g L} \frac{(1-R)}{\sqrt{R}} \sqrt{\frac{\eta_c P_{master}}{P_{slave}}} \quad (8)$$

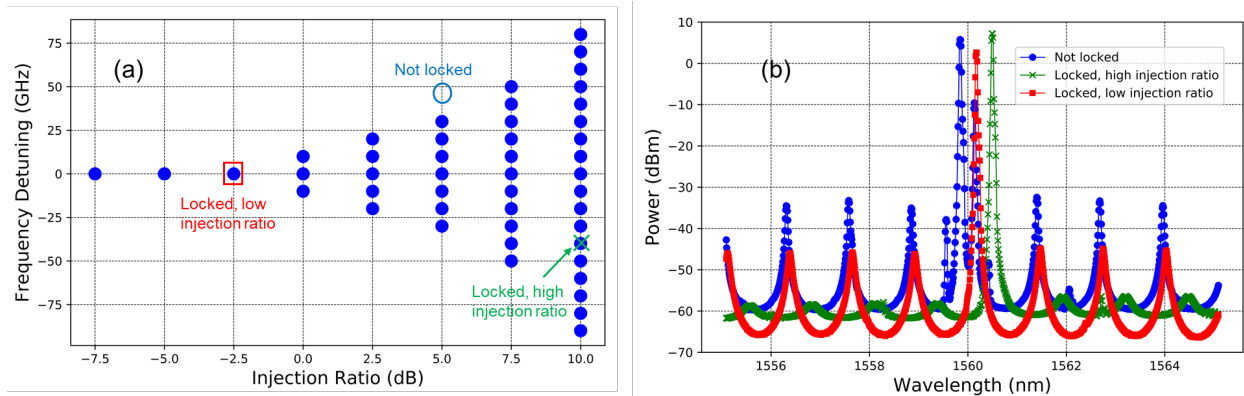
where  $L$  is laser cavity length,  $n_g$  is group index ( $n_g = 3.5$ ),  $R$  is reflectivity at the cleaved facet,  $\eta_c$  is the coupling efficiency of the laser ( $\eta_c = 0.6$ ), and  $\frac{P_{master}}{P_{slave}}$  is the external power injection ratio. These parameters are directly related to the laser physical design. The coupling rate  $k$  can be proportional to the injection ratio through:  $k \sqrt{\frac{S_{inj}}{S_0}} = \frac{\eta_0}{R_{FE}}$ . Here  $S_0$  is the steady state photon number and  $S_{inj}$  is the injected photon number. With the above expressions, the slave laser's detuning frequency can be described as:

$$\Delta f = -\frac{k\sqrt{1+\alpha^2}}{2\pi} \sqrt{\frac{S_{inj}}{S_0}} \sin(\phi_0 + \tan^{-1}\alpha) = -\frac{\eta_0}{2\pi R_{FE}} \sqrt{1+\alpha^2} \sin(\phi_0 + \tan^{-1}\alpha) \quad (9)$$

One can see that the detuning frequency between master and slave lasers is directly related to the steady state phase and the external power injection ratio.


We experimentally measured the injection locking process under various injection ratio and detuning frequencies. In Figure 9 (a), the ECL frequency is adjusted to introduce the frequency detuning from one of the FP laser modes. The injection ratio is defined as the ratio between the master ECL output power and the slave FP-LD output power. The injection ratio is varied by adjusting the master ECL output power via a variable optical attenuator, while the slave FP laser power remains at +5 dBm unchanged. The blue dots in Figure 9 (a) indicate the injection locking process occurs under corresponding injection ratio and frequency detuning. It clearly shows that under higher injection ratio, the injection locking process is more forgiving to frequency detuning between the master and slave lasers. The results are in good agreement with the theoretical calculations reported in reference [4] and [5]. Optical spectrums of the slave FP laser under different locking conditions are shown in Figure 9 (b). The red curve indicates injection locked FP laser spectrum under low injection ratio with minimum frequency detuning, where the green curve is under high injection and negative frequency detuning. The blue curve is showing the condition with medium level injection ratio, when the frequency detuning is too large the slave FP laser is not injection locked and still operating in multi-mode. The corresponding data points of the three curves are labeled in Figure 9 (a) respectively. One can see that under strong optical injection, the side mode suppression ratio (the difference in optical power between the main mode and the largest side mode in decibels) of the injection locked FP laser is much higher compared with that under weak optical injection.





**Figure 9 – Injection locking process study: (a) locking map under different injection ratio and frequency detuning; (b) optical spectrums of the slave FP laser under different locking conditions**

Under 10 dB injection ratio, the SMSR of the injection locked FP laser under various frequency detuning is extracted, as shown in Figure 10. Under positive frequency detuning, the SMSR tends to decrease with increasing frequency detuning. Whereas under negative frequency detuning, the SMSR improves first with increasing detuning, then starts to decrease after the detuning reaches a certain level. This mainly attributes to the fact that the linewidth enhancement factor of the master laser introduces carrier variation which will induce shift of the slave FP laser gain towards longer wavelengths.



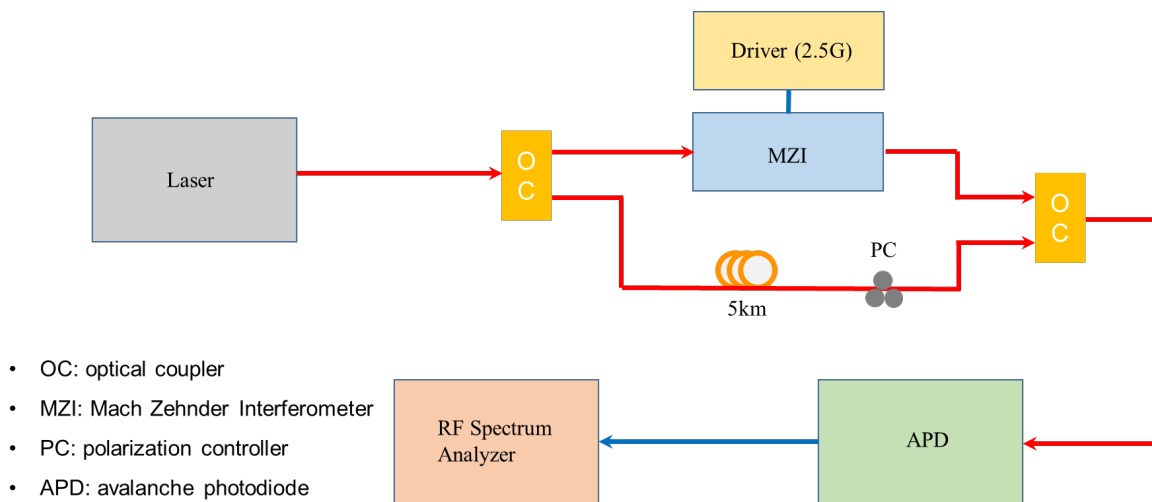
**Figure 10 – Side mode suppression ratio under various frequency detuning**

A key attribute of injection locking is that the detuning frequency range is asymmetrical with respect to the center frequency of the FP-LD side mode. If the master laser frequency is set at the center of the FP-LD side mode, it will achieve injection lock. However, if the master laser frequency is set slightly on the longer wavelength (lower frequency) side, it will have more tolerance to detuning. Therefore, it's more tolerant to maintain locking state in a larger frequency range compared with a “blue-shift” approach. This phenomenon is due to the linewidth enhancement factor of the master laser induced carrier variation which will induce the gain change of slave laser to longer wavelength.

## 5. Experimental Verifications


As the linewidth of the light source plays a critical role in the coherent optical communication systems, high resolution linewidth measurement were conducted on the injection locked FP laser and the comparison with a high quality ECL and a low-cost distributed feedback (DFB) laser was also performed. Due to limited resolution, grating-based optical spectrum analyzers (OSAs) are not suitable for this type of measurement. To achieve the desired resolution for characterizing our laser linewidth, the delayed self-heterodyne measurement technique was adopted here.

The laser linewidth measurement setup is shown in Figure 11. The output light from the laser under test is first split into two paths by a 3-dB fiber optic coupler. One path is sent through a Mach-Zehnder interferometer (MZI) driven with 2.5 GHz. The MZI shifts the detection frequency away from 0 Hz in the RF spectrum analyzer. Another path is sent through a delay line of 5-km SMF-28 fiber with a polarization controller. As a result, the laser light from the two paths are uncorrelated after the long delay. Both paths are then superimposed on another 3-dB fiber optical coupler while the resulting beat note centered at 2.5GHz is recorded with an avalanche photodiode (APD), and the RF spectrum is shown on the RF spectrum analyzer. With the interference between the two optical paths, the photocurrent generated in the APD includes both the direct intensity detection and the heterodyne frequency mixing term. As a result, the optical spectrum of the laser auto correlates with the delayed version of itself, in the frequency domain the detected autocorrelation function has a 3dB linewidth twice of the original laser 3dB linewidth.




**Figure 11 – Delayed self-heterodyne laser linewidth measurement setup**

Based on the delayed self-heterodyne method, the measured linewidth results are shown in Figure 12. The ECL linewidth is within the manufacturer's specifications, and the injection-locked FP laser shows similar performance. Both the ECL and the injection-locked FP-LD have linewidth under 50kHz, which ensures good performance in a coherent system.



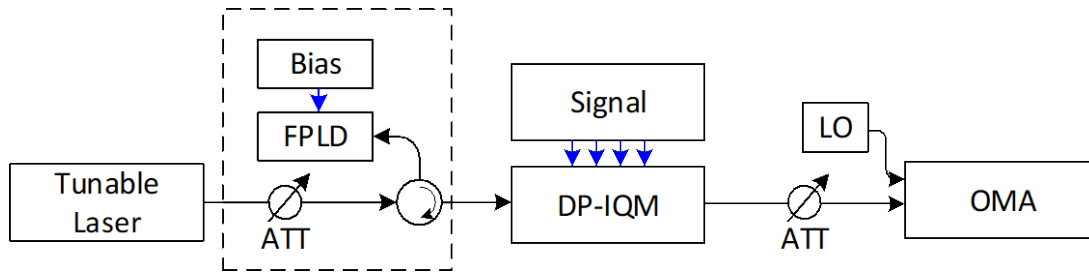
**Figure 12 – Measured ECL and injection-locked FP-LD linewidth**

As a reference, we also measured the linewidth of a DFB laser, with the results are shown in Figure 13. The DFB laser is also a single-frequency laser, and the cost is relatively low compared with the ECL. However, the linewidth of the DFB laser is also much wider compared with either the ECL or the injection-locked FP-LD, in the range of a few MHz.




**Figure 13 – Measured low-cost DFB laser linewidth**

The transmission performance of the coherent optical system with injection-locked FP-LD as the laser source has also been experimentally verified. The system diagram of the coherent optical link is shown in Figure 14.



**Figure 14 – System diagram for BER measurement**


The light from a tunable ECL is attenuated and fed into an FP laser for injection locking. The wavelength of the laser is carefully tuned to be around 1561.3 nm, which is matched with one longitudinal mode of the FP laser. A bias controller is used to control the output power of the FP laser to be around 10 dBm. Then light from the FP laser enters into a dual-polarization IQ modulator where the electrical signals are converted from the electrical to optical domain. A coherent optical modulation analyzer (OMA) is used to receive and sample the signals before offline digital signal processing is used for signal recovery. An optical attenuator is also used here to adjust the received power of OMA. In the experiments, three kinds of data modulation formats are applied, including 32-GBaud dual-polarization quadrature-phase-shift keying (DP-QPSK), 60.375-GBaud dual-polarization 16-ary quadrature-amplitude modulation (DP-16QAM), and 40-GBaud dual-polarization 64-ary quadrature-amplitude modulation (DP-64QAM), which correspond with total data bit rates of 128 Gbit/s, 483 Gbit/s, and 480 Gbit/s, respectively. The measured bit-error rates versus received optical power with injection locking and without it in the ECL case are shown in Figure 15 (a) to (c) in regards of DP-QPSK, DP-16QAM, and DP-64QAM, respectively. It can be observed that the power penalties brought by coherent injection-locking is insignificant for DP-QPSK and DP-16QAM, where the average power penalties are less than 0.3 dB. However, for the 40-GBaud DP-64QAM, the power penalty is slightly increased.



**Figure 15 – BER performance with and without applying COIL for (a) 32-GBaud DP-QPSK, (b) 60.375-GBaud DP-16QAM, and (c) 40-GBaud DP-64QAM**

## 6. Applications and Cost Analysis


The generation of a high spectral purity optical signal with an FP laser has been thoroughly demonstrated in earlier sections. The question remains how this approach can be leveraged to implement low cost coherent systems. The coherent system still needs a source that is spectrally pure which has higher cost. The answer lies in sharing the cost of the higher performance master laser signal. How the master laser is shared has implications in the cost and in the type of applications where this approach can be used. In the access environment, it is common to have hubs with serving areas covering 30,000 to 50,000 households. In a 500 homes passed per node architecture this results in 60 to 100 nodes per hub. In the downstream direction, there are typically dedicated fibers extending to the fiber nodes from a central hub location. It is therefore possible to share a master laser signal by distributing it among the many fiber nodes. Figure 6 shows how you can use a signal reaching the FP at -10 dBm and still injection lock it. Therefore if the master laser is transmitting at +10 dBm optical power, the 20 dB budget can be used to split the signal 64 ways in the hub sending it fiber nodes and remote destinations served from the hub and still have enough power budget for 40 km fiber links. A hub with one or two master lasers can remotely injection locked FP lasers at fiber nodes, base-station and businesses served from the hub (Figure 16).



**Figure 16 – Application I: Multiple fibers to different destinations**

Since the cost of the FP laser is about two orders of magnitude lower than the cost of an ECL, this use case leads to a significant cost reduction of P2P coherent optics links. The cost of the coherent laser could easily drop by a factor of 40. The split ratio of 64 was assumed using commercially available FP lasers. With optimization of the laser parameters such as the cavity length, facet reflectivities can further improve the injection locking effectiveness. A future derivative of this application could be the support of point-to-multipoint scenarios.

So far we have been describing the use of a single master laser with single frequency emission that is split towards multiple destinations. Recent technological advancement have been in developing comb lasers [7-8]. Comb lasers can generate multiple high-fidelity optical signals that are equally spaced in wavelength or frequency. The comb laser can be used to injection locked multiple FP lasers, each using a different tone generated by the comb laser. This in turn can be used to populate a single fiber with DWDM channels (Figure 17).



**Figure 17 – Application II: Comb source for DWDM systems over single fiber**

Using comb lasers with tones separated 100 GHz or 50 GHz apart could respectively generate 40 or 80 DWDM coherent signals in the C-band. Injection locking becomes a strategic tool for cost reduction when the network architecture lends itself to multiple destinations or multiple wavelengths. The high receive sensitivity of coherent optics also allows this sharing to be done remotely so that the cost of remote devices is reduced.

Wavelengths utilized are determined at the hub central location. Since the slave devices just lock to the master laser wavelength, the slave FP laser is not wavelength specific meaning only one part is needed in a remote location regarding which wavelength is to be used.

Systems leveraging optical amplification could achieve some of the functionality described here however cost and simplicity of a Fabry-Perot laser is most attractive.

## Conclusions

Coherent optics technology offers a future-proof solution for cable operators to meet bandwidth demand without the need for retrenching new reinforcement fibers. Cost reduction of coherent laser source is of great interest when bringing coherent optics into access applications. In this paper, a disruptive low-cost injection-locked FP Laser source has been proposed for coherent access networks. Through injection locking, the slave laser closely adopts the optical frequency and linewidth characteristics of the master laser which can be shared between multiple child lasers. As a result, the cost of the coherent optical transceiver can be significantly reduced. This paper introduces the operation principle of FP-LD injection locking as well as design theory. The detuning condition and linewidth characterization have been experimentally carried out with FP-LD samples. The comparison of transmission performance for different laser source is also presented experimentally. This paper also shows the promising applications of the injection locked FP-LD and the frequency comb in coherent access system.

## Abbreviations

APD	avalanche photodiode
BER	bit error rate
CD	chromatic dispersion
CMOS	complementary metal–oxide–semiconductor
CW	continuous-wave

DAA	distributed access architecture
DFB	distributed feedback (laser)
DOCSIS	Data-over-Cable Service Interface Specifications
DP	Dual polarization
DSP	digital signal processing
DWDM	dense wavelength division multiplexing
ECL	external cavity laser
FP-LD	Fabry-Perot Laser Diode
FSR	free spectral range
GHz	gigahertz
HFC	hybrid fiber-coax
ISBE	International Society of Broadband Experts
km	kilometer
LO	local oscillator
MGM	mirror-gain medium-mirror
MHz	megahertz
MZI	Mach-Zehnder interferometer
MZM	Mach-Zehnder modulator
ODC	Optical Distribution Center
OIL	optical injection locking
OLT	optical line terminal
OMA	optical modulation analyzer
P2P	Point-to-Point
PMD	polarization mode dispersion
PON	passive optical network
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying
RF	radio frequency
SMF	single mode fiber
SMSR	side mode suppression ratio
SNR	signal to noise ratio
SCTE	Society of Cable Telecommunications Engineers

## Bibliography & References

- [1] Cable Television Laboratories, Inc. “P2P Coherent Optics Physical Layer 1.0 Specification”, June 29, 2018. <https://specification-search.cablelabs.com/P2PCO-SP-PHYv1.0>
- [2] Cable Television Laboratories, Inc. “P2P Coherent Optics Physical Layer 2.0 Specification”, March 11, 2019. <https://specification-search.cablelabs.com/P2PCO-SP-PHYV2.0>
- [3] Z. Jia, L. A. Campos, C. Stengrim, J. Wang, C. Knittle, “Digital Coherent Transmission for Next-Generation Cable Operators’ Optical Access Networks,” Oct. SCTE/ISBE Cable-Tec Expo’17, 2017.
- [4] Q. T. Nguyen, P. Besnard, L. Bramerie, A. Shen, A. Garreau, et al., “Using optical injection of Fabry-Perot lasers for high-speed access in optical telecommunications”, SPIE Photonics Europe 2010, Bruxelles, Belgium. SPIE (ISBN 9780819481931), 7720, pp.77202D, (2010).
- [5] E. K. Lau, H. Sung and M. C. Wu, “Frequency Response Enhancement of Optical Injection-Locked Lasers,” in IEEE Journal of Quantum Electronics, vol. 44, no. 1, pp. 90-99, Jan. 2008.

- [6] E. K. Lau, L. J. Wong and M. C. Wu, "*Enhanced Modulation Characteristics of Optical Injection-Locked Lasers: A Tutorial*," in IEEE Journal of Selected Topics in Quantum Electronics, vol. 15, no. 3, pp. 618-633, May-June 2009.
- [7] M. D. G. Pascual, V. Vujicic, J. Braddell, F. Smyth, P. Anandarajah and L. Barry, "Photonic Integrated Gain Switched Optical Frequency Comb for Spectrally Efficient Optical Transmission Systems," in IEEE Photonics Journal, vol. 9, no. 3, pp. 1-8, June 2017.
- [8] V. Torres-Company et al., "Laser Frequency Combs for Coherent Optical Communications," in Journal of Lightwave Technology, vol. 37, no. 7, pp. 1663-1670, 1 April, 2019.



# **Cost-Effective, Scalable Quality of Experience (QoE) Monitoring for SD-WAN Networks**

A Technical Paper prepared for SCTE/ISBE by

**Edouard Karam**

Director, Solution Marketing & Product Strategy

Accedian

2351 Blvd Alfred-Nobel, Suite N-410, Saint-Laurent (Montreal), Quebec, H4S 2A9, Canada

1-514-331-6181

ekaram@accedian.com

**Greg Spear**

Senior Solution Manager

Accedian

2351 Blvd Alfred-Nobel, Suite N-410, Saint-Laurent (Montreal), Quebec, H4S 2A9, Canada

1-514-331-6181

gspear@accedian.com

## Table of Contents

Title	Page Number
Introduction	3
1. Executive Summary	3
Content	4
2. SD-WAN Value	4
3. SD-WAN Enterprise Adoption Overview	4
4. Service Providers Adoption Challenges	6
4.1. SD-WAN Adoption Complexities	6
4.2. Increased Network Complexities	6
4.3. SD-WAN Limited Visibility	6
4.3.1. Visibility Limited to WAN Edge	6
5. End-to-End Application QoE	7
6. Adding Active Performance Monitoring for Underlay Visibility	9
6.1. SD-WAN Point-to-Point or Point-to-Multipoint Deployment Model	10
7. Customization for Different Types of Organizations	11
8. Significant Cost Reduction Option for Passive Monitoring Solutions	12
8.1. Cost Optimizations Driving Cost Innovation	12
8.2. Stream to Disk versus Metadata	13
Conclusion	16
Abbreviations	17
Bibliography & References	18

## List of Figures

Title	Page Number
Figure 1 - IDC Worldwide Enterprise SD-WAN survey on cloud usage showing that the majority of enterprise plans to use SaaS within 12 months and a significant percentage of apps are accessed using the internet. (Source: IDC Worldwide Enterprise SD-WAN survey)	5
Figure 2 - SD-WAN vendor solutions monitor limited visibility	7
Figure 3 - Three essential pillars of visibility required to measure application QoE	8
Figure 4 - Complete application QoE visibility	9
Figure 5 - Centralized Gateway SD-WAN Model	10
Figure 6 - Point-to-point or point-to-multipoint business VPN	11
Figure 7 - Ubiquitous application QoE visibility across the entire cloud (prem=on premises, DC = datacenter, MoM=manager of managers)	12
Figure 8 - Most important Goals of Organizations' IT I&O	13
Figure 9 - Stream to Disc vs Metadata operation	14
Figure 10 - Stream to disc versus metadata storage capacity usage for 7 days	14
Figure 11 - Pros and cons of the two main approaches to passive traffic analysis	15

# Introduction

## 1. Executive Summary

Today, SD-WAN vendors offer some type of quality of service (QoS) visibility, but they do not extend this to true application quality of experience (QoE).

While SD-WAN solutions provide visibility into such things as network performance between platforms or bandwidth/capacity usage for top protocols, these metrics are provided only within the walls of the network. As such, they cover application performance from WAN edge to WAN edge only—not to the true edge (the end user's experience).

This paper will focus on the importance of, challenges surrounding, and requirements for true SD-WAN QoE visibility, including a look at costing benefits of metadata-based monitoring compared to traditional stream-to-disc.

Problematic QoE visibility are gaps created by the limitation of SD-WAN solutions:

- No way to pinpoint location/cause of application performance degradations
- Insufficient granularity to perform troubleshooting or optimization across the entire application delivery chain

Areas of visibility needed to measure real user experience include:

- Network performance
- Application delivery (through the network and infrastructure)
- Application transaction delay

Limitations of traditional QoE monitoring solutions using stream to disc include:

- Expensive to deploy and maintain
- Excessively short retention times
- Lack of flexibility to adapt to data growth (storage and interface types)

True SD-WAN QoE monitoring requires:

- Visibility into underlay network
- Visibility into end user application QoE
- Visibility into the root cause of application performance degradation

True, end-to-end QoE SD-WAN monitoring using metadata benefits operators with:

- Efficient capture of only the performance information of interest
- Long-term retention for context and future planning
- Fast performance degradation investigation
- High affordability compared to traditional monitoring

# Content

## 2. SD-WAN Value

The value for software-defined WAN (SD-WAN) varies somewhat depending on whether the organization deploying it is a communications service provider (CSP), managed service provider (MSP), or an enterprise.

For enterprises the ever-growing use of cloud-based applications makes SD-WAN more relevant every day. Software-defined networking (SDN), initially reserved for data center applications—along with a number of other technology enablers—have set the table for SD-WAN to disrupt traditional WAN architectural models prevalent within most enterprises.

From the service provider point-of-view, SD-WAN is appealing for similar reasons, primarily to offer a low-cost bandwidth enhancement to offered services—but also because it unlocks the ability to turn up new features on-demand, enhancing agility and speed for service delivery and the initial service turn-up.

Service providers typically deploy SD-WAN in their network to gain agility and flexibility first and foremost. The software automation at the heart of the SD-WAN solution allows for the creation of fully dynamic networks and give end-users a control into the nature and level of services they require on an ongoing basis. A crucial benefit for service providers is the ability to both deploy and maintain features and services via software deployment. Running software on commercial off-the-shelf (COTS) servers dramatically lowers both risks and costs when compared with the traditional dedicated hardware appliance solutions that required extensive trials to approve and the trained personnel, space, power and cooling to run.

Collectively (across enterprises and service providers), the main benefits of SD-WAN are:

- Access independence
- More bandwidth for less cost
- Cloud migration ease (remote site direct internet connection to the cloud)
- Ability to support hybrid networks (Applications no longer reside exclusively in the data center, and workloads are moving from enterprise data centers to private and public cloud)
- Application performance visibility over the network
- Automated provisioning
- Centralized policy control and management

But, there's a very important item missing from this list: quality of experience (QoE) visibility at the true edge of application delivery—the end user. SD-WAN vendors typically offer some type of quality of service (QoS) visibility, but they do not extend this to true application QoE.


## 3. SD-WAN Enterprise Adoption Overview

Enterprises are adopting SD-WAN as a way to gain access independence, leverage more bandwidth for less money, streamline cloud migration, automate service provisioning, and centralize policy control and management—among other benefits. However, the move to SD-WAN also involves moving a lot of the network control to the service provider's cloud.


The IDC enterprise SD-WAN survey below illustrates the adoption trend: two-thirds of survey respondents indicate they will deploy SD-WAN within the next two years. These results also show that cloud usage continues to rise, as does its importance in WAN technology selection:

- 70% of enterprises currently use infrastructure as-a-service (IaaS) and 90% plan to use platform as-a-service (PaaS) in the next 12 months
- A significant portion of enterprise apps are accessed using the internet


Greater than 70% use IaaS currently;  
90% plan to use PaaS in 12 months.



Q. What type(s) of cloud services or resources is your organization currently using and plans to use in the next 12 months?



Q. What type(s) of cloud services or resources is your organization currently using and plans to use in the next 12 months?



Q. What percentage of your enterprise applications are you currently accessing using the internet?

**Figure 1 - IDC Worldwide Enterprise SD-WAN survey on cloud usage showing that the majority of enterprise plans to use SaaS within 12 months and a significant percentage of apps are accessed using the internet. (Source: IDC Worldwide Enterprise SD-WAN survey)**

By moving to the cloud, enterprises can stand to benefit in several ways: more processing bandwidth for less cost, greater flexibility to scale resources up or down according to organizational needs, and reduced cost by taking away the burden of proprietary hardware or fixed circuits.

## **4. Service Providers Adoption Challenges**

### **4.1. SD-WAN Adoption Complexities**

Service providers are also adopting SD-WAN services as a way to enrich their offerings. But while SD-WAN brings significant, immediate benefits to the enterprise, it makes things more difficult for service providers. One of the largest challenges they face in adopting SD-WAN is dealing with increased complexities, which then translate to increased costs.

Many service providers support multiple SD-WAN solutions, because of preference for certain providers or to continue support for previous deployments. But, each SD-WAN vendor presents its own set of capabilities, management systems, and interoperability issues. This introduces complexities to manage:

- Within a single vendor, there are multiple product generations, models and update cycles to manage.
- Each protocol supports different control and monitoring functions.
- Reporting differs between SD-WAN systems:
  - Performance assurance standards are not uniformly implemented.
  - Different degrees of ‘compliance’ create interoperability gaps.
  - Different monitoring methods deliver different levels of detail.
- Reporting tools are not agnostic: there is no way to compare performance between vendors.

### **4.2. Increased Network Complexities**

In addition to SD-WAN adoption complexities, service providers also need to deal with the increased network complexity as a result of the emergence of virtualization and cloud-native applications. Finally, these complexities are worsened by users’ constant demand for more data, more bandwidth, and a seamless experience. Consequently, the challenge of assuring all parts of the network has never been more difficult for service providers who are already struggling to keep up with user experience visibility.

### **4.3. SD-WAN Limited Visibility**

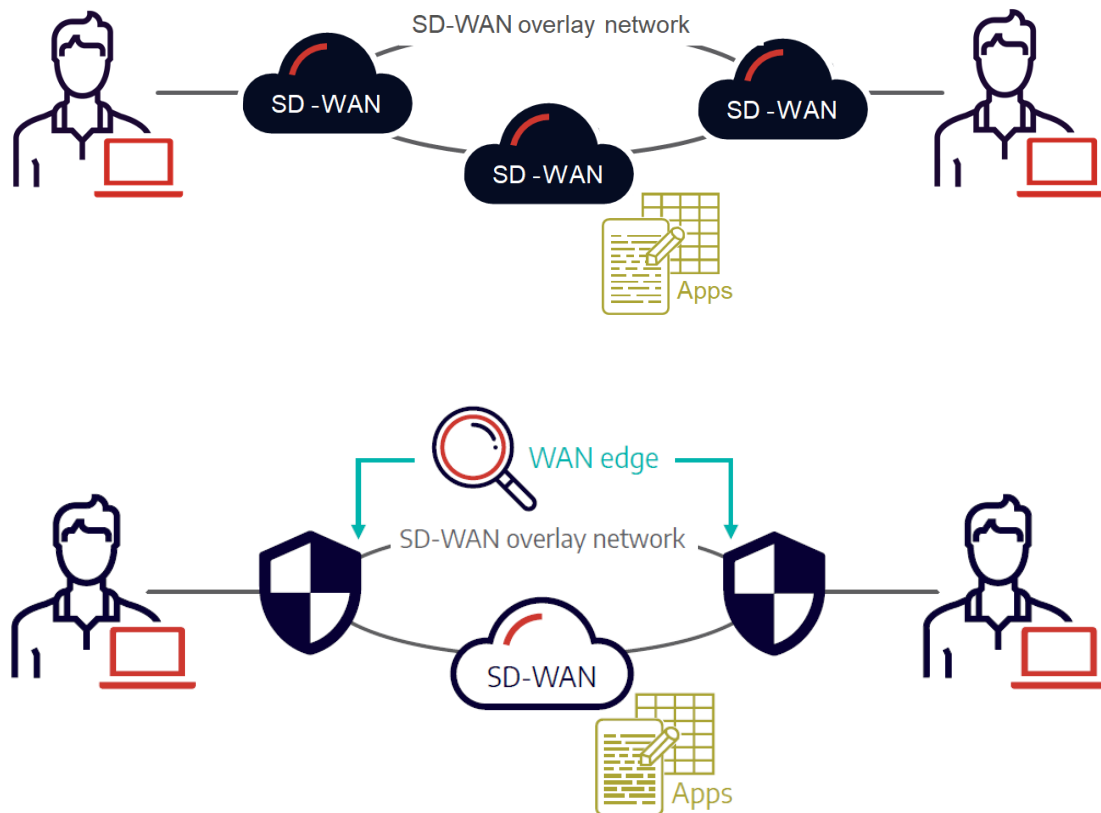
The solutions offered by SD-WAN vendors provide visibility of network performance between SD-WAN nodes (e.g. latency, packet loss) and usage (e.g. top protocols, top talkers). Because these metrics are only provided through Netflow or gateway nodes, SD-WAN solutions fail to offer visibility into network underlay without additional cost (such as additional hardware or software). Network level diagnostics required to provide meaningful data for mean time to resolve (MTTR) add compute burden to SD-WAN platforms (requires x86 resources to monitor the network).

Furthermore, by moving connectivity to the cloud, existing monitoring and performance assurance tools are unable to monitor what happens beyond their own network edge, unable to provide end to end application performance visibility and root cause of application performance degradation.

#### ***4.3.1. Visibility Limited to WAN Edge***

Some SD-WAN solutions provide visibility into such things as network performance between SD-WAN platforms or bandwidth or capacity usage for top protocols. But, because these metrics are provided

within the walls of the network, they cover application performance from WAN edge to WAN edge only—not to the true edge which is the end user’s experience.



**Figure 2 - SD-WAN vendor solutions monitor limited visibility**

As such, service providers experience visibility gaps outside of the confines of their wide area network (WAN) edge and will be unable to see problems that may be affecting users. For example: an inability to pinpoint location or cause of application performance degradations, or insufficient granularity to perform troubleshooting or optimization across the entire application delivery chain, creating problematic QoE visibility breaches. As such, WAN performance indicators may all be green but users could still be experiencing degradations.

Consequently, as more enterprise applications run on SD-WAN overlays, service providers will struggle to assure QoE because they don’t have visibility beyond their existing MPLS networks, and cannot correlate events/issues to ensure any application or network problems are being identified.

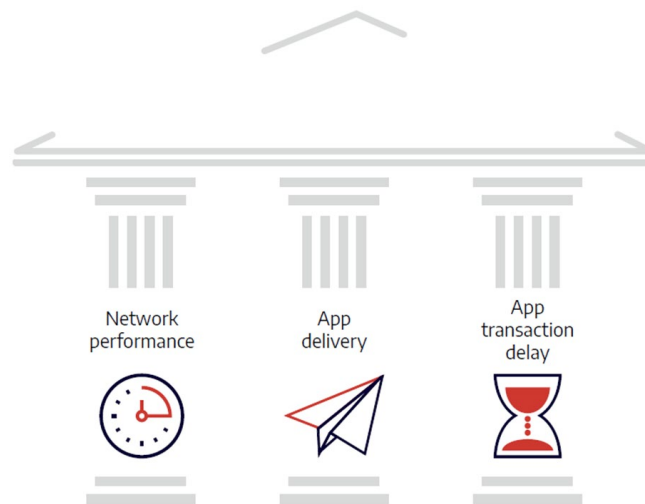
## 5. End-to-End Application QoE

For SD-WAN to show value to service providers and measure the real user experience (QoE), their performance assurance tools need to go beyond the confines of their MPLS or Carrier Ethernet networks, and beyond the overlay monitoring services offered by SD-WAN vendors. It also implies a performance

assurance management tool that is able to monitor the network from end to end and from layers 1 to 7 of the open systems interconnection (OSI) model, while bridging the existing visibility gap between network performance (layers 1-3) and application delivery (layers 4-7).

As such, measuring the real user experience requires full visibility into:

- Network performance
- Application delivery (through the network and infrastructure)
- Application transaction delay



**Figure 3 - Three essential pillars of visibility required to measure application QoE**

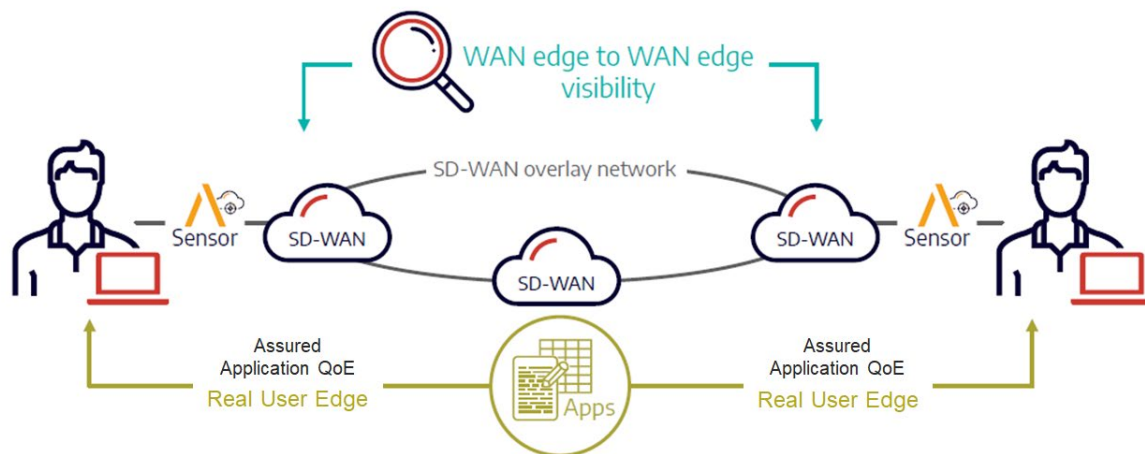
If any one of these is missing, it's not possible to have complete QoE visibility. Most SD-WAN vendor solutions fall short because they do not offer:

- Visibility into the underlay network
- Visibility into end user application QoE
- Visibility into the root cause of application performance degradation

Furthermore, these tools need to go beyond simply monitoring: they also need to be able to identify issues by building a profile of the network and application, what it is doing, and how it is being experienced by the end-user. This unification of network and application performance assurance will provide IT and service provider teams with a single source of truth and will help remove the silos created by having different teams monitor different parts of the network or application chain. Through this single platform, network operations, development, and business line owners can understand the interactions between infrastructure, application and user experience. What's more, this holistic view provides service providers with the opportunity to go beyond mere QoS and towards achieving, real-time full QoE visibility across the entire network chain.

With the right performance management tools that span both the network and application layers, service providers can meet the QoE demands of users, even in an SD-WAN environment.





**Figure 4 - Complete application QoE visibility**

## 6. Adding Active Performance Monitoring for Underlay Visibility

Most SD-WAN customer premises equipment (CPE) solutions provide passive analysis (bandwidth monitoring) of traffic/flows going and very rudimentary (ping / DNS query) methods of active link assurance. No SD-WAN CPE solutions have service activation testing (SAT) type capabilities built-in, which are key in providing underlay visibility. Such SAT and specific key performance monitoring (PM)-type capabilities for active testing include, but are not limited to:

- Service Activation Testing (SAT)
  - L2/L3 RFC2544 / Y.1564 - Standards-based service activation testing supporting commonly employed IEEE RFC-2544 and ITU-T Y.1564 turn-up testing approaches.
  - RFC6349 Framework for TCP Throughput testing
- Performance Monitoring (PM)
  - Y.1731 Ethernet OAM to ensure service availability meets SLA definitions, and to measure continuity and latency using CCM and DMM/DMR messages, respectively.
  - RFC-5357 Two-Way Active Measurement Protocol (TWAMP).

SD-WAN architectures virtualize some or all customer premises functions with a simple COTS server at the customer site. As part of their standard feature-set, SD-WAN solutions implement path monitoring and measurement. However, these measurements are typically insufficient for managed business services over SD-WAN deployments because those service assurance functions implemented purely in software:

- Lack sufficient time stamping precision and packet transmission scheduling control to meet the requirements of:
  - Full line-rate test traffic generation and loopback for SAT and troubleshooting.
  - Precise traffic generation sequencing required by common turn-up test standards (where inter-packet delay needs to be controlled for burst testing, for example).
  - Microsecond-level latency measurement precision required to monitor and report on commercial services SLAs.
- Are subject to the resource-sharing of the x86 system. This causes additional uncertainty in the results by bundling the performance of the x86 system with the performance of the network itself.

In addition, SD-WAN solutions typically use proprietary monitoring and reporting methods that do not interoperate with existing network equipment (or other SD-WAN vendors). Because SD-WAN may only

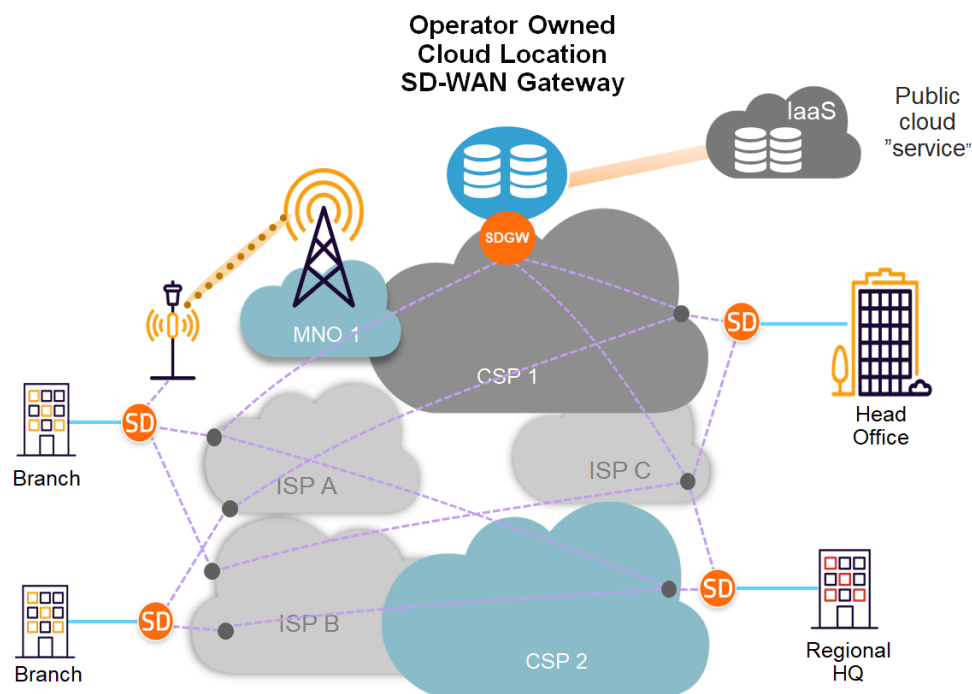
be required in certain locations, any service assurance implementation has to interact seamlessly with the traditional service delivery methods.

Relying on built-in SD-WAN monitoring creates a potential blind spot. This is especially true when considering SAT such as RFC2544 or ITU-T Y.1564 that have no support from SD-WAN vendors. Built-in SD-WAN performance monitoring functions can only provide a top-down view of performance—the over-the-top (OTT) path. This view presents no insight into why a specific path is operating badly; just that it is not performing. Complementing this top-down view with a bottom-up perspective provided by hop-by-hop or layer 2-3 path monitoring (such as SAT or PM) can add the missing pieces to more efficiently run an assured SD-WAN services, enabling detailed troubleshooting and measurable quality improvements

Running a unified assurance solution, which includes both active and passive monitoring, across both the incumbent part of the network and the SD-WAN part of the network also has the benefit of offering a unified level of precision and reporting intervals. As such, pinpointing events and segmenting the network will ease troubleshooting and accelerate mean time to resolution (MTTR) when issues arise.

### 6.1. SD-WAN Point-to-Point or Point-to-Multipoint Deployment Model

Many SD-WAN vendors offer an architecture based on centralized gateways to act as the virtual hub for any number of remote locations (spokes), as shown in Figure 9. The connected sites (branch, head office, HQ) need little hardware and a number of network transports (internet links or traditional WAN links) to establish the overlay network needed for the SD-WAN to operate (purple lines going to the SD-WAN gateway). The overlay network is built by having each remote site establishing encrypted tunnels to the SD-WAN gateway over each provisioned path.



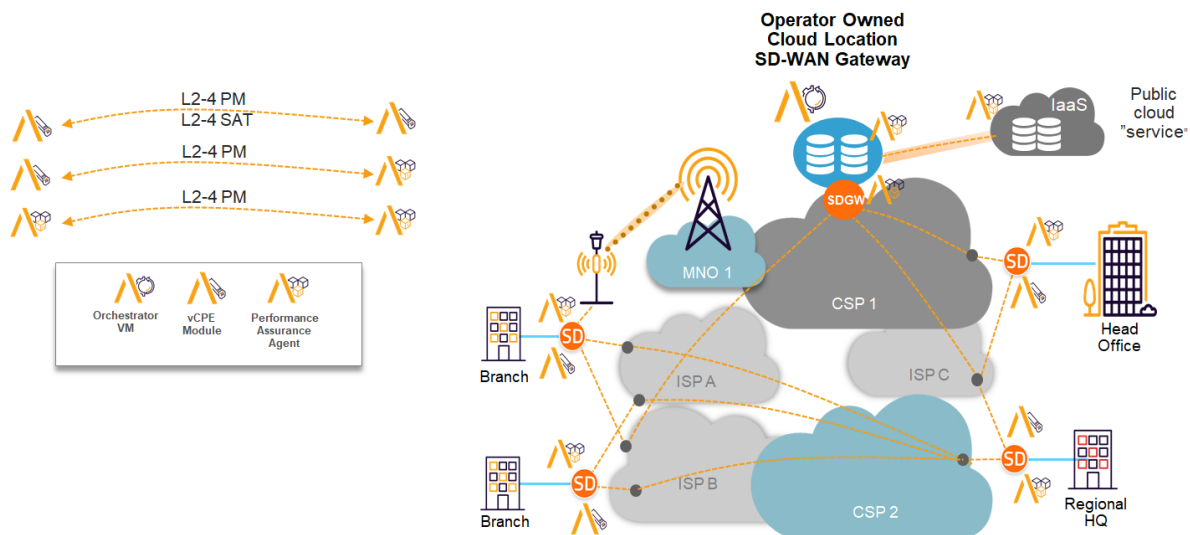
**Figure 5 - Centralized Gateway SD-WAN Model**

In such a model, deploying an active and passive monitoring solution in the SD-WAN CPE enables:

- Core-to-edge or edge-to-edge active PM and SAT
- Cloud gateway location-to-edge active PM and SAT

And in addition, for the cloud service(s):

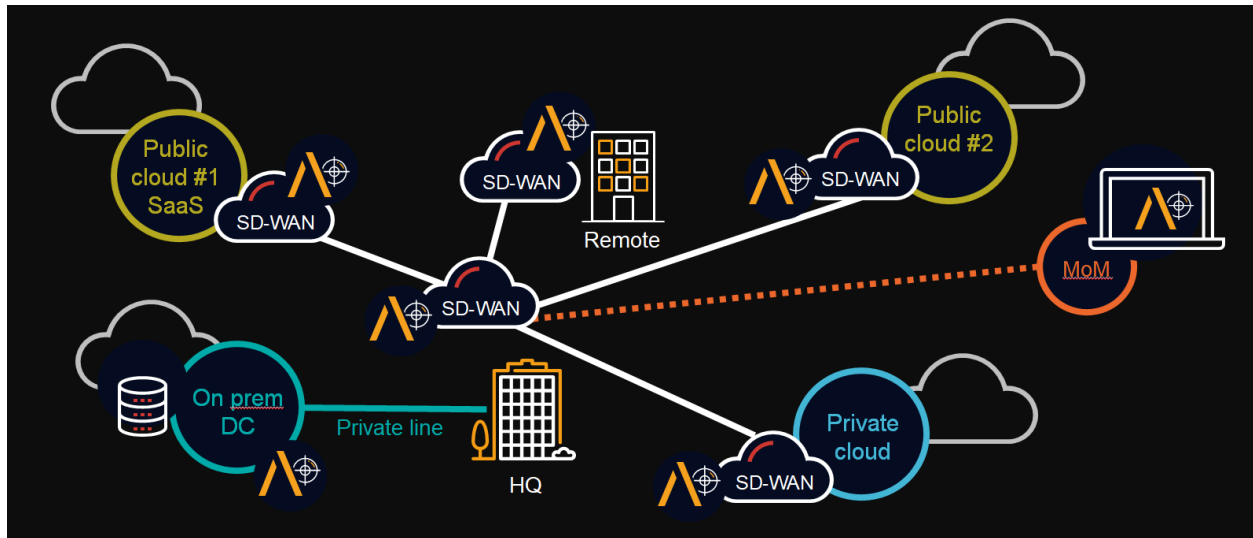
- Cloud gateway location-to-cloud service location
- Cloud service location-to-edge active PM and SAT



**Figure 6 - Point-to-point or point-to-multipoint business VPN**

## 7. Customization for Different Types of Organizations

Most organizations that move to SD-WAN have some reliance on the cloud but still haven't solved their monitoring problem. Fundamentally, monitoring user experience of network and cloud applications requires new methods and different metrics. Without these in-place before, during, and after an SD-WAN deployment, IT teams are left with little visibility and big headaches. This is where new solutions need to provide organizations with the insights and edge that they need, filling in the gaps, extending performance visibility to the real edge and going beyond QoS to true QoE.



**Figure 7 - Ubiquitous application QoE visibility across the entire cloud**  
(prem=on premises, DC = datacenter, MoM=manager of managers)


In short, such tools serve as a complement to SD-WAN monitoring capabilities, adding depth (underlay network monitoring) and breadth (application QoE) performance visibility and offering one single source of ‘truth’ about user experience.

## 8. Significant Cost Reduction Option for Passive Monitoring Solutions

### 8.1. Cost Optimizations Driving Cost Innovation

According to Gartner, 60% of enterprises will have implemented SD-WAN by 2024 (compared with 20% today), a change made to increase agility and enhance support for cloud applications. But, budgets are not expected to keep pace: the forecast for IT budget growth in 2019 is 2.7%, down from 3.1% in 2018, and network budgets are roughly flat. Meanwhile, 71% of IT budgets are dedicated to “running the business,” so that tends to be where the large focus of efficiency gains lie. With network budgets essentially flat, organizations need to do the proverbial “more with less.”

The figure below shows the top three goals of I&O leaders based on Gartner’s Infrastructure and Operations (I&O) Executive Leaders 2018 Survey. Lowering cost was the second most important goal for I&O leaders in 2018 and remains a critical goal still today.

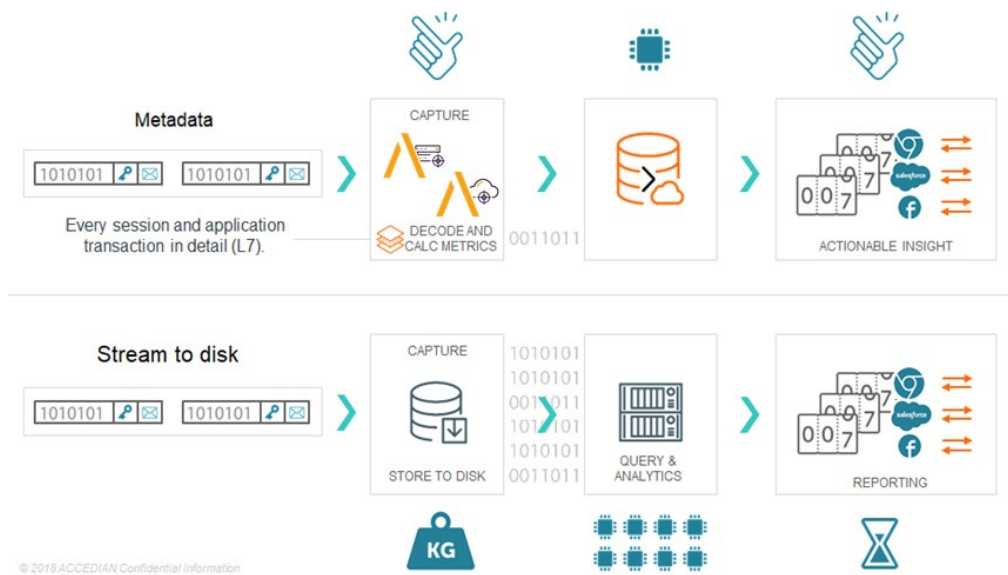


**Figure 8 - Most important Goals of Organizations' IT I&O**

Cost optimization will also be apparent with new on-demand consumption models, inspired by closer alignments with cloud principles. This will result in a migration from traditional capital-expenditure (capex)-centric “buy and manage” models to opex-centric services, where nothing is “owned” by the enterprise. This also offers the flexibility to scale up/scale down without cost penalties.

## 8.2. Stream to Disk versus Metadata

One often overlooked option that reduces the total cost of ownership (TCO) for passive application and network monitoring is leveraging metadata versus stream to disk methods. Below is an overview of how the two methods operate.



**Figure 9 - Stream to Disc vs Metadata operation**

The two technologies can provide similar results but their requirements and cost can vary greatly. For example, the large storage and CPU requirements for a stream to disk solution is much larger compared to metadata. As shown in the comparison table below, even though metadata storage does vary depending on type of requests, it still uses 20x to 100x less storage capacity when compared to stream to disk.

Storage	Stream to Disk	Metadata
12Gbps at <b><u>TCP level</u></b> for 7 days	<b>154 TB</b> 12Gbps x 7 days x 128B/750B	<b>2 TB - 8 TB</b>
12Gbps at <b><u>Application level</u></b> for 7 days	<b>907 TB</b> 12Gbps x 7 days x 750B/750B	<b>9 TB - 45 TB</b>

$$\text{Bandwidth}(\text{bps}) \times \text{RetentionTime}(\text{t}) \times \text{SlicingKey} / \text{AveragePacketSize} = \text{Storage}$$

**Figure 10 - Stream to disc versus metadata storage capacity usage for 7 days**

METHOD	PROS	CONS
Stream to Disk	<ul style="list-style-type: none"> <li>• Packet level analysis</li> <li>• Compliance</li> </ul>	<ul style="list-style-type: none"> <li>• Short retention</li> <li>• TCO+</li> <li>• Dynamic environment / cloud</li> </ul>
Metadata	<ul style="list-style-type: none"> <li>• Long retention</li> <li>• TCO</li> <li>• Speed of investigation</li> <li>• East-West visibility</li> </ul>	<ul style="list-style-type: none"> <li>• Packet level analysis</li> </ul>

**Figure 11 - Pros and cons of the two main approaches to passive traffic analysis**

Stream to disk retention time is often short; increasing retention time requires more storage, consequently increasing TCO. Furthermore, stream to disk is not fit for dynamic environment or cloud environments. When performing stream to disk, users will target specific points for traffic capture in their environment, defining the points that will capture traffic. This is harder to do in an environment where points are always changing such as a cloud environment.

Application and TCP monitoring using passive monitoring is the tool of choice when you want to get the most accurate picture of the user experience. But, you must be able to answer two critical questions: *What data do I need?* and *how far back in time do I need to look?* When it comes to retention periods, the bottleneck becomes storage space. As shown above, with stream to disc, storage space increases drastically. Storage comes at a cost, so TCO is always a concern.

To make this fit the TCO model, metadata is the preferred solution. The depth and breadth of metadata today will more than satisfy the requirements for TCP monitoring and, more importantly, application monitoring down to the transaction level for months rather than days. Metadata makes it much simpler and faster to retrieve, present, and correlate data. Plus, this makes it possible to go back in time, monitor progress during transitions, and provide data for collaboration.

Stream to disk does have its place (for compliance/regulation, for example) but is probably not the effective solution to monitor QoE for SD-WAN networks. It is simply too expensive, and retention times are not adequate for application and network troubleshooting. Metadata offers a lightweight software solution that aligns with the SD-WAN solution, and also offers the longer retention time that is critical for monitoring moves and changes.

To keep SD-WAN a cost effective, flexible and programmable solution, network and application performance monitoring tools should provide detailed information that aligns with the limited IT and network budgets of I&O leaders today.

# Conclusion

Today, SD-WAN vendors offer some type of quality of service (QoS) visibility, but they do not extend this to true application quality of experience (QoE).

While SD-WAN solutions provide visibility into such things as network performance between platforms or bandwidth/capacity usage for top protocols, these metrics are provided only within the walls of the network. As such, they cover application performance from WAN edge to WAN edge only—not to the true edge (the end user’s experience).

These limitations create problematic QoE visibility gaps:

- No way to pinpoint location/cause of application performance degradations
- Insufficient granularity to perform troubleshooting or optimization across the entire application delivery chain

Measuring the real user experience requires full visibility into network performance, application delivery (through the network and infrastructure) and application transaction delay. If any one of these is missing, it is not possible to have complete QoE visibility. SD-WAN vendor solutions fall short because they do not offer:

- Visibility into the underlay network
- Visibility into end user application QoE
- Visibility into the root cause of application performance Degradation

In addition, cost optimization goals are driving on-demand consumption models, inspired by closer alignments with cloud principles, resulting in a migration from traditional capital-expenditure (capex)-centric “buy and manage” models to opex-centric services, where nothing is “owned” by the enterprise. This also offers the flexibility to scale up/scale down without cost penalties.

Traditional application monitoring solutions measuring QoE are expensive as they typically use stream to disc capabilities—with short retention times, lack of flexibility to adapt to data growth (storage and interface types), and a high TCO. However, new solutions are available which use metadata (capturing only the information of interest). Such solutions offer longer-term retention and increased investigation speed at a fraction of the cost, aligning with the cost reduction that SD-WAN solutions deliver.

Metadata is best solution for SD-WAN, because it delivers depth and breadth of metadata to more than satisfy IT and network operator requirements delivering TCP monitoring and, more importantly, application monitoring (down to the transaction level for mounts not days). Metadata also addresses the limited IT and network budgets of I&O leaders today.



# Abbreviations

CBS	Committed Burst Size
CCM	Continuity Check Message
CFM	Connectivity Fault Management
CIR	Committed Information Rate
CoS	Class of Service
COTS	Commercial off-the-Shelf
CPE	Customer Premises Equipment
CPO	CoS Performance Objectives
C-VLAN	Customer VLAN
CE	Carrier Ethernet
DMM	Delay Measurement Message
DMR	Delay Measurement Response
DSCP	Differentiated Services Code Point
EBS	Excess Burst Size
EIR	Excess Information Rate
EMIX	Ethernet Mix
EVC	Ethernet Virtual Connection
FDV	Frame Delay Variation
FEC	forward error correction
FL	Frame Loss
FLR	Frame Loss Ratio
FTD	Frame Transfer Delay
Gb	Gigabyte
GUI	Graphical User Interface
HD	high definition
HFC	hybrid fiber-coax
Hz	hertz
IR	Information Rate
ITU-T	International Telecommunication Union – Telecommunication
KPI	Key Performance Indicator
LBM	Loopback Message
LBR	Loopback Reply
M Factor	Margin Factor
MAC	Media Access Control
Mbps	Megabit per second
MEF	Metro Ethernet Forum
MSO	Multiple Systems Operator
MTU	Maximum Transmission Unit
NID	Network Interface Device
NMS	Network Management System
NFV	Network Function Virtualization
NFVI	NFV Infrastructure
OAM	Operations, Administration and Maintenance
OSS	Operational Support System

PCP	Priority Code Point
QoS	Quality of Service
QoE	Quality of Experience
SaaS	Software-as-a-Service
SAT	Service Activation Testing
SCTE	Society of Cable Telecommunications Engineers
SLA	Service Level Agreement
SLS	Service Level Specifications
SOAM	Service OAM (IEEE Y.1731)
S-VLAN	Service VLAN
TWAMP	Two Way Active Measurement Protocol (ITU-T RFC-5357)
vCPE	Virtual CPE
VLAN	Virtual LAN
VM	Virtual Machine
VNF	Virtual Network Function
AP	access point

## Bibliography & References

1. *CE 2.0 Service Management Life Cycle White Paper*, July 2014; Metro Ethernet Forum
2. G.8013/Y.1731: *OAM functions and mechanisms for Ethernet-based networks*, November 2013; ITU-T
3. MEF 23.1: *Carrier Ethernet Class of Service – Phase 2*, January 2012; Metro Ethernet Forum
4. Y.1564: *Ethernet Service Activation Test Methodology*, March 2011; ITU-T
5. RFC-5357: *A Two-Way Active Measurement Protocol (TWAMP)*, October 2008, IETF
6. Vanilla+, January 2019: SCPs add value to SD-WAN by achieving full QoE Visibility  
<https://www.vanillaplus.com/2019/01/03/44207-csps-add-value-sd-wan-achieving-full-qoe-visibility/>
7. Gartner 2019 Strategic Roadmap for Networking.

# **Fifty Shades of Grey Optics:**

## **A Roadmap for Next Generation Access Networks**

A Technical Paper prepared for SCTE•ISBE by

**Venk Mutalik**

Executive Director  
Comcast

1401 Wynkoop Street, Denver  
860-262-4479  
Venk\_Mutalik@Comcast.com

**Bob Gaydos**

Fellow  
Comcast

1800 Arch Street, Philadelphia  
267-286-3214  
Robert\_Gaydos@Comcast.com

**Dan Rice**

VP

Comcast

1401 Wynkoop Street, Denver  
267-286-3214  
Dan\_Rice@Comcast.com

**Doug Combs**

Architect

Comcast

1800 Arch Street, Philadelphia  
267-286-3214  
Doug\_Combs@Comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
1. Distributed Access Architecture .....	4
2. The Axioms .....	5
3. Grey Optics .....	6
4. Multiwavelength and Singlewavelengths Systems .....	7
Distributed Access Architectures .....	9
5. Basic RPD DAA Architecture .....	10
6. Grey Optics Aggregation .....	11
GOA Performance .....	13
7. Downstream .....	14
8. Upstream .....	15
9. Additional GOA Options .....	16
GOA Operations .....	17
10. RF Levels across the Plant .....	17
11. Fiber Connectivity .....	18
12. Ingress Monitoring and Control .....	19
13. Monitoring the GOA and GOTs .....	20
Upgrade Options .....	20
14. Uneven Traffic Growth .....	21
15. Even Traffic Growth .....	22
16. Full Duplex in GOA Architecture .....	23
Economics of the GOA .....	23
17. Density and Aerial/Underground Plant .....	24
18. Economics and Discussion .....	25
GOA Trial .....	26
19. Trial Area Selection and Basic Illustrative Rules .....	26
Switch On A Pole .....	28
Conclusion .....	31
Acknowledgements .....	31
Abbreviations .....	31
Bibliography & References .....	32

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Illustrating Cost and Time Metrics for Fiber Construction and Electronics Deployment.....	6
Figure 2 – Illustrating the Optical Spectrim in a single fiber.....	7
Figure 3 – Taxonomy of Optical Impairments.....	8
Figure 4 – Typical MWL and SWL Performance.....	9
Figure 5 – RPD DAA Architecture .....	10
Figure 6 – GOA DAA Architecture.....	11
Figure 7 – Typical GOA Node (Left) and GOT Node (Right).....	12
Figure 8 – Block diagram of the GOA Node .....	13
Figure 9 – Measured MER values for RPD, GOA-GOT links and node.....	14
Figure 10 – Optical AGC in the DS Link .....	15
Figure 11 – GOA US Performance.....	15
Figure 12 – GOA Operations: RF Levels in the Plant .....	18
Figure 13 – Fiber Connectivity in the GOA architecture .....	19
Figure 14 – RPD, GOA, Analog and GOT DAA Nodes in use.....	21
Figure 15 – Upgrade path for uneven traffic groth .....	22
Figure 16 – Illustrating Even Traffic Growth.....	22
Figure 17 – Mapping Density in Cable Networks .....	24
Figure 18 – Total cost of DAA for RPD and GOA architectures .....	25
Figure 19 – GOA Trial Area Description.....	28
Figure 20 – Illustrating the Switch on a Pole (SOAP) Architecture .....	29
Figure 21 – Cascaded SOAP and the Converged Network.....	30

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – DRT RF Levels in the US .....	16
Table 2 - Example Table of monitorable and control parameters.....	20
Table 3 – Analyzing GOA Cost Reductions over RPD .....	26
Table 4 – Preliminary GOA Design Rules .....	27

# Introduction

Access network technology is evolving at an ever-increasing rate and the devices that are deployed in the outside plant may need to outlast multiple generations of change. In addition, as demand for bandwidth increases and the number of subscribers per node decreases, more access nodes are deployed, causing the cost per subscriber to increase. This requires an architecture that maximizes invested capital yet allows for flexible adoption of new technology.

This paper proposes a roadmap of ‘Grey Optical Aggregation’ (GOA) for the outside plant that will lower the cost of distributed access networks. As a rule of thumb, grey optics work well when distances are short and fiber is plentiful. Colored optics, like those used in DWDM systems, are best when the opposite is true. Also as a general rule, and comparatively, grey optics are inexpensive while colored optics are not. The challenge to the HFC industry is to optimize the use of each technology, deep in the plant, so as to maximize the fiber asset while keeping costs at a minimum. There are many ways to do this and hence there are many shades of grey optics.

The GOA architecture begins with lower cost grey optical nodes that are aggregated together at an RPD node location, allowing the subtended nodes to share the capacity of the 10Gbps DWDM Ethernet link. The roadmap culminates in a low powered, environmentally hardened ‘Switch On A Pole’ (SOAP) that multiplexes multiple 10Gbps grey Ethernet optics and leverages Coherent Optical links of 100Gbps and beyond to extend the headend into the outside plant as close to customer as possible. This allows the operator to pivot between or use multiple access technologies at the very end of the network easily.

In this paper, we begin with a description of grey optics and the benefits and tradeoffs relative to DWDM optics. We will then describe the process of incorporating grey optics aggregation in DAA networks and demonstrate the benefits of the GOA architecture. We then discuss operational aspects of this new architecture and the various upgrade options. We then describe the SOAP architecture and provide a stable roadmap towards supporting ever growing demands of the future while utilizing multiple access network technologies such as PON, DOCSIS and Ethernet.

## 1. Distributed Access Architecture

Broadband access networks continue to experience substantial growth in High Speed Data (HSD) capacity demands year over year. Estimates indicate that the Compounded Annual Growth Rate (CAGR) for HSD downstream (DS) is around 35%; for the upstream (US), CAGR is around 20%. Additionally, Comcast is seeing increases in its HSD customer base as well. All in all, prescient analysis of this trend has led to the development of the Distributed Access Architecture (DAA), which has enabled the company to keep up with the capacity needs while also enhancing customer satisfaction. Operational benefits also accrue due to the architecture [1,2].

DAA architecture development began in Comcast around 2015. The architecture called for the elimination of RF Amplifiers in the Outside Plant (OSP) and the reduction of Households Passed (HHP) per node by driving fiber deeper into the network. Initially, this was achieved by using multi-wavelength optimized Analog Optical Transmitters in the DS and Digital Return Transmitters in the US. This called for an optimized wavelength plan and optical passives that mitigated optical non-linearities and fiber effects. Extensive architecture rules and play books were developed to determine node placements and other plant adjustments so as to prepare the network for the future. To date, over a million HHPs have benefited from DAA; simultaneously, the customer experience improved and operational expenses reduced.

Around 2017, Comcast began the move from Analog DAA to Digital DAA. This involved virtualizing the CMTS, and investing in high speed connections from the Primary Headend (PHE) to the Secondary Headend (SHE). At the secondary headend, white box Distributed Access Architecture Switches (DAAS) were placed with 10Gbps DWDM connections to individual fiber nodes. These fiber nodes had the same envelope as the previous analog DAA nodes and were placed at similar locations, as determined by the aforementioned DAA architecture rules and playbooks. These fiber nodes now have a Remote PHY Device (RPD) that takes in the 10Gbps signal and transmits an analog signal in the DS to the home. Similarly, in the return path, the RF signals are combined and fed to the RPD located inside of the node. An architecture of this kind provides high Modulation Error Ratio/(MER)/Noise Power Ratio (NPR) in the DS/US and enables higher orders of QAM modulation, thus leading to increased overall capacity. Additionally, better monitoring capabilities and dashboards enhance the already substantial operational benefits across the board.


In this paper, we describe the Digital DAA architecture in some detail, and explain new architecture variants that can substantially reduce cost, preserve upgrade options and help drive fiber deeper.

## **2. The Axioms**

It is axiomatic that labor and infrastructure costs increase over time and electronics cost decrease over time. Furthermore, construction has various dependencies in the form of permits, crew availability, weather and materials availability. As such, there is an inherent strategic approach that needs to be taken in respect of new infrastructure deployment. Construction is an inherently slow and deliberate process.

Electronics, by contrast, change fast. With Moore's Law, computing speeds increase every 18 months. Dennard's Law predicts lower power consumption in later generation devices, and Koomey's Law predicts miniaturization in succeeding generations of electronics. Designers taking these benefits into account either offer lower cost or higher functionality -- or both -- by manipulating space/power/speed combinations.

It is worth noting that within 2 years of the initial analog DAA program, Comcast has evolved to the next generation of the DAA program, which is to harvest the benefits of speed/space/speed to virtualize and distribute the CMTS, and essentially put a substantial part of the CMTS functionality into the fiber node. Further work along these lines could move us towards Full Duplex DOCSIS (FDX), as well as high split or extended spectrum options. Additionally, there are convergence possibilities between 5G and broadband services in the edge.



**Figure 1 – Illustrating Cost and Time Metrics for Fiber Construction and Electronics Deployment**

Driving fiber deeper is of strategic importance to many service providers as a way to further enable multiple services. That said, today's fiber deployments also involve significant electronics dependencies in the form of virtualizing the CMTS, and deploying RPDs and DWDM optics, which may be an overkill for current capacity needs.


There is therefore an urgent need to orthogonalize fiber/infrastructure deployments and electronics deployments. Doing so would enable service providers to construct fiber and accrue economies of efficiency by deploying electronics when the technology is mature and the capacity needs are manifest.

This paper addresses this issue and proposes an architectural concept of Grey Optics Aggregation (GOA), which will be described next.

### 3. Grey Optics

Figure 2 depicts the spectrum in an optical fiber. Typically it stretches from 1260nm to 1620nm, and is divided in 20nm bands called the Coarse Wave Division Multiplex (CWDM) bands. Various Comcast entities have different optical assets utilizing several parts of the spectrum. However, the wavelength range from around 1525-1570 nm is especially heavily utilized due to the wide availability of erbium doped fiber amplifiers (EDFAs). Therefore this band is further divided in many smaller bands of approximately 0.8nm width or 100GHz of optical spectrum – there are 48 such wavelengths (and sometimes with 50GHz spacing, in which case there are more than 96 such wavelengths) in that band alone. The wavelengths in this band are therefore called the Dense Wavelength Division Multiplex (DWDM) band.





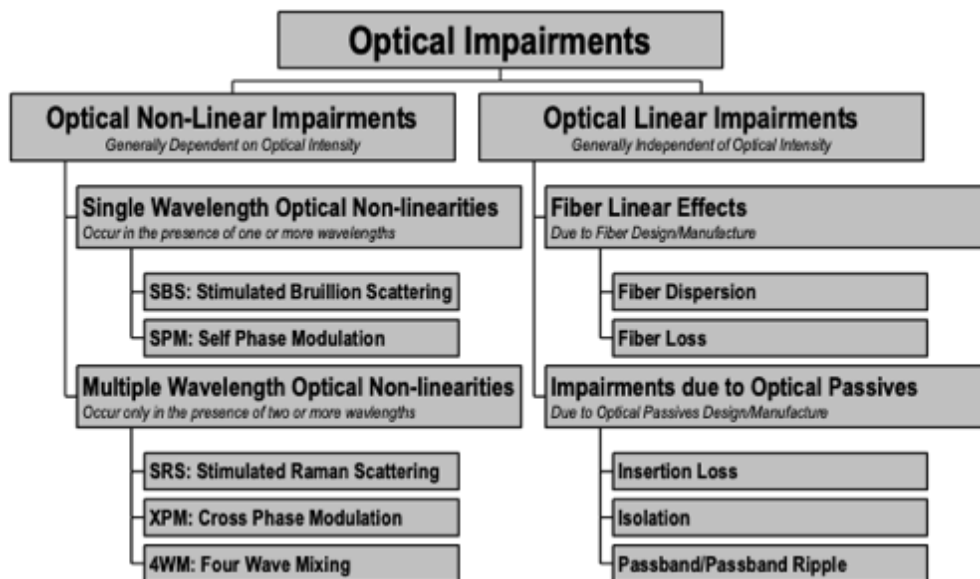
**Figure 2 – Illustrating the Optical Spectrum in a single fiber**

DWDM Optics have a high requirement to maintain the optical wavelength of their optics. Typically they have to maintain their wavelengths to within  $\pm 0.1$  nm from beginning of life (BOL) to end of life (EOL). Designing and operating optical equipment with this stringent requirement is what increases the cost of DWDM optics. In common parlance these are also called Colored Optics. As difficult as the manufacture of DWDM Optics is, they are used routinely now for long haul optics and in trunk fibers. This is because the higher cost of DWDM optics is fully justified, in that the cost of fiber construction is even higher for those trunk links.

Grey Optics, on the other hand, do not maintain optical wavelengths in so tight of an optical range when the density is not required. Grey Optics wavelengths can move 10 - 50 times in spectrum location more than the Colored Optics. This enables significantly lower costs ( $\sim 10\times$  lower) for these, compared to the DWDM optics. Since each of these wavelengths is unconstrained, just one or very few of these wavelengths can be used in one fiber, which is why these are best suited for small distances in the edge, where new fiber construction for DAA type architectures deploy high fiber count cables.

## 4. Multiwavelength and Singlewavelengths Systems

Presented below is a simplified taxonomy of optical impairments in an optical fiber [3]. When multiple wavelengths (MWL) course through an optical fiber, and at high optical levels, many optical effects and non-linearities emerge and can impact performance. For this reason, Analog DAA, which uses up to 16 wavelengths, uses a carefully optimized industry standardized wavelength plan. In addition, launch optical power, fiber reach, receiver input power, choice of optical passives and transmitter design all have major impacts on performance.



**Figure 3 – Taxonomy of Optical Impairments**

Multiwavelength systems are affected by multiple linear and non-linear impairments, which can include Shot noise, laser RIN, Receiver EIN, EDFA noise, optical passives ripple and fiber dispersion effects. In addition, impairments from fiber crosstalk, due to imperfect passives isolation, and optical non-linearities such as Stimulated Raman effect (SRS), Cross Phase modulation (XPM) and Four Wave Mixing (4WM), can occur. The resultant MER of DS wavelengths are relatively modest. Indeed, the ability to identify all impairments and manage their impact simultaneously is at the heart of making robust multiwavelength systems work.

Plant Composition vs. Plant Density Matrix		
	Multiwavelength System 16WLs + Passives	Single Wavelength System C-Band
Short Reach <2km	MER ~40-42dB	MER ~46-48dB  Well Suited for GOA
Long Reach >30km	MER ~38dB  Used in Analog DAA	MER ~40-42dB

**Figure 4 – Typical MWL and SWL Performance**

However, when single wavelength (SWL) systems are used, and for a very limited reach, system design is much simpler and can be shot noise limited. As a result, much higher MER values are obtained. It is not the intent of this paper to describe MWL and SWL system design, but this important outcome is highlighted here because this result enables us to use analog optics without sacrificing DS performance when we next describe the Grey Optics Architectures.

## Distributed Access Architectures

Hybrid Fiber Coax (HFC) has been the mainstay of Cable Networks and continues to offer significant benefits over other forms of last mile access, such as twisted pair, commonly used by telcos. Typical HFC systems use fiber nodes and RF amplifiers. While RF amplifiers help drive signals, they are numerous in number, require maintenance, consume power and vary over temperature. For this and other reasons, Comcast began to eliminate RF amplifiers and move to ‘N+0’ or “node plus zero” amplifiers. N+0, also called DAA interchangeably, reduces the number of actives in the plant by half, strategically drives fiber deeper into the network and reduces HHP/node, increasing capacity. Since the fiber is driven deeper into the network, more RF spectrum and the move to DAA is often associated with increasing the DS to 1 GHz or 1.2GHz (from 750MHz-860MHz) and to 85MHz in the US (from 42MHz) typically found in HFC. This increase in spectrum and the decrease in HHP/Node together give a very robust runway for future capacity growth, while the reduction of OSP actives and increased MER reduces maintenance and enhances the customer experience.

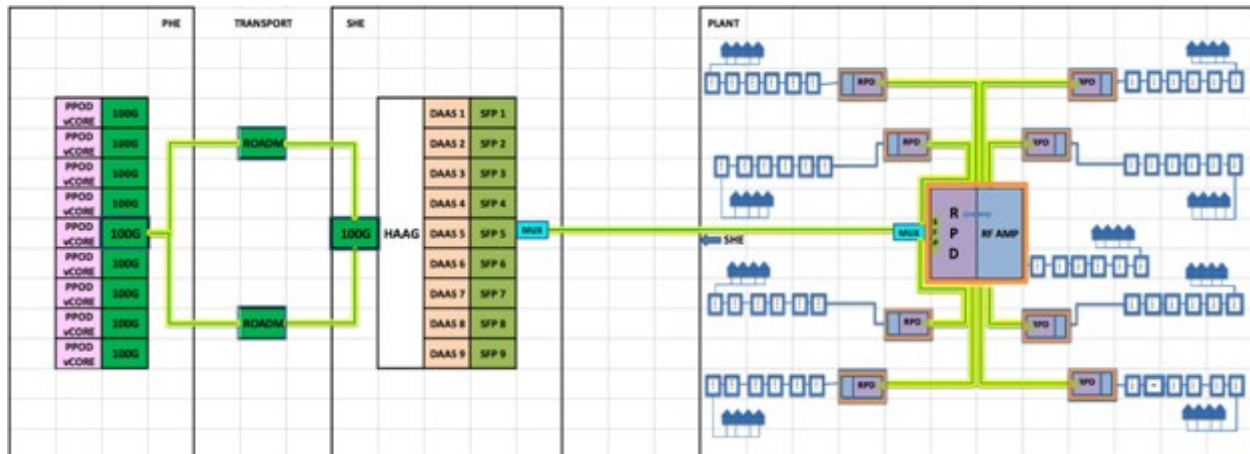
An RPD is a Remote PHY Device that accepts 10Gbps signals optically and converts the signals to conventional RF output. Since this conversion is done using advanced techniques such as Direct Digital Synthesis (DDS), the ensuing MER out of the RPD is very high -- as high as 48-50dB with full 1.2GHz

loading. This is headend-grade signal quality, right at the node, and is the primary draw for a move to DAA and to incorporate RPDs in fiber nodes.

## 5. Basic RPD DAA Architecture

A strawman architecture of RPD DAA is presented in Figure 5. Typically, a conventional HFC ‘Parent’ node is divided into many smaller DAA nodes. Normally, near the parent node location, an optical passive is placed that has up to 48 ports. Then fibers are drawn to each of the nodes which have an RPD with DWDM Small Form Factor Pluggable (SFP) modules connected to optical passive in the parent node location. The RPD consumes 10Gbps of data and creates RF signals for distribution in the DS. In the US, burst mode RF from all of the ports is combined and fed to the RPD, which takes the analog input, demodulates it, converts it into a digital stream, and sends back to the headend via the DWDM SFP. Although the picture below shows just 9 RPDs, it is typical to see ~12 RPDs on average and as many as 24 RPD DAA nodes in one parent node location.

The muxed optical signals are then de-combined at the SHE (secondary headend). The SHE comprises an optical passive that is a mirror of the one at the parent node location, which then feeds as many SFPs as were used in the node of the outside plant. Each of these SFPs are housed in the DAAS switch with at least as many ports as SFPs. Oftentimes these DAAS ports are aggregated into a switch called a Headend Aggregation (HAAG) switch, and the data is connected to the PHE via 100-400Gbps high speed coherent links that may traverse thru multiple ROADMs (reconfigurable optical add drop multiplexers). At the PHE (primary headend), these links are then connected on a one-to-one basis to virtualized CMTS also called vCOREs or vCMTS that are located in server racks, or Pods called PPODs (Primary Pods) that mirror their OSP configurations.



**Figure 5 – RPD DAA Architecture**

It is easy to see in Figure 5 that each RPD requires 2 SFPs, two pairs of optical passive ports, one DAAS port, access to the HAAG and 100-400 Gbps optics along with a connection to one vCORE. While RPD DAA was designed with about 80-100 HHP/Node in mind, it is common to see much lower densities of around 55 HHP/Node in the outside plant. When this happens, the number of RPD nodes increases proportionally, and along with the cost of RPDs in the OSP, one would also encounter significant costs in the ISP in the form of SFPs, DAASs, HAAGs and vCOREs, and the consequent increases in critical infrastructure (air/power and space) in the ISP locations.

Typically, DAA conversions are done on a secondary headend to secondary headend (SHE to SHE) basis, to enable the move from conventional CMTSs to vCMTSs. This helps the SHE to maintain just one architecture in its inside plant (ISP), and is a more sustainable way to support future capacity needs. Furthermore a move of this kind all but eliminates RF combining mechanisms prevalent in SHEs, and frees up the critical infrastructure for other uses. Perhaps most importantly, this is the only way to plan, build and commission the OSP. OSP construction is a complicated effort, with multiple moving parts from designs to walkouts to permits to traffic control and weather and seasonality. Aerial construction is comparatively simpler and less expensive than underground construction, which could be much more expensive (10-15x) and time consuming.

When density in the OSP turns out to be low, then the consequent increases in RPD node counts can extract substantial costs across both OSP and ISP. Even if the density on average is high, such as that in an urban or suburban location, there are still going to be several pockets where the density is much lower than optimal.

Electronics in the RPD that cost a lot today will cost substantially less in a couple of years. Tying up capital in electronics for small node sizes, where such capacity is not currently needed, would be much better spent on fiber construction activities, since fiber is an enduring asset that lasts for decades. Finally, technology itself is changing fast and the competitive landscape can call for FDX, ES, High Split, 5G front and back haul, EPON/FTTH or SOAP, or a combination all above.

There is thus a need to have an architecture option that enables stable, predictable and cost effective DAA implementation across various densities, and one that has multiple upgrade paths to accommodate capacity increases and transparent to technology changes.

## 6. Grey Optics Aggregation


Consider the architecture presented in Figure 6. From previous discussions, we know that trunk fiber is scarce with a long reach, and thus is suitable for long reach MWL DWDM 10Gbps optics. The plant fiber from the primary node location to all the child node locations is very short (0-2km), and more importantly plentiful, since it has been freshly laid out, and is thus suitable for SWL Grey Optics. Especially in low- to mid-density areas, where the number of HHPs/Node could be very low, this architecture provides a way to aggregate a constellation of DAA nodes.



Figure 6 – GOA DAA Architecture

To do this, an RPD Node is modified to have some of its RF output connected to a DS Grey Transmitter. The cable industry has for a long time been using DWDM Grey Analog transmitters for US transmission. Furthermore, there has been a 4x miniaturization of DS DWDM transmitters over the last decade. In this case, both techniques are used together and transmitter bandwidth is enhanced to cover a range of 100MHz to 1.2GHz while maintaining its temperature performance over the industrial temperature (I-Temp) range (-40 to +85C). Furthermore, the wavelength control over the transmitter output is relaxed, which, when combined with the very small link spans, creates the perfect opportunity for a downstream Grey transmitter (SWL, analog) with robust MER and low cost. Typically, these transmitters have 10dBm of output power, which, when split 8 ways and needing to cover <2km, typically arrives at the constellation node receiver at around 0dBm.

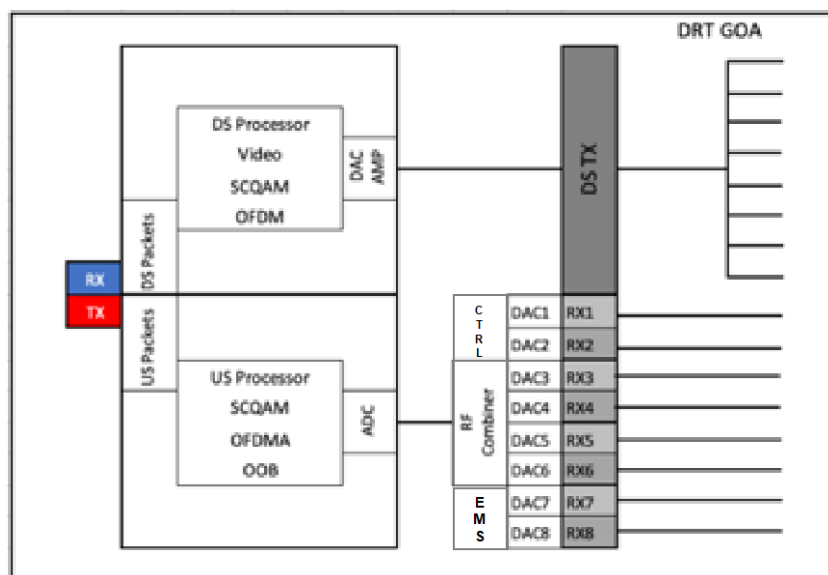
We call this central node the Grey Optical Aggregating Node (GOA), and the constellation nodes the Grey Optical Terminating Nodes (GOT). With this terminology settled, and with a friendly reminder that “GOT” in this case does not stand for “Game of Thrones,” we can now describe downstream and upstream transmission.



**Figure 7 – Typical GOA Node (Left) and GOT Node (Right)**

The GOT Nodes have analog receivers, just like in Analog DAA nodes. However, unlike the Analog DAA nodes, the GOT nodes include optical receivers with Optical AGC functionality. Although typical optical input coming is around 0dBm, as indicated earlier, the Optical AGC enables the receivers to put out a constant RF output, regardless of minor changes in received power. This concept will be described later in the paper.





**Figure 8 – Block diagram of the GOA Node**

In the US, the GOT Nodes have digital return transmitters (DRTs) just like in the analog DAA node. However, unlike the analog DAA node, the GOT Node DRTs have Grey Optical SFPs, which, again are small-form pluggable modules, and a key reason these architectural advancements are even possible. This reduces the cost of the GOT node considerably, while enabling it to span the modest 0-2km links to traverse to the GOA Nodes. At the GOA Node, each of these DRTs are received in separate photo-diodes, not unlike the ones used in headends to receive the DWDM inputs for the Analog DAA system. However, the RF combining is considerably simpler, and these are then combined together with the RF from the GOA node port itself and presented directly to the RPD. In Analog headends, considerable critical infrastructure is given to RF combining. By minimizing and simplifying this, considerable power and space is conserved, and all 8 of the GOT nodes are thus combined in the GOA Node.

Because the RF combining in the US is done via the DRTs, the US system is exceptionally robust and high performing. Figure 7 shows a typical GOA node with two Quad GOA modules that help to provide a level of modularity while still gaining the benefits of aggregation.

Furthermore, a change in the PHY layer technology from the RPD to vCORE does not affect any of the GOT devices downstream of the GOA as long as the RF technology from the GOT node on down remains the same. This feature offers an element of flexibility in accommodating technology changes on the trunk route.

Finally, each of the GOT Nodes has a one-way element management system (EMS). This enables the entire set of critical information from the GOT nodes to be projected to the GOA node. The GOA Node itself has the I2C data bus enabled, so that the RPD is able to connect to and query all results from the GOA module, as well as measure and control the elements in the GOA node itself.


## GOA Performance

With the GOA architecture now clear, we can now look at its downstream and upstream performance.

## 7. Downstream

In the DS, the critical elements in the chain are the RPD, the analog GOA-GOT link, and the performance of the Node itself. The RPD MERs are generally very high, but it is also dependent somewhat on the RF load required of the RPD. Typically the MER of the RPD is in the vicinity of 48-50dB for the 1 to 1.2GHz load. The RF portion of the Node itself contributes to the MER, especially at 1.2GHz loading and its high DAA RF levels (73.6dBmV TCP); for 1GHz loading, the impact of the node is very modest at the DAA levels (69dBmV TCP).

The MER of the GOA-GOT link is dominated by shot noise and laser RIN, since the input power to the receiver is generally quite high and the fiber distance is quite small (<2km and around 0.5km on average). This MER is around 46-48dB for the 1-1.2GHz load, depending upon the optical input to the GOT receiver. For this reason, we expect little meaningful impact to performance due to the GOA-GOT link, which has been borne out by several tests as shown in Figure 9.



**Figure 9 – Measured MER values for RPD, GOA-GOT links and node**

Figure 9 shows the measured MER values of the entire chain of RPD, GOA-GOT links and the node itself. At input powers of around 0dBm, it is seen here that the performance of the GOA system is quite robust at 1GHz and at 1.2GHz (again, close attention to the TCP, unit-to-unit variations and temperature performance should be taken into account). In any case, the MER values have a healthy margin over the 4K OFDM QAM.

Figure 10 illustrates the measured RF output of an optical AGC receiver over the range of +3dBm to -4dBm and the Total Composite Power (TCP) of the RF output. It is seen that the receiver established good control over its output and held the node to its stated output power over very large optical input levels. In reality, the optical levels from node to node are much more modest, and this feature goes a long way to making the GOA architecture plug-and-play.






Figure 10 – Optical AGC in the DS Link

## 8. Upstream

In an earlier section, we explained how the upstream path of this system is based on DRT links. As is well known, the RF levels of the DRT link are independent of optical levels, as long as the link is functional. The optical noise power ratio (NPR) is independent of the optical links as well, again as long as the link is functional. Therefore one gets true plug-and-play link performance, as well as a high NPR, while using DRT links.




Figure 11 – GOA US Performance

Presented above is a measured NPR of GOA-GOT system. A typical DRT over an upstream mid-split (5-85MHz) would have an NPR curve with a peak NPR of ~54dB and a dynamic range ~23dB for 33.5dB, sufficient for OFDMA 1K Operation. When more DRT links are combined, there is a progressive reduction in the NPR values. With all 8 DRT links from the 8 GOT nodes combined, the NPR curve still has a peak of ~48dB and a dynamic range of ~15dB. Furthermore, at a typical 8dBmV RF input to the node, the NPR is ~40dB, has ~9dB protection from clipping, and a 6dB margin over the OFDMA 1K performance and could support more efficient higher capacity modulation.

**Table 1 – DRT RF Levels in the US**

GOT	GOA in	RPF Rfin
GOT 1	-4.0	7.8
GOT 2	-11.9	7.8
GOT 3	-15.0	7.8
GOT 4	-8.7	7.7
GOT 5	-3.7	-
GOT 6	-11.7	-
GOT 7	-14.3	-
GOT 8	-8.5	-
<b>Median</b>	<b>-10.2</b>	<b>7.8</b>
<b>Max</b>	<b>-3.7</b>	<b>7.8</b>
<b>Min</b>	<b>-15.0</b>	<b>7.7</b>

Several tests have been conducted over optical levels from -15dBm to -4dBm for various individual GOTs. Unsurprisingly, the RF levels and the NPR performance was completely unaffected. A complete system, built with GOA, multiple GOTs, Cable Modems (CMs) and traffic generators connected has been used to determine that the performance is just as solid over packet and frame losses, over the entire range of testing.

## 9. Additional GOA Options

GOA nodes with Analog Links instead of the DRT links were also evaluated. Constructing an Analog GOA was with an MDR (multi diode receiver), similar to the ones used in current day RFoG links to eliminate Optical Beat Interference (OBI) [4]. In our evaluations, the US links were at 1610nm; the MDR multiplexed the DS and US lightwaves and used just one fiber for the GOA-GOT links. Since the combining is only over 8 GOTs, instead of the traditional 32 encountered in RFoG links, the performance was quite good and matched that of the DRT links. In addition, the MDR-based approach offers wider BW, such as might be useful with a move to a high split or even FDX. However, and as expected, there is a lot of upstream RF level variation. Comcast has a very wide and diverse footprint and for this reason, we selected the DRT-based approach for our GOA architecture.

The GOA node was designed with only two RF blocks in the base. This is done to accommodate two GOA modules and the DS Tx and maintain the power envelope of the DAA node. Alternately, a single Quad GOA receiver module and 3 RF blocks could be configured in the base as long as the power envelope is maintained. This is particularly useful in moderate density environments where the additional RF block can cover more HHP and may be useful depending on RF coaxial cable routing.

Another variation that can be considered is to configure no RF blocks in the GOA node at all. We call this case the “Pure-GOA”, in that the GOA does not cover any HHPs at all and simply performs the aggregating function. While this will reduce the power dissipation of the GOA node and might be a good option to streamline design as well, it comes at a cost. One additional GOA will require an additional strand or pedestal location, and will require a GOT node nearby depending on existing hardline cable routing. In many cases, the cost and complexity of trying to arrange for two DAA sized nodes in close proximity is not justified and so, we have chosen the GOA approach with RF capability.

Fiber management and ergonomics of the GOA construction is an important topic. Certainly, optical splitters and fibers can be accommodated in the fiber tray of a DAA node. However, doing so would require opening the GOA node and risk performance impairments to the GOT constellation for fiber maintenance issues. For this reason, we have uniform sized splice enclosures for each of the GOA and GOT nodes. An added benefit of doing so is the ability to upgrade GOT nodes to RPD nodes with relative ease and is described later in the paper.

It is possible that GOA architecture can accommodate RF amplifiers in the plant. This might be especially useful in underground plant, however any decision to use RF amplifiers has the unfavorable side effect of dramatically increasing the number of field actives. In our analysis, we noticed a doubling of field actives for N+1 type architectures relative to N+0 type architectures. This result must be carefully weighed with known cost of maintaining the additional field actives and the cost and ease of field upgrades to accommodate traffic over the life of the plant and adoption of technology such as FDX.

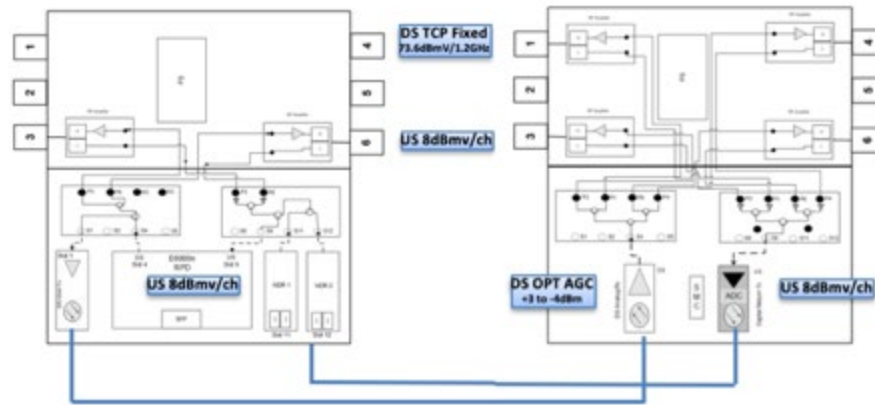
Finally, creating a headend architectural equivalent of the GOA can enable an easy 4x4 split of an existing node. Innovative options that utilize the essentials of GOA architecture to address legacy plant node splits while moving towards vCMTS and consolidating headend RF combining are worthy of further analysis.

## GOA Operations

As impressive as the performance of the GOA-GOT links are, additional considerations must be kept in mind, in terms of operations, monitoring and ergonomics. Since the optics in all the GOT nodes are Grey Optics, there is a benefit in their cost and simplicity in their procurement.

### 10. RF Levels across the Plant

Presented in Figure 12 is a schematic of one GOA, one GOT and all the RF levels in the system. The RF output levels of the GOA and GOT exactly match the RF levels of an RPD node or of an analog DAA node. As described previously, the combined effect of the GOA-GOT links with optical AGC ensures that all GOT nodes put out the same set RF levels and promote Plug-and-Play operation in the DS.




**Figure 12 – GOA Operations: RF Levels in the Plant**

For the Upstream, the RF levels for the mid-split continue to follow the set 8dBmV/Ch standard for both the GOA and GOT. The RF levels out of the GOA receiver modules are then combined in return combining boards and transformed so that a Unity Gain circuit is maintained from the GOT input to the RPD input, just as conceived in current RPD node standards.

The combined effect of the plug-and-play modules at the DS and US operations, and the modular GOA architecture, enables easy and simple upgrades in both even and uneven traffic growth conditions that will be discussed later in the paper.

## 11. Fiber Connectivity

Providing fiber connectivity from the GOA to GOT is depicted here. Typically there would be an optical passive with multiple ports (12 ports seems most reasonable, as typically more than one GOA exists within one Parent node domain) in a splice enclosure near to the GOA or a parent node location. The SFP in the RPD is connected to the SHE thru the splice enclosure passives.



**Figure 13 – Fiber Connectivity in the GOA architecture**

The RF level out of the RPD is distributed to the RF ports in the GOA node and to the DS transmitter. The output of the transmitter is then brought to the splice enclosure and split up to 8 ways; fibers are then connected to the GOT node. Each GOT node receives the same copy of the DS. In the US, the DRT fiber is connected to the splice enclosure and routed thru to the GOA module. Fiber entry service cables and a fiber connection standard to identify GOT nodes connected to GOA modules have been architected to enable these connections between the splice enclosure and the GOA node.

## 12. Ingress Monitoring and Control

While the GOA architecture offers great performance, it is at its heart an aggregating network. As with all aggregation points, there is opportunity for any one single element to pump in noise and impact the entire system. In real life, this could have an unfortunate effect: the field is unable to know which GOT was the offending node. As a consequence, a visit to the GOA would be required to open the node and fiber pull (pinch) to isolate ingress and identify the offending GOT node. Then the node would be closed, the ingress hunted down and mitigated; a return visit to the GOA is likely required to get the all clear.

To avoid these trouble shooting challenges, the GOA architecture provides powerful tools to remotely identify ingress effects, with the help of the monitoring and control design decision logic indicated in earlier sections, and described in the next section in more detail. Each GOA-GOT link can be either shut down or attenuated by 6dB with an RPD command. This level of remote control gives one the ability to isolate ingress or noise remotely and enable the operations technicians to identify the GOT node affected and avoid opening the GOA needlessly.

## 13. Monitoring the GOA and GOTs

Because there are many remote elements connected to the one RPD in the GOA node, there is a great desire to obtain, collate, curate data from all the GOT nodes. This is effectively done by the one way EMS typically offered by the DRT-GOA link and picked up by the GOA module which is connected to the RPD. The following table provides a few representative parameters that can be monitored.

**Table 2 - Example Table of monitorable and control parameters**

<b>GOA: Monitorable and Control Parameters</b>	<b>GOT: Monitorable Parameters</b>
<ul style="list-style-type: none"> <li>– General Module Information <ul style="list-style-type: none"> <li>• Model Number,</li> <li>• Serial Number</li> <li>• Node/FTX/GOA/</li> <li>• RFAs/PSU</li> </ul> </li> <li>– Downstream Transmitter <ul style="list-style-type: none"> <li>• Optical Tx Power (dBm)</li> <li>• Laser Temp (C)</li> </ul> </li> <li>– RF Amplifiers <ul style="list-style-type: none"> <li>• Amplifiers Bias</li> <li>• Ingress Control (On/-6dB/Off)</li> </ul> </li> <li>– Powering <ul style="list-style-type: none"> <li>• AC Input Voltage (V)</li> <li>• DC Output Voltage (V)</li> </ul> </li> <li>– Digital Return Receiver(s) <ul style="list-style-type: none"> <li>• Optical Input Power (dBm)</li> <li>• Ingress Control (On/-6dB/Off)</li> <li>• SFP Temperature (C)</li> <li>• GOA-SFP Model Numbers</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>– General Module Information <ul style="list-style-type: none"> <li>• Model Number,</li> <li>• Serial Number</li> <li>• Node/FRX/DRT/</li> <li>• RFAs/SMC/PSU</li> </ul> </li> <li>– Forward Receiver <ul style="list-style-type: none"> <li>• Optical Rx Power (dBm)</li> </ul> </li> <li>– RF Amplifiers <ul style="list-style-type: none"> <li>• Amplifiers Bias</li> </ul> </li> <li>– Powering <ul style="list-style-type: none"> <li>• AC Input Voltage (V)</li> <li>• DC Output Voltage (V)</li> </ul> </li> <li>– Digital Return Transmitter <ul style="list-style-type: none"> <li>• Optical Output Power (dBm)</li> <li>• SFP Temperature (C)</li> <li>• SFP Model Number</li> </ul> </li> </ul>

The ability to dashboard all this information can also help additional systems such as traffic analysis tools, to identify individual GOT nodes and understand their compositions for traffic allocation, which enables more effective capacity management in upgrade scenarios.

## Upgrade Options

Figure 14 shows nodes used in analog DAA, RPD DAA, the GOA and the GOT. The point here is that GOA and the RPD nodes have a similar power consumption profile, and Analog DAA node and the GOT nodes have a similar power consumption profile. In a system that has many more GOTs than the GOAs, the power consumption is always going to be less than that of an RPD-only deployment. However, systems are designed with the full RPD power budget in mind to enable future upgrades if and when required.

Also notice that converting a GOT node to an RPD node is as simple as taking out the DS receiver, the DRT and the SMC card and replacing it with an RPD. Note that since the RF levels were all mirrored, it



should be fully plug and play. These DS receivers and the DRTs are all relatively inexpensive Grey devices and can be reused elsewhere.


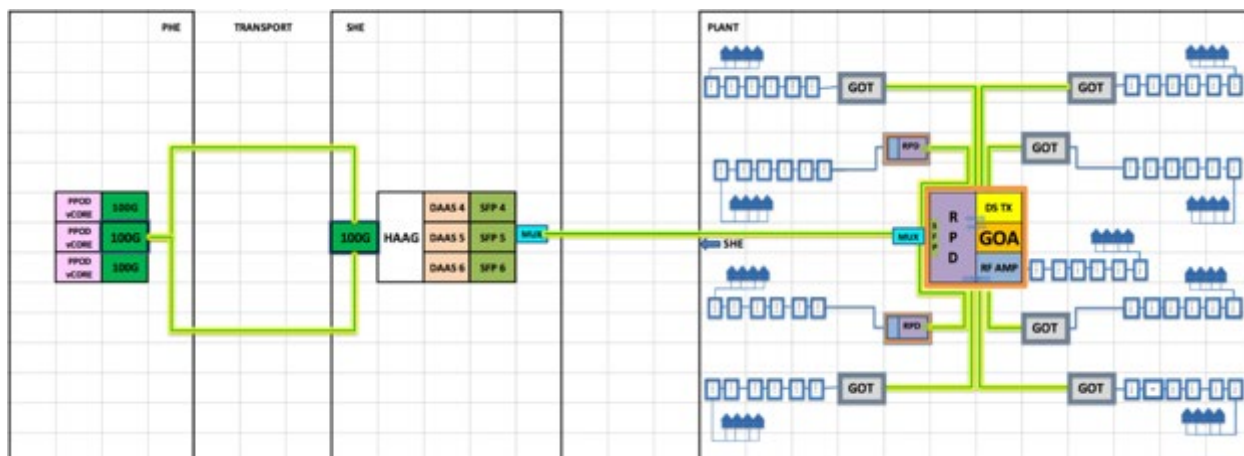


Figure 14 – RPD, GOA, Analog and GOT DAA Nodes in use

## 14. Uneven Traffic Growth

We begin with GOA architecture in a modestly dense location. However if one or possibly two GOT nodes experience huge traffic growth, such as might happen if an MDU were connected in a GOT node location, then that GOT node is converted to GOA. To do this the DS receiver, SCM and DRT are removed and replaced with an RPD. Then the fibers connecting the Receiver and DRT connect instead to the RPD. At the original GOA location, the fibers leaving the splitter for the GOT node and the incoming fibers going to the GOA are connected to the optical passives ports and linked back to the SHE. At the SHE, SFPs accepting the OSP are connected, a DAAS port is allocated and a vCMTS microservice is spun up in the PHE PPOD. Thus, deploying the compute and electronics just in time, taking advantage of Moore's Law improvements at the time the capacity is required.



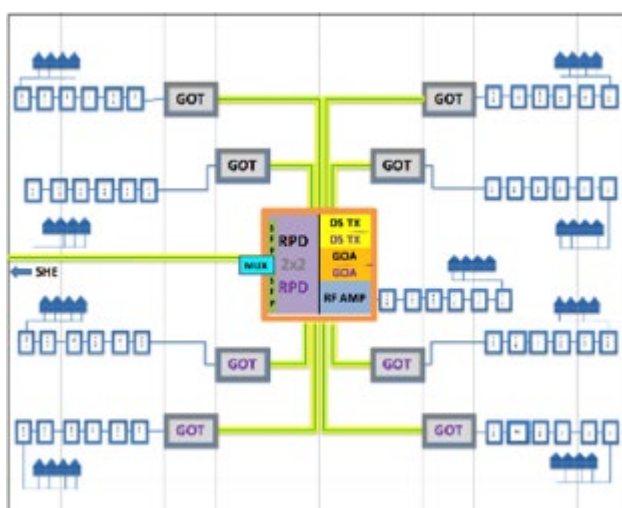
**Figure 15 – Upgrade path for uneven traffic growth**

It is also possible that a spare set of fibers from the GOA location to each GOT can be allocated and spliced during the original deployment, thus saving a visit to the GOA location when an upgrade is required. This is a time value cost and labor tradeoff, but if the original GOT locations are less likely to be upgraded, one may use the minimal cost approach and use the first revisit to the GOA system as an opportunity to add the additional optical equipment and invest in the labor to provide connectivity to the entire set of GOTs at that time.

It is important to note that the RPD that is deployed for additional traffic is success based. The intent is to protect initial capital outlay; the new RPDs have had the benefit of Moore's, Dennard's and Koomey's laws, collectively making them less expensive and more capable at their time of deployment.

## 15. Even Traffic Growth

Comcast has great visibility into keeping track of traffic on an RPD basis that extends all the way to GOTs. Thus the ability to perceive and predict creeping traffic increases is well understood.



**Figure 16 – Illustrating Even Traffic Growth**



When critical alarms appear, an elegant way of enhancing traffic overall would be to add an additional DS transmitter and disaggregate the two Quad GOA modules. Then a 2x2 RPD can be deployed with distinct connections to the DS transmitters and quad GOA modules to double the capacity for the entire GOA system. An important part of this type of upgrade is that technicians need not visit the GOT locations at all. A 2x2 RPD, when deployed, also fits in with the Moore's, Dennard's and Koomey's laws, and will likely be more efficient and cost effective when required.

## **16. Full Duplex in GOA Architecture**

The 10G initiative in the industry led by Comcast envisions the ability to use Full Duplex (FDX) type of transmission in the RF plant. In this architecture, the band of frequencies from 85MHz to 684Mhz can be used for both US and DS transmission at the same time on the RF plant. At the heart of this technology is Echo Cancellation of DS and US signals so as to preserve sufficient MER/SNR for adequate demodulation. Additionally, the concept of interference Group (IG) and transmission group (TG) are invoked in CMTS scheduling algorithms to further enhance overall capacity in the network.

A move to FDX is fairly easy in the GOA architecture once the underlying technology options have been standardized. In the simple case, each of the GOT nodes would be converted to FDX RPD nodes and the fiber links established as described in the previous sections.

There is however additional work in the industry to accommodate RF Amplifiers in an FDX environment. The idea is to distribute echo cancelling in the RF amplifiers and use appropriate IG and TG groups to maintain thruput capacity. In this approach, each of the GOT nodes could be treated similarly and each GOT enclave be part of an IG or TG. We note however that this approach is still under development as part of our FDX design activities.

## **Economics of the GOA**

When one looks at the make of the various nodes in discussion, it is easy to see that the GOA node may be more expensive than a regular RPD node due to the addition of the DS transmitter and the GOA aggregation receiver modules, even if a couple of the RF blocks are taken out. The RPD node, in turn, is more expensive than an analog DAA node, after accounting for the DWDM SFPs in both nodes and due to the more mature components used in the analog DAA node. The GOT node is less expensive than the analog DAA node since it uses Grey optics for the DRTs, which are generally less expensive than the DWDM counterparts and do not need optical DWDM passives to enable their transmission. Therefore the OSP cost is less for the GOA and GOT combination than a straight deployment of any other individual architecture.

From an inside plant perspective, the RPD architecture provides great benefits of virtualization and white box switching over the more conventional analog DAA architecture. Better monitoring and visualization of the network is also possible here. The RPD architecture also saves much needed critical infrastructure in the primary and secondary headends. Since the GOA architecture utilizes the same RPD plant -- but quite a bit less of it -- there is quite a bit of GOA savings in the inside plant as well, relative to all other architectures as shown when comparing Figures 5 and 6.


While it is true that the RPD architecture provides smaller service group sizes, the same virtue becomes vice in low and moderate densities.

## 17. Density and Aerial/Underground Plant

Outside plant is a complex entity; an adequate understanding requires years of experience. In every case, touching the plant requires permitting, construction, and traffic management, among other challenges. For this reason, DAA build guidelines prescribe expected RF levels, performances and practices that are conducive for the long term plant viability.

For purposes of this discussion, we classify plant in terms of density and composition. We selected 4 separate systems with varying densities and plant compositions. We then designed the system as RPD or GOA and compared attributes of the design.

The plant may have higher or lower density, and it is measured in terms of HHP/Mile or as HHP/Node. There is a good correlation between the two, as shown in Figure 16. However, a note of caution is in order for the HHP/Mile parameter. Sometimes this parameter can have very high values if measured in a predominantly rural area with one or two MDUs. In this case, it skews the numbers, since the relatively small mile footprint of the MDU sometimes gives can suggest that the plant is much more dense than that of the plant median. For this reason, the HHP/Node could be a better number, however even that is prone to design guideline changes that might come but infrequently. For example, dropping the RF levels of nodes could reduce HHP/Node, or reducing drop levels could increase the HHP/Node metric, but in either case the HHP/Mile number is unchanged.



**Figure 17 – Mapping Density in Cable Networks**


Figure 17 depicts the four systems with their HHP/Mile numbers drawn against the HHP/Node numbers. While the trend of higher HHP/Mile leading to higher HHP/Node is unmistakable, there is considerable variation. It is best to use multiple metrics to understand the OSP, as well as people who are skilled in the art, to initiate and complete designs that involve complex and interdependent parameters.

Like most service providers, Comcast's footprint is diverse, in terms of mixed aerial and underground plant. Aerial plant is both easier to build and considerably less expensive than the underground plant. As rule of thumb for DAA, in aerial plant construction labor cost and material costs are similar, but in case of underground, construction and labor is generally higher than material cost.

In Figure 17, for the two low density plants, one each was selected in a predominantly aerial area while the other was in a predominantly underground area. The same is the case with the two high density areas, where one each was in aerial predominant and underground predominant areas.

## 18. Economics and Discussion

Presented below is a graph of the total cost in arbitrary units of the design of various RPD and GOA designs in low density high aerial, high density low aerial, high density high aerial and low density low aerial areas. The design was first done using standard DAA guidelines. Note here that the construction costs are similar for RPD and GOA architectures. This is because both of them are high performing fiber deep nodes and have identical reach. The significant reduction in cost is due to the OSP and ISP material cost. This includes lower cost DWDM SFP and Passives, and fewer RPDs, lower cost GOT nodes, DAAS, transport elements and vCOREs. In this example we have not used the added benefits of Critical Infrastructure reductions and is a subject of evaluation internally.



**Figure 18 – Total cost of DAA for RPD and GOA architectures**

It is seen here that the cost of RPD and the GOA both decrease as the density increases and also as the plant is predominantly aerial. It is also seen that while the GOA is always less expensive than the RPD, its attractiveness grows at lower densities.

**Table 3 – Analyzing GOA Cost Reductions over RPD**

Plant Composition vs. Plant Density Matrix				
	Plant Composition Mostly UNDERGROUND (>85%)		Plant Composition Mostly AERIAL (>80%)	
LOW Plant Density <60HHP/Mile	Total GOA Cost -14%		Total GOA Cost -17%	
	GOA Material Cost -41%		GOA Material Cost -36%	
	GOA Construction Cost 0%		GOA Construction Cost 0%	
	Mostly UG Plant => Construction Costs High Low Density =>Material Costs High		Mostly AR Plant => Construction Costs Low Low Density => Material Costs High	
HIGH Plant Density >100 HHP/Mile	Total GOA Cost -6%		Total GOA Cost -15%	
	GOA Material Cost -27%		GOA Material Cost -29%	
	GOA Construction Cost 0%		GOA Construction Cost 0%	
	Mostly UG Plant => Construction Costs High High Density => Material Costs Low		Mostly UG Plant => Construction Costs Low High Density => Material Costs Low	

The summary quad chart above shows that in all quadrants, the benefit of GOA is manifest. However in three of the 4 quadrants there are double digit reductions in cost. This enables the deployment of fiber deeper into the network for the same capital investment by reducing material cost until it is required in the future gaining the benefits of the improvements in silicon technology. The combination of high Aerial plant with lower densities is especially suitable for GOA. In the underground plant with higher densities, construction cost dominates, which may make it advantageous to deploy RPDs on day one.

## GOA Trial

Even though the science, technology and economic trends are favorable, it is always best to run a limited trial to verify the numbers and to gain valuable experience that may help to further fine tune the architecture as we move forward. To this end, a trial in the late summer/early fall timeframe is planned. Our partners in the Central Division have selected an area that enables us to build and stress the GOA architecture in multiple ways, which will help us to further understand operational issues, verify performance numbers, do further economic analysis and provide firmer basis for future deployments.

### 19. Trial Area Selection and Basic Illustrative Rules

As part of the trial, the following basic apriori rules have been established as illustrated in the table below. If a node were to have more than 100HHP, then it will be constructed as an RPD node as the node density

justifies this option. Furthermore, we limit the number of HHP/GOA to no more than 256HHP with no more than 8GOTs per GOA. These limits will be refined as we proceed thru trail.

**Table 4 – Preliminary GOA Design Rules**

<p><b>Max HHP/GOT = 100 ~ 128</b></p> <ul style="list-style-type: none"> <li>• If HHP/GOT exceeds 100, consider putting in an RPD-DAA Node just for it</li> </ul> <p><b>Max HHP/GOA = 256 1.2/85 ~ 304 860/42</b></p> <ul style="list-style-type: none"> <li>• Recommend additional GOA once the HHP/GOA exceeds 256</li> </ul> <p><b>Max Subs/GOA = 128 ~ 152</b></p> <ul style="list-style-type: none"> <li>• Assuming average 50% Penetration</li> </ul> <p><b>Max CPE/GOA-RPD or vCMTS = 256 ~ 304</b></p> <ul style="list-style-type: none"> <li>• Assuming Up to 2 two-Way CPEs/Sub (XB6+X1, 45M/26M all over Comcast)</li> </ul> <p><b>Max GOT+GOA = 8 GOTs + 1 GOA</b></p> <ul style="list-style-type: none"> <li>• GOA system would have Max of 34 RF Ports</li> </ul>
---

Ideally, a trial of the GOA architecture would encompass moderate to low density areas, require multiple GOA nodes within one legacy node footprint, and enable GOA nodes to have the variety of GOT nodes per GOA node, including a full complement of 8 GOT nodes per GOA, and have a variety of HHP/GOA. Fortunately, the first trial of the GOA architecture attains to reality with precisely such attributes.

Figure 19 shows a geographical sketch of the trial site. Currently, a legacy node of 551 HHP is split up into 21 DAA nodes, averaging 26HHP/DAA node, for which fiber construction has been done. These 21 DAA nodes are now divided within 3 groups, each fed by a GOA node. The first GOA node has 240HHP and supports 8 GOT nodes, the second GOA supports 202HHP and 7 GOT nodes while the third GOA node supports 109HHP and 3 GOT nodes.

This spread of GOT nodes and HHP/GOA will give us a good opportunity to study various aspects of the architecture. As part of the trial, it is proposed to analyze and track a matrix of predicted vs. actual performance, operational issues related to deployment and maintenance, historical traffic for the legacy node vs. the actual in the GOA architecture, and ISP cost saving in respect of critical infrastructure.



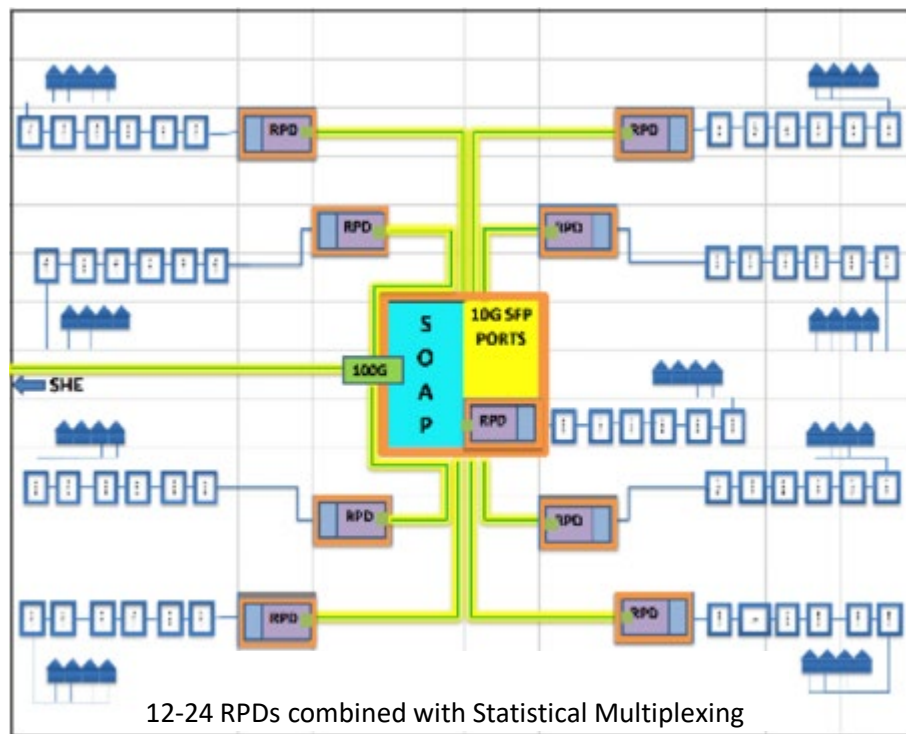
**Figure 19 – GOA Trial Area Description**

Over time, this will also be used to analyze ingress monitoring and mitigation as well as the process of upgrading individual GOT nodes to RPD nodes as traffic dictates. Furthermore, we can track the operational benefits of fiber predominant architecture and compare it to known trends of already deployed analog and RPD based DAA links.

## Switch On A Pole

The GOA roadmap culminates in a low powered, environmentally hardened ‘Switch On A Pole’ (SOAP) that multiplexes multiple 10Gbps grey Ethernet optics and leverages Coherent Optical links of 100Gbps and beyond to extend the headend into the outside plant as close to customer as possible. This allows the operator to pivot between or use multiple access technologies at the very end of the network easily such as Remote PON OLT or wireless nodes.





**Figure 20 – Illustrating the Switch on a Pole (SOAP) Architecture**

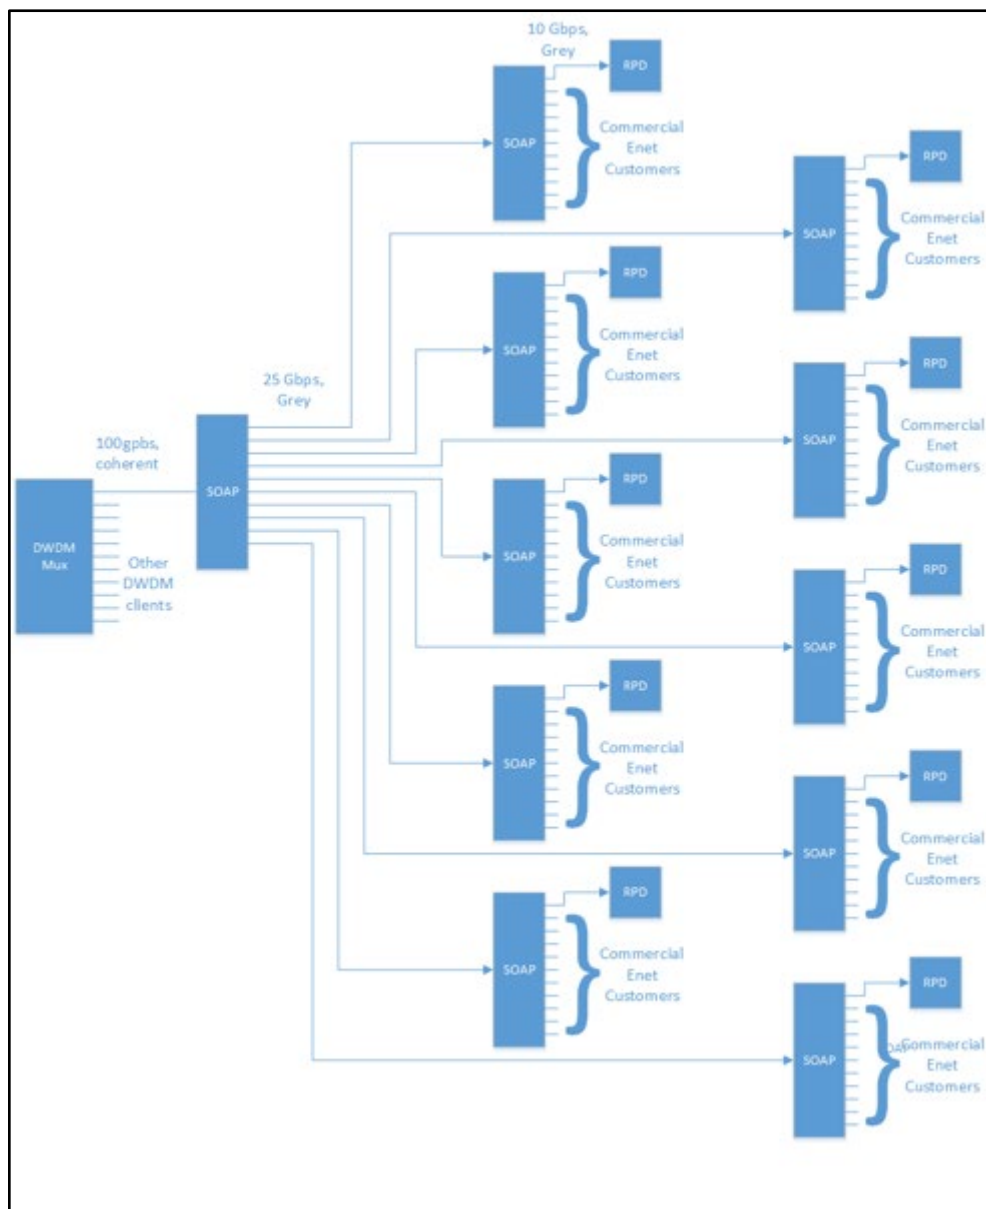
As previously mentioned, GOA is especially attractive in low to moderately dense areas. In higher density areas, after factoring in the labor and depreciation for the hardware, we will probably begin with RPDs instead of the GOA. There are still ways to optimize the costs of this solution using grey optics. One of the biggest “costs” of doing DWDM is the opportunity costs of the fiber and the wavelengths in the fiber. The fiber from SHE to node is one of the most valuable assets a cable operator owns and therefore maximizing its potential is paramount to long term shareholder value.

Cable operators have utilized statistical multiplexing, aka “over-subscription” techniques for decades. Rather than use individual wavelengths to address RPDs, a method of multiplexing at layer 2 is required and that is accomplished by pushing a switch deeper into the network using cost optimized commodity switching silicon solutions hardened for industrial temperatures.

The simplified picture above shows only 8 RPD nodes connected to a 100G link, however with the over-subscription model, ideally, the switch would support between 12 and 24 grey optical ports that would feed access nodes, each requiring a pair of fibers. In addition the switch would support a pluggable, dense, coherent capable DWDM wavelength that is 100Gbps (or more) capable of reaching lengths of 40-80km (or more). While this technology has been in existence for quite some time, meeting the low power and harsh temperature requirements of an access network node has been difficult. Recent innovations in opto-electronics are allowing for I-Temp capable 100G Coherent pluggables to be economically manufactured. Merchant silicon network chips have also matured and allowed multiple non-traditional network device manufactures to include switches in their products. In other words, node manufacturers can now use plug and play long haul, coherent optics, grey short haul optics, and a switch chip with their existing housings and power supplies.

By having an Ethernet switch deep in the network, cable operators can pivot between access network technologies quickly with little cost. In fact the switch allows the MSO to offer a multitude of services to

consumers in the area. A SOAP device could support low cost, remote OLTs, 5G radios, RPDs, and direct Ethernet switches in an oversubscribed mode. In other words, if only a few customers were using a new technology while the bulk were on another legacy technology, there would be no need to dedicate an entire wavelength to each service leaving one empty and one full. Eventually, the traffic will shift from one to the other. A grey optic in the switch allows for lower deployment costs of the new technology which allows for a faster innovation cycles.



**Figure 21 – Cascaded SOAP and the Converged Network**

SOAP devices can be cascaded without adding expensive optics or consuming multiple fibers and wavelengths, as is the current case with direct Ethernet services. For example, in order to accommodate 100 direct Ethernet customers today, 200 wavelengths across multiple fibers would be required from the SHE to the neighborhood. Also, multiple fiber splices and muxes would be required. Alternatively, 11 SOAPS could be deployed using one as a parent which aggregates 10 deeper in the network, which in turn



aggregate 10 customers a piece. The result would only consume 1 long haul optic, and 2 wavelengths and still support 20 RPDs. Combining residential and commercial services in such a way would allow for rapid commercial service growth without consuming very much optical capacity.

## Conclusion

In this paper we proposed a roadmap of ‘Grey Optical Aggregation’ (GOA) for the outside plant that will lower the cost of distributed access networks deployment -- while still providing the goodness of DAA performance and the potential for lowering operations costs. Just as importantly, the pivot to GOA enables a more effective use of trunk fibers from the SHE to the field, and conserves critical infrastructure at the headends, while driving fiber deeper into the plant.

We described several GOA options that demonstrate the flexibility of the architecture, as well as upgrade roadmaps for uniform and localized traffic growth. Economic analysis of the architecture in aerial and underground areas indicates savings across the board, but most especially in low to moderately dense areas. A field trial of the GOA in 2019 was described which will help in refining deployment and operational rules.

We described the evolution to an innovative SOAP architecture that leverages a low powered, environmentally hardened ‘Switch On A Pole’ (SOAP) that multiplexes multiple 10Gbps grey Ethernet optics and leverages Coherent Optical links of 100Gbps and beyond to extend the headend into the outside plant as close to customer as possible. An innovative cascaded-SOAP architecture allows the operator to pivot between or use multiple access technologies at the very end of the network easily.

## Acknowledgements

It is now our pleasure to acknowledge the help, support and encouragement of Tony Werner, Elad Nafshi and the senior leadership at Comcast. Tom Bach’s insights into operational aspects of GOA made this architecture more robust. Thomas Carroll of the Central Division has earned our gratitude with excellent selection of trial site and for his support. Finally we thank the folks at ARRIS/CommScope, especially Brent Arnold, for their support in the development of GOA technology.

## Abbreviations

AGC	Automatic Gain Control
AP	Access Point
BOL/EOL	Beginning of Life/End of Life
bps	Bits per Second
CAGR	Compound Annual Growth Rate
CMTS	Cable Modem Termination System
CWDM	Coarse Wave Division Multiplexing
DAA	Distributed Access Architecture
DAAS	Distributed Access Architecture Switch
DOCSIS	Data Over Cable Service Interface Specification
DRT	Digital Return Transmitter
DS	Downstream
DWDM	Dense Wave Division Multiplexing

EDFA	Erbium-Doped Fiber Amplifier
EIN	Equivalent Input Noise
EPON	Ethernet Passive Optical Network
ES	Extended Spectrum
FDX	Full Duplex DOCSIS
FTTH	Fiber To The Home
FEC	Forward Error Correction
GOA	Grey Optical Aggregation
GOT	Grey Optical Terminating Node
HFC	Hybrid Fiber-Coax
HHP	Households Passed
HD	High Definition
Hz	hertz
IG	Interference Group
ISBE	International Society of Broadband Experts
ISP	Inside Plant
MER	Modulation Error Ratio
MWL	Multiple Wavelength
NPR	Noise Power Ratio
OBI	Optical Beat Interference
OSP	Outside Plant
PON	Passive Optical Network
PHE	Primary Headend
PPOD	Primary Pods
RIN	Relative Intensity Noise
RPD	Remote PHY Device
SOAP	Switch On A Pole
SCTE	Society of Cable Telecommunications Engineers
SFP	Small Form Factor Pluggable
SHE	Secondary Headend
SRS	Stimulated Raman Scattering
SWL	Single Wavelength
TG	Transmission Group
US	Upstream
vCMTS	Virtual Cable Modem Termination System
vCORE	Virtual CMTS Core
XPM	Cross-Phase Modulation
4WM	4 Wave Mixing

## Bibliography & References

1. *Distributed Access Architecture – Goals and Methods of Virtualizing Cable Access*, Nagesh Nandiraju et. al., SCTE EXPO 2016
2. *Aboard the Technology Wave: Surf Report*, Rob Howald 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

# **Disaggregated, Coherent DWDM Solution at Shaw's Newest Cloud Datacentre Interconnect**

An Operational Practice prepared for SCTE•ISBE by

**Michael Ting Wang, P. Eng.**

Network Architect III

Shaw Communications Inc.

2728 Hopewell Place NE ▪ Calgary, AB. T1Y 7J7 ▪ Canada

403-303-4054

Michael.Wang@sjrb.ca

# Table of Contents

Title	Page Number
Table of Contents .....	2
1. Introduction.....	4
2. Legacy DWDM Systems .....	4
2.1. Monolithic Chassis.....	4
2.2. Fixed Optics On The Network Side .....	5
2.3. Complexity Of Planning And Provisioning Of New Services .....	5
2.4. Non-Coherent Fixed Grid.....	5
3. Modeling of The Cloud Data Centre Interconnect .....	6
3.1. Pre-defined Objectives For Modeling .....	6
3.2. Underlying Assumptions For Modeling .....	6
3.3. Modeling Results.....	6
3.3.1. Disaggregated, Pizza-Box Type, Modular Hardware Architecture .....	7
3.3.2. Small Form-Factor 100G/200G Pluggable Optics on the Network Side.....	8
3.3.3. Simplicity of Planning and Provisioning of New Services .....	10
3.3.4. Coherent Flexgrid .....	10
3.4. Analysis of Modeling Results .....	11
3.4.1. Ultra-High Bandwidth Capacity.....	11
3.4.2. Scalability and Adaptability.....	11
3.4.3. Efficient Use of Space and Power .....	12
3.4.4. Low Cost Per Gigabit of Bandwidth .....	12
3.4.5. Full Range Optical Power Auto-Adjustment .....	12
4. The DCI Solution And Its Advantages For All Operators .....	12
4.1. Physical Layout of A Standard Data Centre .....	12
4.2. The Cloud Data Centre Interconnect Solution .....	13
4.3. The Advantages For All Operators.....	14
4.3.1. Efficient Scaling .....	15
4.3.2. Optimal Utilization of Space and Power .....	15
4.3.3. Feature and Function Flexibility.....	15
4.3.4. Everything Pluggable for Transponders.....	15
4.3.5. Operational Simplicity .....	16
Conclusion .....	16
Abbreviations.....	17
Bibliography & References .....	17
Acknowledgements .....	17

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – A Typical Monolithic DWDM Device .....	5
Figure 2 – Aggregated/Monolithic VS Disaggregated/Modular.....	7
Figure 3 – A typical 1RU ROADM blade.....	8
Figure 4 – A typical 1RU Splitter/Coupler blade.....	8
Figure 5 – A typical 1RU Transponder Blade .....	8
Figure 6 - CFP2, Block Diagram .....	9
Figure 7 – Dimensions of CFP, CFP2, CFP4 .....	10
Figure 8 – Planning and Provisioning Process .....	10
Figure 9 – Standard Physical Layout of A Typical Data Centre.....	13
Figure 10 – DCI Architecture On Modular Hardware – Year One.....	13
Figure 11 – DCI Architecture On Modular Hardware – Three Years Later.....	14

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Legacy Paradigm of Non-Coherent Fixed Grid .....	6
Table 2 - New Paradigm of Coherent Flexgrid.....	11

# 1. Introduction

The phenomenal growth in Cloud services is driving the need for higher bandwidth within datacentres. To keep up with the global demand for higher bandwidth, datacentre operators need an interconnect solution that is dense, economically efficient and secure.

Shaw Communications is an MSO in western Canada that serves approximately two million Internet customers. Our newest LEED Gold certified Calgary Data Centre is a densely populated facility which offers hyperscale Cloud services. The opening of this state-of-the-art facility required an upgrade to the legacy low-density DWDM systems, which no longer met the requirements of Cloud Interconnect.

This paper outlines the work to create a modular coherent Cloud Data Centre Interconnect (DCI), which is technically agile, superbly scalable, operationally simple and economically efficient. The insights described in this paper into the DCI analysis and actions are intended to help other MSOs to achieve significant benefits when undergoing a similar Data Centre DWDM Interconnect modernization program.

## 2. Legacy DWDM Systems

Legacy DWDM technology has four major characteristics:

- It is based on monolithic chassis.
- It uses fixed optics with a maximum bandwidth of 10Gbit/s per client port.
- The planning and provisioning of new services are complex.
- It is non-coherent with fixed grid.


These characteristics will be discussed in detail below.

### 2.1. Monolithic Chassis

All manufacturers of legacy DWDM devices opted for monolithic chassis. Figure 1 below shows a typical monolithic DWDM device with the following dimensions (H x W x D) 22.5'' x 21.5'' x 11.5''. Most datacentre racks (45RU) have a height of 78.75''. The chassis in this example, with a height of 22.5'', occupies a significant portion of the rack. For its operation, the chassis requires interface cards in a wide variety of configurations and firmware with each card being separately managed by the chassis.

At the initial stage of a deployment, the chassis is usually only partially populated with cards. The monolithic chassis could represent a significant CAPEX expenditure that often is not fully utilized. In some cases, it might take a few years before a chassis can be fully populated. A large initial investment for such chassis simply denies precise financial planning.

After some time, once some chassis are fully populated. Network growth imposes the need for additional chassis, then the cycle starts again. The second chassis might be only able to be partially used for a long time. The chassis backplane capacity and features are physically fixed. There is no granularity for growth or new features. The paradigm of monolithic chassis fundamentally conflicts with agile DevOps concept.



**Figure 1 – A Typical Monolithic DWDM Device**

## **2.2. Fixed Optics On The Network Side**

Legacy DWDM line-side optics are fixed on the transceiver cards (transponders). Packing a 10G, 40G or 100G DWDM module into pluggable form is technically very challenging. DWDM transmission can involve multiple wavelengths that are relatively close to one another, the laser's wavelength must be tightly controlled to prevent it from wavering. Usually, wavelength lockers are needed in order to get wavelength stability and the traditional wavelength lockers are too bulky to fit in pluggable packages.

For a fixed optics card, a failed optic means a full card replacement. The process itself is wasteful. In addition, these cards tend to be bulky and have low port density.

## **2.3. Complexity Of Planning And Provisioning Of New Services**

The amplifiers in most legacy DWDM systems do not support auto-configuration and/or optical power auto-adjustment. Legacy DWDM systems often require the use of a modeling and simulation tool for their design. The inputs of such a design tool include the results of OTDR and the dispersion parameters of a fibre link. The output is configuration scripts or commands to configure and commission the system. It normally takes a long time to learn how to properly use such a tool. In many cases, minor changes to parameters in the resulting scripts or discrepancies between them and the hardware will prevent the system from turning up. Therefore, the turn-up process takes longer than required and often requires assistance from the vendor.

For a DWDM engineer of legacy systems, the use of planning and provisioning tools is cumbersome, and it requires extensive training.

## **2.4. Non-Coherent Fixed Grid**

Legacy DWDM technologies use non-coherent optics in a fixed grid with 100Ghz or 50Ghz spacing. As all major vendors are moving away from the paradigm/mode, it is unnecessary to discuss its technical details in depth. Instead, we will list features and the corresponding drawbacks of these obsolete paradigm in Table 1 below for reference purposes.

**Table 1 – Legacy Paradigm of Non-Coherent Fixed Grid**

Features	Drawbacks
On/Off Modulation; Direct Detection	Works well up to 10 Gbps. On/off modulation has low spectral efficiency at higher speeds. Very sensitive to optical propagation impairments. May need DCM which incurs higher latency.
Fixed Grid	Bandwidths cannot be adjusted flexibly. Lower spectrum efficiency when compared to flexible grids.

### 3. Modeling of The Cloud Data Centre Interconnect

The transport networks that relay information between data centres and to/from Internet Exchange Points (IXPs), are called “Data Centre Interconnect (DCI)”. DCI traffic predominantly flows across point-to-point connections, the technological platform for DCI is DWDM. The growing demand for cloud-based content delivery is putting pressure on the DWDM DCI. There is now an urgent need to boost the capacity and speed of the DWDM transport networks.

Data centres must use high-end DWDM optical network solutions for transporting multiple 10G, 40G, 100G, 200G services. At the same time, there is a heavy downward pressure on cost, power consumption, and the amount of physical space occupied by data centre equipment in general.

When the construction of the newest LEED Gold certified Calgary Data Centre was completed in March 2017, we set out to determine the best fitting and easily scalable DWDM paradigm. We decided to model out the possible solutions in a comprehensive study using growth patterns of our other data centres and evaluating multiple DWDM platforms.

#### 3.1. Pre-defined Objectives For Modeling

We determined the following key technical and business goals for modeling

1. Ultra-High Bandwidth Capacity
2. Scalability and Adaptability
3. Efficient Use of Space and Power
4. Low Cost Per Gigabit of Bandwidth
5. Full Range Optical Power Auto-Adjustment

#### 3.2. Underlying Assumptions For Modeling

Our modeling was based on several important assumptions. Each assumption is illustrated here:

- The anticipated traffic growth for a data centre was pegged at a compound growth rate of 49% per year per industry average (*ref. [1]*). Following the standards of the telecommunications industry, we use total cost per gigabit of bandwidth as a measure of the cost-effectiveness of a given solution.
- In line with industry standards, the DWDM DCI should only take 1.5% of the total available space and power of the data centre. 95% of space and power should be reserved for servers. The remaining 3.5% is reserved for related routers and switches. The servers are considered as revenue generating while DWDM, routers, and switches are considered as service-related expenses.


#### 3.3. Modeling Results

We will first present the results of the modeling without detailed analysis. In a latter part of this paper, we will carefully analyze the results. Essentially, the modeling results in a collection of mandatory DWDM product features and criteria for the selection of a vendor’s product line for the DWDM DCI.



### 3.3.1. Disaggregated, Pizza-Box Type, Modular Hardware Architecture

We need to choose a disaggregated DWDM platform for DCI. Disaggregation in optical networking allows operators to be much more flexible than they were with monolithic solutions. “Disaggregated Optical Transport” has two meanings in the industry. The first meaning is full-fledged white box platforms for ROADMs, for splitter/coupler functionality, and for coherent transponder. The service provider brings and/or builds its own software for the white boxes. The second meaning is the separation of ROADMs, splitter/couplers and transponders into stackable modules. Each functionality would be provided separately by individual modules and potentially from different vendors. For the purposes of this paper, the second is used. Figure 2 shows the concepts of aggregated/monolithic and disaggregated/modular.




**Figure 2 – Aggregated/Monolithic VS Disaggregated/Modular**

The monolithic chassis is shown on the left side of Figure 3, while the modular chassis is shown on the right side. The monolithic chassis must have a common copper backplane for control cards to manage line cards as well as transponders, this technical complexity results in high costs. The aggregated platform is space and power hungry and it is difficult to scale in different form factors. A monolithic chassis is the integration of multiple technologies advancing at different speeds, all of them anchored to the initial design constraints of the chassis.

The modular platform’s simplicity significantly reduces costs. Operations are also simplified compared to vertically integrated transport platforms using monolithic chassis. Small form factor means there is no wasted chassis capacity. Each module’s power usage is low, and the modularity allows for pay-as-you-grow. The monolithic chassis is very rigid for software/firmware upgrades, the entire chassis will either be upgraded or not upgraded. But for a modular platform, decisions whether to upgrade or not can be made for every individual module. If a feature is needed for a specific module, it can be upgraded. Otherwise, the upgrade can be postponed.


The disaggregated nature of modular, pizza-box type equipment allows the vendor to quickly innovate. The vendor will have shorter, more precisely targeted development and release cycles. This shorter innovation cycle and the constraints on rack space are the chief reason why it is no longer optimal to use a platform that packs multiple functions onto a single monolithic chassis. The new approach is to separate functional capabilities into 1RU standalone devices, known as “blades”. For example, based on the main building blocks of a DWDM system, there are 3 types of blades: ROADM blade (Reconfigurable Optical

Add/Drop Multiplexer), splitter/coupler blade and transponder blade. Figure 3 shows a typical 1RU ROADM module.




**Figure 3 – A typical 1RU ROADM blade**

The ROADM-on-a-blade provides wavelength selective switching and amplification. Our data centre requires up to eight ROADM blades (one main blade and up to seven tributary blades) that can be interconnected as single Network Element. Figure 4 shows a typical 1RU splitter/coupler module.



**Figure 4 – A typical 1RU Splitter/Coupler blade**

A splitter/coupler blade has the main function of channel add/drop. The splitter/coupler replaces legacy static filters in CDC (Colorless, Directionless, Contentionless) configurations. Figure 5 shows a typical 1RU transponder blade.



**Figure 5 – A typical 1RU Transponder Blade**

The transponder blade is a transceiver with very high port density on both network and client sides.

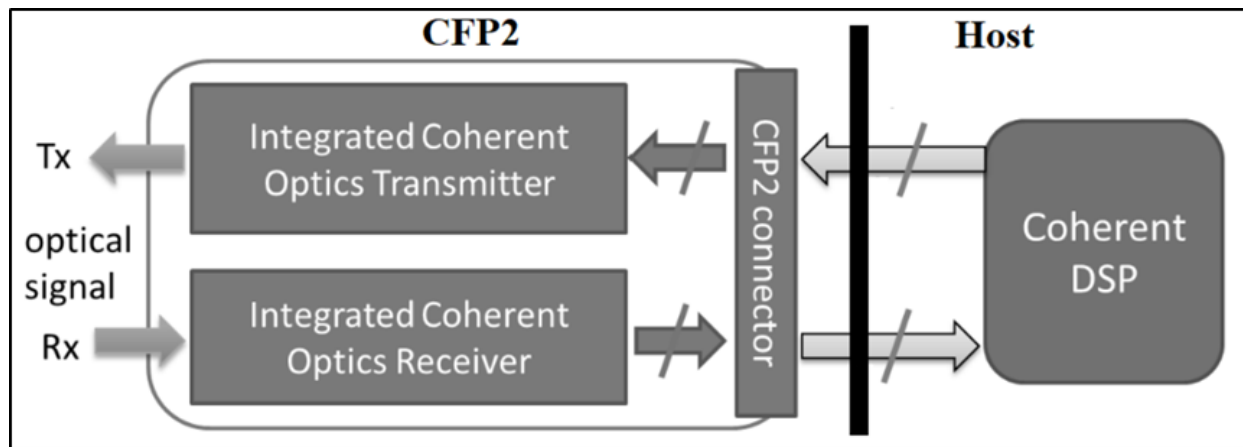
### **3.3.2. Small Form-Factor 100G/200G Pluggable Optics on the Network Side**

A key supporting ingredient to lowering 100G/200G costs for Data Centre Interconnect is the availability of pluggable transceivers. Most industry experts consider 100G/200G CFP2-DCO as well as CFP2-ACO transceivers a key enabling technology for smaller, lower cost solutions. Both types of pluggable technology have increased in importance in recent years. Without getting into too many technical details,

ACO is analog and first-generation technology while DCO is digital and more advanced. But ACO is acceptable to us. ACO's price is lower than that of DCO.


We have to select a product line with pluggable 100G/200G optics on the network side. The key value offered by pluggable optics is the ability to easily replace or repair a link without having to replace the entire transponder. Pluggable optics also provide the flexibility to use various physical media depending on fibre type and reach requirements. Additionally, the disaggregation of line side optics from the rest of the system enables upgrading optics to higher bandwidth (upgrade-as-you-grow) without large upfront CapEx. The pluggables also introduce the possibility of using third party optics.

Until recently, DWDM optics for 40Gb and 100Gb have been fully integrated to host systems due in part to signal integrity concerns between the optics and digital signal processor (DSP). With the standardization of electrical interfaces and advancements in coherent optics, the time is now right for pluggable 100G/200G DWDM applications. Figure 6 shows the block diagram of CFP2 and host. The CFP2 in the diagram can either be fixed or pluggable. It just shows general arrangement of the CFP2 in relation to host.



**Figure 6 - CFP2, Block Diagram**

The Optical Internetworking Forum (OIF) has standardized Analog Coherent Optics (ACO), as pluggable optics for 100G/200G DWDM. ACOs separate the DSP from the optics and provide an analog interface to the pluggable Coherent optics, offering the benefits of pluggable optics without compromising the system performance. Today, coherent interfaces with integrated DSP are available in CFP and CFP2 formats. Figure 7 shows the dimensions of CFP, CFP2 and CFP4 network-side 100G/200G pluggable optics.



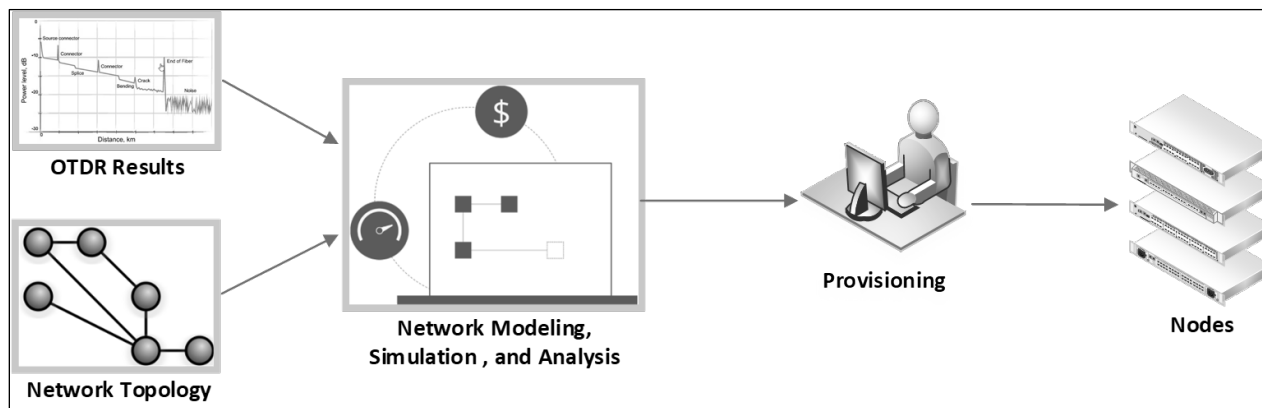
**Figure 7 – Dimensions of CFP, CFP2, CFP4**

Reprinted from medium.com (2017). "Aria Zhu, CFP Wiki: CFP /CFP2/CFP4 Transceiver Module Overview"  
<https://medium.com/@AriaZhu/cfp-wiki-cfp-cfp2-cfp4-transceiver-module-overview-6837f85e1797>

### 3.3.3. Simplicity of Planning and Provisioning of New Services

The availability, quality and necessity for a Design/Planning tool offered by DWDM equipment manufacturers was an important factor in platform selection for the DCI. While designing DWDM networks has always been a complex task because of the intrinsic nature of optical transmission, which is based on optical power level and non-linear effects of fibres. For the DWDM signal to travel longer distances, the optical power level must be sufficiently high. On the other hand, high optical power will trigger non-linear effects which will result in severe signal distortion. A designing/planning tool will find the optimal power required to ensure operation at the target distance while keeping the distortion in check.

We must select a product line that has the capability to adjust its power automatically based on its ability to detect fibre characteristics in the field and does not require the use of a design/planning tool. Figure 8 is the flow chart of planning and provisioning process.



**Figure 8 – Planning and Provisioning Process**

### 3.3.4. Coherent Flexgrid

The advanced modulation, coherent detection, and flexgrid are now the industry norm. Table 2 below lists the main features and benefits of coherent flexgrid.

**Table 2 - New Paradigm of Coherent Flexgrid**

<b>Features</b>	<b>Benefits</b>
Advanced Modulation; Coherent Detection; Digital Signal Processing; Soft-decision Forward Error Correction	High-bandwidth transmission. It has the required spectral efficiency with great noise tolerance. It compensates for optical impairments without requiring any regeneration or dispersion compensation of the signal on the link.
Flexgrid	It squeezes more channels into the existing fibre. It increases spectral efficiency and allows for narrower channel spacing within the existing C-Band.

### **3.4. Analysis of Modeling Results**

In this section, we will analyze the modeling results in the context of the pre-defined objectives listed in section 3.1. The entire analysis is also based on the assumptions discussed in section 3.2.

#### **3.4.1. Ultra-High Bandwidth Capacity**

Let's look at the first objective of section 3.1 – “Ultra-high Bandwidth Capacity”. Global Cloud data centre interconnection bandwidth will grow explosively, which outpaces Internet traffic significantly. Service providers are projected to grow their cloud interconnectivity capacity enormously as they work to provide additional and better cloud-native digital services.

The bandwidth criteria to choose modular blades are:

- 3RU rack space should be able to serve the first 100G wavelength. Each extra 1RU should be able to serve an extra 100G wavelength.
- The ROADM blade's port facing outside fibres must be able to support at least 100G per wavelength.
- The ROADM blade's total number of add/drop ports and passthru ports should not be less than 5.
- The transponder blade's line side must be able to support at least 100G
- The transponder blade's client side must be able to support at least 10x10G ports.

#### **3.4.2. Scalability and Adaptability**

We now move to the second objective of section 3.1 – “Scalability and Adaptability”. DWDM network scalability is the ability to easily add optical bandwidth. DWDM network adaptability is the ability to easily adapt to topology changes by adding or reducing degrees (directions). The initial traffic of a new data centre is usually quite low while the traffic growth will happen in spurs. In each spur, the growth rate can be shockingly high.

The scalability/flexibility criteria to choose modular blades are:

- When a transponder blade is added to the DWDM network element for bandwidth growth, it should not interrupt the existing traffic of the network element.
- When a ROADM blade is added or removed from the network element due to topology change, the traffic on other ROADM blades should not be impacted.
- Low cost of entry.
- Modular hardware platform that supports pay-as-you-grow model.

### **3.4.3. Efficient Use of Space and Power**

The next focus is the third objective of section 3.1 – “Efficient Use of Space and Power”. Most of a data centre’s space and power should be reserved for revenue generation servers and storage devices, not transport equipment. A typical datacentre only assigns 5% floor area for transport.

The space/power criteria to choose modular blades are:

- Per each 100G wavelength service, the minimum rack space saving should be 75% of the reference legacy system.
- Per each 100G wavelength service, the minimum power reduction should be 35% of the reference legacy system. Power savings for modular systems are not as significant as rack space savings, because digital signal processors (DSPs) of coherent receivers consume more power than their legacy counterparts.

### **3.4.4. Low Cost Per Gigabit of Bandwidth**

The fourth objective is “Low Cost Per Gigabit of Bandwidth”. The cost criterion is simple:

- The cost per gigabit of bandwidth should be lower than 50% of the reference legacy system

### **3.4.5. Full Range Optical Power Auto-Adjustment**

Let’s discuss the last objective of section 3.1 – “Full Range Optical Power Auto-Adjustment”. The optical power criteria to choose DWDM platform are:


- The DWDM platform needs to have the capacity to automatically adjust optical power based on its detection of fibre characteristics. Manual configuration of optical power should only be needed in the event of failures or during troubleshooting.
- The amplification stage should use built-in variable optical attenuators capable of full range adjustments to avoid the need for adding extra fixed attenuators.

## **4. The DCI Solution And Its Advantages For All Operators**

### **4.1. Physical Layout of A Standard Data Centre**

First, let’s briefly review the physical layout of a data centre. The Main Distribution Area (MDA) is defined as the location where traffic from all areas of the Data Centre converges. It is the hub of the cabling system and transport equipment. This is the central point of distribution for the data centre and every data centre has at least one. The Equipment Distribution Area (EDA) houses storage devices and application servers, this is where minimum 50% of data centre’s total floor area should be allocated to. An EDA is also called a “Data Hall”. The main distribution area may serve one or many equipment distribution areas (EDA) within the data centre. Usually, the Main Distribution Area has two rooms, MDA1 and MDA2. The Equipment Distribution Area contains multiple Data Halls.

The combined floor area of all Data Halls takes up the majority floor space for the data centre. The MDA only takes a small section of the data centre. All transport equipment must fit into MDA. The horizontal distribution area (HDA) may or may not be a separate area. In many cases, the HDA is part of EDA. Figure 9 shows a standard physical layout of a typical data centre.

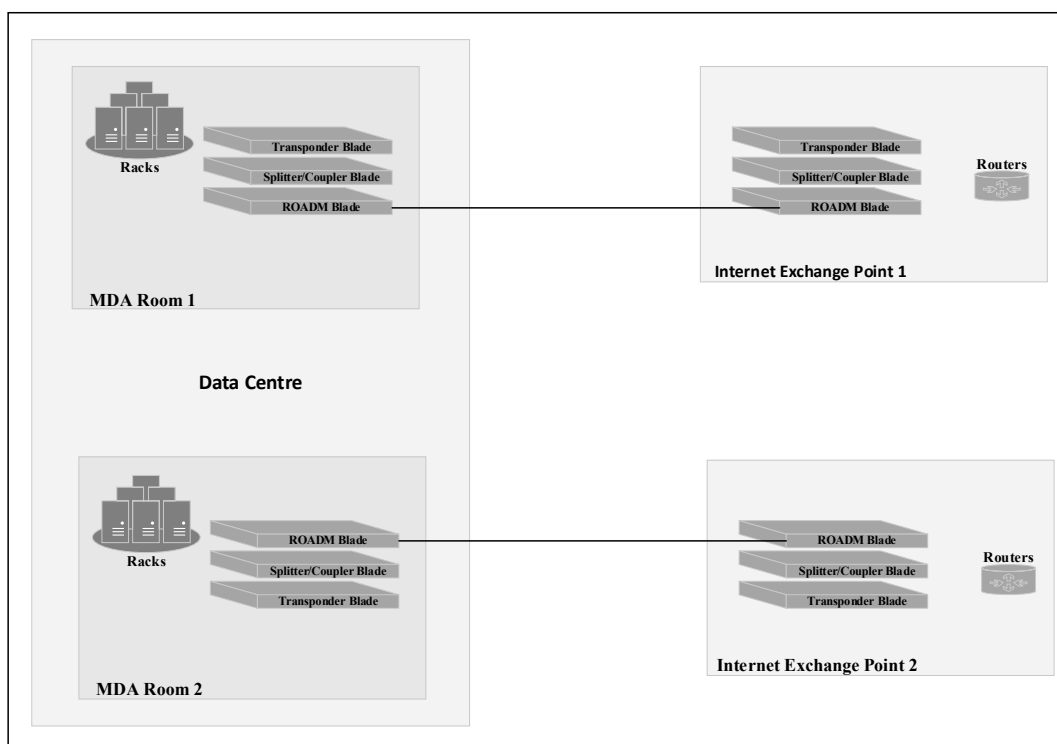


**Figure 9 – Standard Physical Layout of A Typical Data Centre.**

Reprinted from community.fs.com (2016). “Pre-terminated Cabling System for Data Centre Structured Cabling”  
<https://community.fs.com/blog/pre-terminated-solutions-for-data-centre-md-a-hda-and-eda-cabling.html>

## 4.2. The Cloud Data Centre Interconnect Solution

At the Campus Data Centre, there are two MDA (Main Distribution Area) rooms. Figure 10 below shows the current architecture of the DWDM interconnect. Each MDA room has its own set of one ROADM blade, one splitter/coupler blade, and one transponder blade. Inside each set, the three blades were connected to one another by inter-blade fibres. There are two types of inter-blade fibres. The first type is for managing multiple blades as one single Network Element. The second type is for user traffic.



**Figure 10 – DCI Architecture On Modular Hardware – Year One**

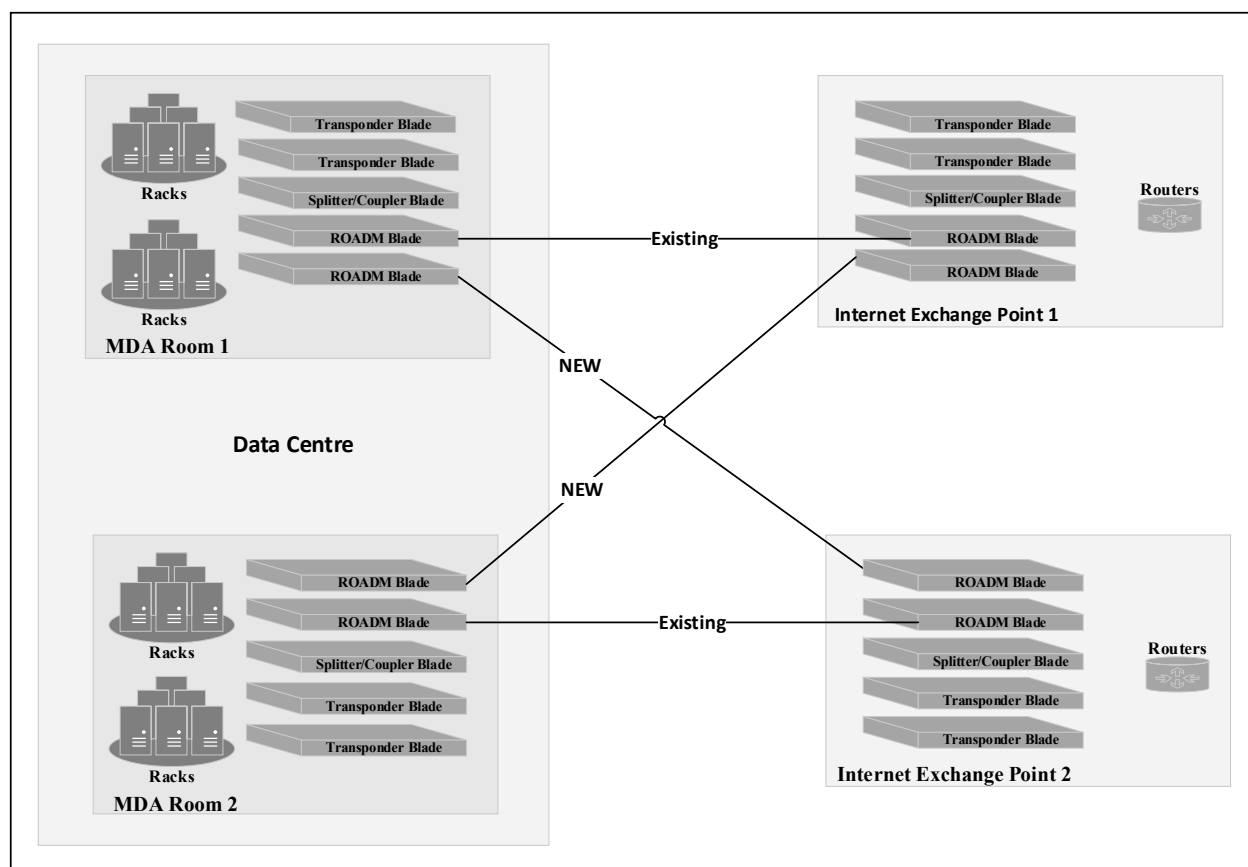
The Campus Data Centre is a brand new, state-of-the-art Cloud-service data centre. For the initial 3 years, the architecture above should be enough to meet the bandwidth demand. Each MDA room in the data centre

has one DWDM degree (direction) to an Internet Exchange Point (IXP). The data centre will have two connections to two different IXPs for redundancy. In each IXP site, the DWDM equipment set precisely mirrors that of the corresponding MDA room.

As it is shown above in Figure 9, there are no DWDM connections between the two MDA rooms. Inter-MDA connections should never be needed. If somehow the operators find that the inter-MDA traffic exists, it is a sure indicator that the data centre's design of storage/server placement is not optimal, and the design must be revised.

The Figure 11 below shows the future architecture of the DCI (Data Centre Interconnect) three years later. A new degree (direction) will be added in the form of a new ROADM blade in each MDA room. The new ROADM blade will be connected to the alternate Internet Exchange Point for protection. An additional transponder blade will also be added to each MDA room due to traffic growth. The additions of extra ROADM blades and transponder blades will not have any impact on the existing traffic.

The DWDM equipment set in an Internet Exchange Point does not have to mirror the MDA room. We have the flexibility to determine and deploy additional blades for each specific site based on the traffic demand.



**Figure 11 – DCI Architecture On Modular Hardware – Three Years Later**

### 4.3. The Advantages For All Operators

The successful implementation of a disaggregated Data Centre Interconnect has important positive implications and advantages for all operators. Traditional legacy DWDM systems have generally been



designed as vertically integrated, fixed systems. The vertical integration usually results in systems that have a large number of individual traffic and control cards for each of its various functions. These include traffic line cards, amplifiers, mux/demuxes, network management cards, and cards for dispersion compensation. Such systems are far too rigid.

The implementation of disaggregated solution has shown five main benefits over traditional monolithic systems that should be applicable to most operators. Operators should put DWDM disaggregation high on the agenda of cloud data centre interconnect.

#### **4.3.1.      *Efficient Scaling***

The disaggregation paradigm offers the advantage of efficient scaling for all operators. Scaling efficiently is one of the key network requirements for operators in the cloud era. The building block approach to hardware allows for a low initial spend for year-one deployments with the ability to grow incrementally as traffic increases and more capacity is required. Many converged, monolithic-chassis DWDM systems, by contrast, are able to handle future traffic volumes on day one, but also require a large up-front payment for that capacity even when the capacity may not be needed for several years. This is particularly true for chassis equipped with specialized hardware such as terabit scale switching fabrics, as the fabric is part of the initial installation, even though transponders may be added over time.

#### **4.3.2.      *Optimal Utilization of Space and Power***

The blade-centric solution delivers significant advantage for all operators in terms of optimal space and power management. In data centre environment, space and power are always at a premium. The blade-centric architecture eliminates rack partitioning completely. There is no need to find contiguous space, since the 1RU pizza-box blades are compact and can be distributed wherever there is space, including in different racks.

While DSPs do consume more power than their legacy direct detection counterparts, the blade-centric paradigm isolates the power-hunger to transponder blades. The vendors can concentrate on independent improvements of power without the constraints of traditional converged shelves. Most operators should be able to achieve an overall 30% power savings per 1G bandwidth comparing to legacy systems.

#### **4.3.3.      *Feature and Function Flexibility***

The modular product line provides the major advantage of feature/function flexibility for all operators. Disaggregation allows the vendors to rapidly implement and deploy features and functions when needed and in the amount that is needed. Key to this value proposition is the separation of functionalities into different hardware and the separation of the hardware development cycle from the software cycle. Because of the modularity and lower constraints of disaggregation, vendors tend to be much more responsive to feature/function requests.

Separating into modular blades from a shelf approach provides for independent blade improvements as functionality advances. Independent modular design facilitates agile development, continuous performance enhancement, and continuous technology innovation.

#### **4.3.4.      *Everything Pluggable for Transponders***

The coherent platform has the advantage of pluggable optics for transponders from network side to client side. This is significant for all operators. Until recently, the focus of pluggable optical module specifications and form-factors has been for client optics, where the predominant client specification is Ethernet.

Without exception, the DCI solution uses standard pluggable client optics such as SFP+ (10G) and QSFP28 (100G). Our Data Centre Interconnect shows that the pluggable line optics provide substantial benefits to operators. It gives operators the pay-as-you-grow benefit of only adding the expensive 100G coherent optics when traffic exists to support the additional expense.

#### **4.3.5. Operational Simplicity**

Operators should be able to take advantage of the operational simplicity of the DCI solution. The auto-adjustment platform offers operational simplicity by providing a high level of automation that facilitates network planning, engineering, configuration and deployment as well as accelerates the setup of end-to-end services. The intelligent software performs network optimization and accelerates the end-to-end provisioning/addition of wavelengths.

## **Conclusion**

This paper has reviewed the greenfield deployment of a modular Flexgrid DWDM system for Cloud Datacentre Interconnect. The advanced modulation techniques used, in conjunction with coherent detection and digital signal processing, has proven to be the most viable solution for Cloud Interconnect. Not only does it have the required spectral efficiency, it delivers increased Optical Signal-to-Noise Ratio and decreased Bit Error Rate by effectively compensating for fibre impairments.

The new DWDM system demonstrates superb scalability and density. It delivers 75% space savings and 30% reduction in power, on top of a 90% increase in available system bandwidth, in comparison to legacy DWDM systems. The system is also ready to support next generation modulations and automation to allow us to further scale it easily and economically.

Monolithic chassis are operationally cumbersome. A disaggregated, colourless (any wavelength) and directionless system can simplify operations. In addition, the system's ability to automatically adjust and optimize the line to the ever-changing characteristics of the fibre helped to dramatically simplify the design and maintenance of the Cloud interconnects. The new system results in excellent operational simplicity.

Monolithic chassis represent a significant initial capital outlay and the chassis backplane is inherently inflexible. On the contrary, a blade-based modular system requires a smaller, granular initial investment while allowing for a pay-as-you-grow approach. With single rack unit sized blades, the system is also space efficient and flexible. In general, it is very cost effective.

## Abbreviations

DWDM	Dense Wavelength-Division Multiplexing
DCI	Data Centre Interconnect
IXP	Internet Exchange Point
RU	Rack Unit
ROADM	Re-configurable Optical
OTDR	Optical Time-domain Reflectometer
DSP	digital signal processor
MDA	Main Distribution Area
EDA	Equipment Distribution Area
HDA	Horizontal Distribution Area
DCO	Digital Coherent Optics
ACO	Analog Coherent Optics

## Bibliography & References

- [1] Ian Redpath, “Global Data Centre Interconnect & Purpose-Built DCI Forecast Report: 2017–22”, *ovum.informa.com*, 2018
- [2] Harj Ghuman, “Coherent Access Applications for MSOs”, *SCTE•ISBE*, 2018.
- [3] Ian Betty, “Implementation Agreement for CFP2- Analogue Coherent Optics Module”, *OIF*, 2017
- [4] Heidi Adams, “Trends in Metro Optical Networks”, *IHS Markit*. 2017.
- [5] Peter Winterling, “Optical Modulation Methods”, *JDSU White Paper*, 2008.
- [6] Sterling Perrin, “Bring Disaggregation to Transport Networks”, *Heavy Reading*, 2015
- [7] James H. Chien, “Flex Coherent DWDM Transmission Framework Document”, *OIF*, 2017
- [8] Simon Stanley, “The Rise of PAM4 and 64QAM: A Competitive Analysis of Optical Modules & Components”, *Heavy Reading*, 2019.
- [9] L. Alberto Campos, “Proactive Network Maintenance Evolution to the Optical Domain in Coherent Optics”, *SCTE•ISBE*, 2018.

## Acknowledgements

*I would like to express my special thanks of gratitude to my VP Damian Poltz (Shaw Communications Inc.) who helped me a lot in finalizing this paper within the limited time frame, as well as my manager Felipe Arroyo (Shaw Communications Inc.) who gave me the golden opportunity to do this wonderful project on the topic and guided me through the project.*

# **Operational Transformation**

## **Modernizing Field Operations**

A Technical Paper prepared for SCTE•ISBE by

**Derek Strauss**

Director, Plant Operations  
Shaw Communications Inc.  
630, 3rd Avenue SW, Calgary Alberta  
+1-403-473-0004  
derek.strauss@sjrb.ca

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Modern Field Operations .....	3
1. Modernization Drivers .....	3
2. Key Metrics.....	4
3. Technology and Timing.....	5
4. Detailed Procedures.....	6
5. Results .....	11
6. Challenges and Opportunities .....	12
Conclusion .....	13
Abbreviations.....	13
Bibliography & References .....	14

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Shaw Field Operations Modernization Drivers.....	4
Figure 2 - Sample Program Metrics (Source: Shaw Communications).....	5
Figure 3 - Shaw Iterative Program Delivery.....	7
Figure 4 - Sample of multiple Shaw CAD-based HFC network record systems.....	7
Figure 5 - Shaw Predictive Network Maintenance coordination group.....	8
Figure 6 - Sample of Shaw work ticket.....	9
Figure 7 - Sample view of Shaw scheduler for field technicians.....	10
Figure 8 - Sample focus on change support and people first practices .....	11

# Introduction

Changing market demands require MSOs to deploy and operate smart, low-touch networks which leverage the constantly evolving tools, industry training and work processes. There is an increasing need for MSOs to modernize field resource activity to efficiently operate hybrid fibre-coax (HFC) networks and 4G LTE small cell deployments— while preparing field resources for emerging technologies such as distributed access architectures (DAA) and Full-Duplex DOCSIS (FDX).

This paper/presentation outlines the potential challenges and opportunities in modernizing field operations based on insights from Shaw Communications' recent integration and deployment of single-source geographic information systems (GIS), proactive network maintenance systems (PNM), a centralized workflow system, and a field service management system under one program.

Topics discussed will include efforts to support GIS data migration and upgrades, the process of PNM operationalization, amalgamation workflow activity into a singular system, the optimization of field management services, and essential operationalization and change management considerations.

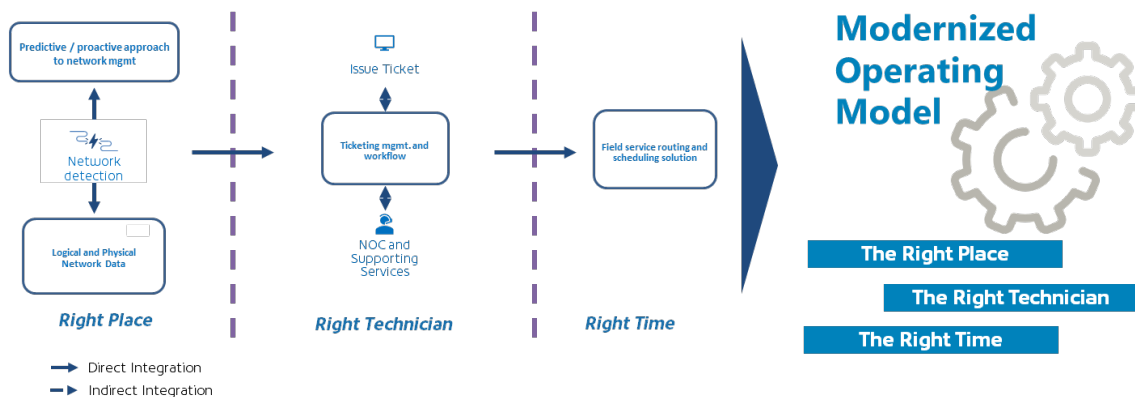
Examples of tool migration challenges, operational improvements and setbacks, training considerations, and field technician engagement will highlight the topics that must be addressed for the industry to modernize field operations.

## Modern Field Operations

A solid investment in field operations is instrumental in building a modern method of deploying networks and delivering services. Amalgamation of workflow systems and optimization of field management services can lead to significant operationalization and change management challenges, specifically around tools, training and engagement. To address these challenges, organizations should focus on the simplification of operational improvements and opportunities for modernization of field activity.

### 1. Modernization Drivers

Emerging technologies, software systems, resourcing changes and training needs have created opportunities to enhance the customer and employee experience. Consumers are now demanding the ability to self-serve and self-manage their installation and network connectivity experiences. To meet these demands, organizations must equip employees with simplified tools, streamlined processes and training on emerging technologies. Modernization of a converged workflow system enables a customer to receive support from the right technician, with the right equipment quickly, within a low-contact environment. Modernization also allows Operations to deliver operational efficiency programs across a wide variety of employee touch points, with a focus on employee engagement and simplicity.



**Figure 1 - Shaw Field Operations Modernization Drivers.**

## 2. Key Metrics

Delivering operational efficiency in Field Operations can be overwhelming, especially when you consider the multitude of factors that can impact day-to-day operations. Focusing on larger stacks of measurables allows for the long-term evaluation of modernization programs. As shown in Figure 2, Shaw Communications placed focus on Cost Efficiency, Work Quality and People. Under these larger stacks, smaller metrics were formulated based on long term focus areas supporting operational efficiency. The highlighted areas are key opportunities we are focusing our efforts towards as the program matures.

Operational reduction areas to be considered when integrating a modernization program include:

1. Technical Service Representative (TSR) contacts
2. Customer Service Representative (CSR) contacts
3. MTTR (mean time to repair)
4. Tech to tech referrals and a
5. Overall service calls and truck rolls.

Key metrics should also focus on:

1. Increase in Network Design efficiencies
2. Network/ Plant Technician efficiencies
3. In-home/In-Business Technician efficiencies
4. Increase in customer self-install/healing
5. Improved capital spend efficiency

## Sample | Program Metrics

	Metric	Description
<b>Cost Efficiency</b>	Cost per customer	Total ops spend / Total accounts
	Support Concentration	# of tech service related calls and tech service truck rolls per account
	Self Install %	% Self-install / Total Installs
	Self-Install Failure Rate	Failed / Total Self-installs
	Resourcing ratio	Ops employees (FTE) per 1,000 accounts
	Support interactions ratio	Total support interactions / Total accounts
	Truck roll ratio	Total truck roll / Total accounts
<b>Work Quality</b>	Service calls (30 days of install) ratio	Service truck rolls / Total accounts
<b>People</b>	Internal First - Operations	(Lateral moves + promotions) / Total Ops hires

**Figure 2 - Sample Program Metrics (Source: Shaw Communications)**

### 3. Technology and Timing

Converging technology platforms must be considered when modernizing field operations activity. Shaw Communications is focused on the integration and deployment of single source GIS, PNM, a centralized workflow system, and a field service management system under one operational program.

The first key technology deployment involved merging our legacy physical access (HFC) network records, with our logical fibre network within a geophysical information system (GIS) for our Operations teams to utilize in deploying and managing the (HFC) network. Our HFC records were managed across four systems requiring significant swivel chair activity to support the deployment and management of the HFC and Fibre Networks. Moving our HFC records into a single source of truth allows for the sunseting of multiple platforms and legacy swivel chair processes.

The second key technology deployment involved implementing a proactive network maintenance (PNM) system. Launching a PNM program enabled Shaw to move away from a legacy preventative maintenance (PM) program that required a significant level of labor and capital intensity to support HFC health and reliability. The new PNM program provides the flexibility to TSR, Network Operations Centre (NOC) and field staff to predict, and proactively correct, potential HFC network and in-home network failures while improving overall poor performance. Moving to a PNM technology provides the opportunity for Shaw to reduce multiple legacy tool platforms that field operations used to support HFC and in-home networks.

The third area of focus is a modern digital workflow system to move Field Operations work through the organization. This enables TSR'S and CSR 's to connect our customer journey with the activities of our field staff on the HFC network. The technology creates a single ticketing system with the ability to reference and link all field-related work activity to a customer experience. Allowing real-time communication between our NOC, TSR and CSR and digital experience teams and the customers they are



serving. Deploying a digital workflow system also allows Shaw to reduce and streamline hundreds of processes and administrative applications to manage work activity on the HFC Network.

The fourth area of technology is the enablement of an enterprise-wide field service management (FSM) software system. FSM technologies allow for the linkage of schedules, service routing, dispatch support, knowledge sharing, parts management and technician-to-customer communication. Advancements in FSM technologies allows field technicians and field operations support to have the most relevant information at their fingertips, while allowing customers to have insight to field technician arrival in real time.

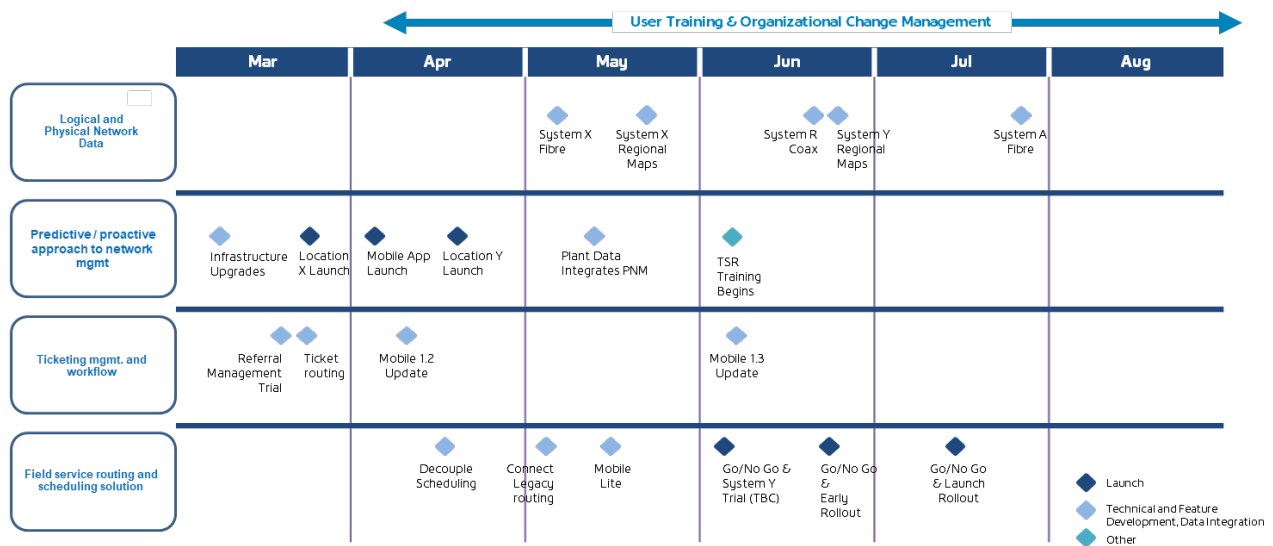
## **4. Detailed Procedures**

When converging technology platforms, we contemplated several critical levers to move our field operations team onto a condensed set of digital platforms. We knew it had to be simple for the field operations teams to absorb, operate and sustain. The program focused linking our GIS, PNM, Centralized digital workflow and FSM systems and Operational processes into a single field operations optimization program. This led to the development of a program with five unique work streams to support the following areas of modernization:

1. GIS Operations workstream
2. PNM Operations workstream
3. Centralized digital workflow workstream
4. Field Services Operations workstream
5. Operational Change Management workstream

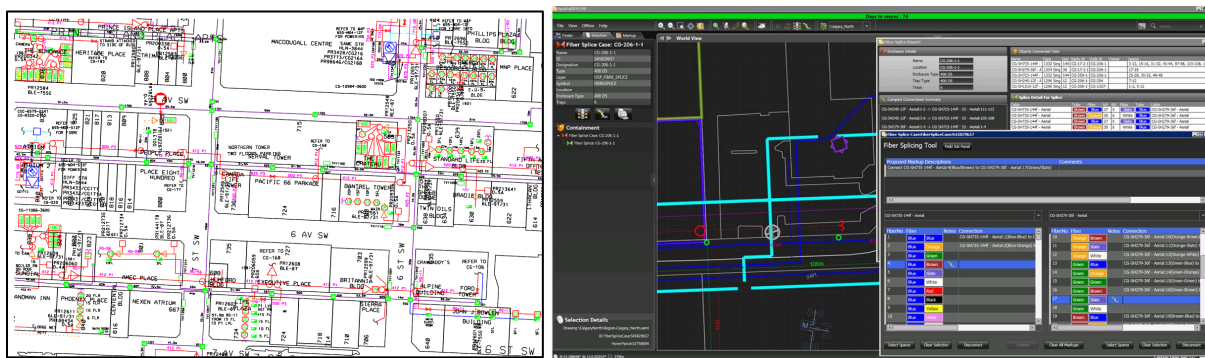
Operational Change Management (OCM) was introduced to support the significant changes and impacts resulting from the four digital platforms. OCM was implemented to support the organizational training, communications, change support and best practices around people – all required to support our field operations teams.

After the workstreams were connected, formalization of the program continued with budget development, key performance indicators (KPI) development, and program cadence. Technology and Operations teams were integrated to ensure alignment on the delivery schedule for each program. This collaborative model with four work streams allowed for information sharing, developer integrations, prioritization of requirements and preparation of operational change.



**Figure 3 - Shaw Iterative Program Delivery**

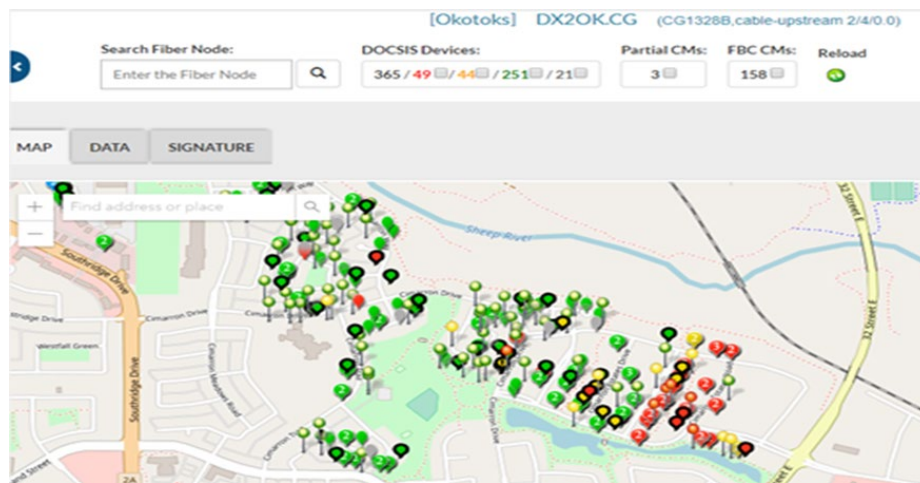
1. **GIS Operations workstream.** The mapping system has several data entry methods which require significant effort to manage physical network inventory. Two CAD-based drafting systems were used in tandem to support fiber design and HFC design, while an internal software suite supported fiber network management. We knew that merging these four systems into a single source of truth for the physical network inventory would be critical for supporting HFC design, management and health. Shaw has started the integration of the fiber network into a single CAD system. Alongside our partner integrator, we have tested the system's capability to migrate both the fiber record and the entire suite of HFC network records into a single CAD-based drafting system, while reducing labor requirements and long transition times. Preliminary results are favorable and show potential to migrate the physical network design and management into a single source of truth efficiently. The efficient migration sets up a key capability for our PNM Operations.



**Figure 4 - Sample of multiple Shaw CAD-based HFC network record systems**

2. **PNM Operations workstream.** Our Proactive Network Maintenance journey started with the operational requirement to leverage CableLabs technology for analyzing pre-equalization data in our cable modems to locate impairments in our coaxial networks. As the technology matured, demand for a PNM grew, especially for Shaw's Plant Operations Technicians, Network Operations Centers, Technical Service Representatives and In-home Technicians. The biggest

undertaking was operationalizing the technology for daily use. We designed customized training programs for each user group, leveraging what we knew about groups' abilities to triage the data on hand. Once the tool was integrated into daily Operations workflows, we began the process of overlaying our PNM data onto our GIS-based system for Operations to visualize correlation groups of modems compared to physical network assets. The collaboration of our GIS workstream and our PNM workstream is critical to achieve the targeted KPI's in MTTR and overall truck roll reductions.

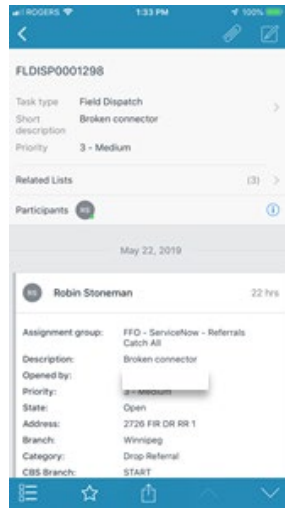


**Figure 5 - Shaw Predictive Network Maintenance correlation group**

3. **Digital Workflow and Ticket Management workstream.** A significant deep dive into workflow and process reviews was required to understand all work activity on the HFC network. To complete this, we had to adopt principles and best practices to improve operational efficiencies. Our operating ecosystem acknowledged multiple ticketing systems, process systems and work requests that did not connect activity on the HFC network to our customer experiences. The multiple workflows also compromised our ability to organize the dependencies of each system impacting our customer experience. We focused our efforts on moving all operational workflow onto a single ticketing system. The single ticket system has allowed us to prioritize activity and level of effort with overall technical integration and operational changes. The first area of focus was to provide our Plant Technicians the ability to open and close tickets automatically from the field when making changes on the HFC network. This reduced the need to call into our NOC for ticket management, and administrative assignment. The second area we focused on was creating the ability for the same Plant Technicians to self-assign tickets from our NOC, allowing the right technician with the right skillset to pick up the ticket and manage it through the process until closure based on their location and the tools and equipment they have at their disposal. The third area of focus is the operational sunseting of several internal ticketing systems and moving our field operations onto a single platform for internal referrals. Field referrals between technicians is a significant workflow and the focus is to reduce our referral touch points into a single, traceable work request. This function allows Shaw to correlate customer impacting field referrals into a trackable workflow for our CSR agents to provide insight through to our customers.


The key driver of the Ticket Management workstream is now focusing on moving all customer requests and operational requests into the same ticketing platform, connecting customer needs

and network management onto a digital platform, and routing through our field services management system.



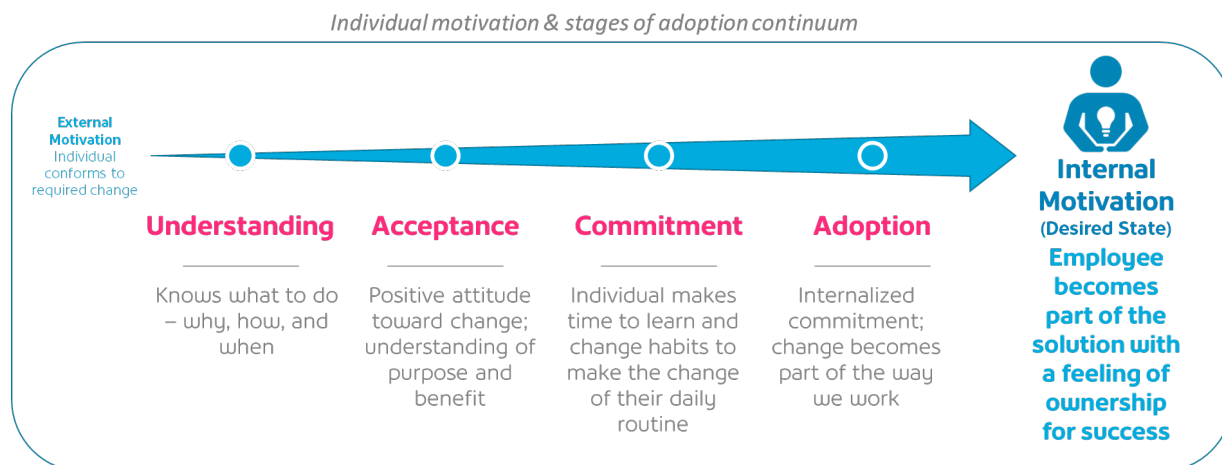
**Figure 6 - Sample of Shaw work ticket**

4. **Field Service Management workstream.** The ability to tactically move our field operations teams through the customer and plant ecosystem presented several unique operational challenges. Our in-home technicians were routed to customer appointments through older technologies, limiting our ability to support customer needs efficiently. Those limitation also extended to same day support from CSR, customer appointment notification, and the flexibility in customer changes. Network, plant and construction resources were supported through manual processes, with multiple tools being used to support employees, contractor movement and communication. Moving to a single Field Services platform allows the ability to support all field staff scheduling, workflow support, tech-to-tech communication, knowledge and inventory management. The platform allows for a new routing engine that uses traffic analytics, guiding the right technician to the right location as efficiently as possible. The integration work that is currently underway allows Shaw to move all tickets into our FSM system for routing or workforce management needs.



5. **Operational Change Management (OCM) workstream.** Significant focus was placed on the role operational change plays in the lift of tool sets and process changes. Linking four large programs together under one umbrella drives significant organizational alignment, communication, training and change support. When modernizing operations and tool sets, our initial focus was driven by the ability to operationalize the tool and ensure associated integration work between each tool set. As we matured the program integrations, we acknowledged quickly the need to frame up operational change support for employees. Our first leverage point was employee enrolment on tool selection. We enrolled field operations employees to participate in vendor selection committees. The participation in those committees allowed for a deeper understanding for the purpose of the tool, and how it would impact their day-to-day activity.

Our next step was to interview a cross-section of field teams for feedback on how this tool could improve their day-to-day activity and what process changes would be necessary to accommodate the tool set. This process helped build a foundation for initial adoption and drove field staff involvement and excitement surrounding the modernization of field operations. Instead of hiring traditional OCM resources, we focused internally within our field resource pool. Highly engaged team members from across our field operations footprint enrolled themselves to be our change management team. We divided the pool of resources across each workstream and these teams drove peer-to-peer OCM delivery for understanding, acceptance, commitment and adoption of each tool set and work process change (as shown in Figure 8).



**Figure 8 - Sample focus on change support and people first practices**

**Other Initiatives.** The program has recently absorbed the network monitoring program into the iterative delivery process. Adding network monitoring into the digital workflow, creates additional workflow improvements for technical teams when managing incidents and outages. A clear opportunity to reduce MTTR KPI's. The program also has a clear mandate to ensure all tool sets having the ability to be accessed from mobile devices, where field teams can easily access their work and associated tools remotely, with ease.

## 5. Results

Results of the modernization program are preliminary, however very promising. Connecting four very strategic programs has allowed us to start the preliminary measurement of KPI's as we move into our next fiscal cycle. Immediately, we see favorable results in technician engagement, technician efficiency, customer access improvements, and network troubleshooting efficiencies.

1. **Engagement.** The ability to leverage our people in decision making was critical outcome. Enrolling our people early in the program and having their input in vendor selection, tool integration, process changes, peer-to-peer communication, training development and operationalization has proven to be a winning formula for transformational modernization.
2. **Time Savings.** Early results indicate a cumulative time savings of up to 25% per week for our technical teams as we enable new technologies and their associated work processes. These initial time savings are driven from less time triaging customer issues, locating HFC issues, technician drive time, administrative activity, automated access to knowledge while on site, and the automation or elimination of legacy work processes. A simple change in ticket self-assignment saved over 16,000-man hours a year.
3. **Same day support.** The initial launch of our field services routing engine has allowed us to route the closest technician to meet customer needs in real time. The ability to support customer needs same day with the appropriate skillset allows for capacity efficiencies and the ability to reduce lost sales opportunities and reduce churn due to delayed technical support.
4. **HFC Troubleshooting.** Immediate ability to locate and correct HFC impairments on the wireline network and in the home. Our PNM suite is only in soft launch; however, the technical teams have begun to leverage its ability to sort through coloration groups and start correcting HFC impairments directly, with a significant reduction in troubleshooting time.

5. **Tool efficiency.** We are actively reducing our suite of tools for supporting customer experience and network connectivity. We are looking at reducing our operational tool needs by 50% over a three-year period. The reduction has a direct positive impact on our teams supporting the legacy tools and reduces potential operational costs to support tools field teams are not using.

The results mentioned above are the immediate wins within a new program. As we mature the modernization program, our field and technical teams see significant opportunities to layer in new deployment architecture and the potential operational efficiencies realized by the program.

## 6. Challenges and Opportunities

There are challenges when integrating four large projects under one modernization program. The most significant entails merging the project as an integration of technologies and workflow. We had started the integration of technologies early in the process; however, we were late on the operationalization of workflow needed for the field technician or the customer to realize the benefit of the technology transformation. Connecting technology teams and operations teams was challenging due to a barrier in understanding “technical jargon and operational process”. Once overcome, other areas of focus became critical for success.

1. **Program Integration.** Ensuring the projects within our modernization program had the integration visibility required to properly deliver each tool or process. It was imperative we had a clear line of site on what integrations were needed before committing to the deliverables. This activity could become challenging, and not having enough resource to support the integrations can slow the delivery of the program.
2. **Collision of Activity.** Collision of activity continues to be one of our larger focus areas. A program like modernization can have a significant amount of change associated with the program. Ensuring we don’t have change exhaustion or collision within the program or alongside other programs can be difficult, and we have had to shuffle deliverables to ensure experiences are kept as optimum as possible.
3. **Connecting Leaders.** Enrolling leaders at all levels can be overwhelming. Adding modernization programs to the existing workload of day-to-day operations should be examined very closely to ensure programs don’t get diminished in their importance. Encourage program leads to enroll leaders at all levels early in the program. This allows for the socialization with their teams earlier in the process and starts the change process very early.
4. **Communications and Training.** We paired communications and training together within Shaw’s modernization program. We found the challenges in delivering communications and training activity similar, due to the volume of information our technical teams would need access to. We have started using modern means to communicate program cadence and provide the ability for simple access to knowledge or training for each program. Allowing field operations to self-serve information and training on the program they are interested in provides an immediate human efficiency and unclogs traditional means of communications.

# Conclusion

A solid investment in field operations is instrumental in building a modern method of deploying networks and delivering services. Several areas within operations can benefit from process automation and modernized tool sets deployed in tandem with strong change support. At Shaw Communications, we focused our efforts on integrating the repository for physical network data, the PNM system, the ticketing system, and our field services management system into one workflow program. This has allowed us to be very clear on our ability to move work through field operations with a committed impact on our key metrics over a three-year plan.

Key metrics for field operations modernization:

1. Increase in Network Design
2. Network/ Plant Technician efficiencies
3. In-home/In-Business Technician efficiencies
4. Increase in customer self-install/healing
5. Improved capital spend efficiency

Through this modernization process we encountered integration and operational challenges which have impacted our ability to simplify activity as first described. The challenges in amalgamating workflow systems can be complex, specifically the change management needs and potential to disengage field staff. A solid partnership between OCM and any automation process is critical and should be treated as a key component of any modernization program.

The key ingredients for a successful modernization program include:

1. Engaging field operations and industry leaders to help define the tools and process they need to support positive customer experiences for the future;
2. Merging technical and operations teams together to work through integration and deployment challenges; and
3. Integrating OCM early into the modernization program.

Once the program is operationalized, field teams can quickly absorb the required architectures, technologies and customer requests with simple tools and processes.

# Abbreviations

DAA	distributed access architecture
FDX	full duplex docsis
GIS	Geographic information system
HFC	hybrid fiber-coax
PNM	proactive network maintenance
TSR	Technical Service Representative
CSR	Customer Service Representative
MTTR	mean time to repair
KPI	key performance indicators
FSM	field service management
OCM	operational change management



# **Bibliography & References**

All information sources and references are from Shaw Communications.

Cable lab doc CM-GL-PNMP-V02-110623 Proactive Network Maintenance Using Pre-equalization doc

# **DOCSIS® 4.0 Technology Realizing Multigigabit Symmetric Services**

## **Migration Scenarios for Multigigabit Return Services**

A Technical Paper prepared for SCTE•ISBE by

**Doug Jones**  
Principal Architect  
CableLabs  
858 Coal Creek Circle, Louisville, CO. 80027  
303.661.9100  
d.jones@cablelabs.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
Content .....	4
1. Coaxial Cable Spectrum.....	4
1.1. Forward (Downstream) Spectrum .....	4
1.2. Return (Upstream) Spectrum .....	5
1.3. Spectrum, Capacity and Speed .....	5
1.4. Traditional 750 MHz cable plant.....	6
1.5. 1.8 GHz, Return Spectrum and Symmetry .....	6
1.6. DOCSIS 4.0 .....	8
2. Enabling Technologies.....	8
2.1. Legacy Service and Spectrum Limits .....	8
2.1.1. Digital Video Set-top Box Spectrum.....	9
2.1.2. DOCSIS 2.0 and Earlier Spectrum .....	9
2.1.3. DOCSIS 3.0 Spectrum .....	9
2.1.4. DOCSIS 3.1 Spectrum .....	9
2.1.5. DOCSIS 4.0 Spectrum .....	9
2.2. Switched Digital Video .....	9
2.3. 750 MHz spectrum example .....	10
2.3.1. 750 MHz with SDV.....	10
2.3.2. 750 MHz without SDV .....	11
3. MultiGigabit Return Examples .....	12
3.1. 3 Gbps Return Example .....	12
3.1.1. 3 Gbps Return Example with SDV.....	12
3.1.2. 3 Gbps Return Example without SDV .....	13
3.2. 4 Gbps Return Example .....	13
3.2.1. 4 Gbps Return Example with SDV.....	14
3.2.2. 4 Gbps Return Example without SDV .....	14
3.3. Other Return Examples .....	15
4. Multigigabit Forward Example .....	15
Conclusion .....	16
Abbreviations.....	17
Bibliography & References .....	18

## List of Figures

Title	Page Number
Figure 1 – North American spectrum usage; 42 MHz return split and forward to 750 MHz .....	6
Figure 2 – Comparison of 750 MHz and 1.8 GHz HFC Networks .....	7
Figure 3 – 750 MHz plant spectrum allocation with SDV .....	11
Figure 4 – 750 MHz plant spectrum allocation without SDV.....	12
Figure 5 – 3 Gbps return with legacy spectrum allocation with SDV .....	12
Figure 6 – 3 Gbps return with legacy spectrum allocation without SDV.....	13
Figure 7 – 3 Gbps return with modified legacy spectrum allocation without SDV.....	13

Figure 8 – 4 Gbps with legacy spectrum allocation with SDV .....	14
Figure 9 – 4 Gbps with legacy spectrum allocation with SDV .....	14
Figure 10 – 4 Gbps return with legacy spectrum allocation without SDV .....	14
Figure 11 – Symmetric Multigigabit Service Example .....	15
Figure 12 – Symmetric Multigigabit Service Example .....	16
Figure 13 – DOCSIS 4.0 Technology Spectrum Plan .....	16

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Historical Cable Plant upper limits .....	4
Table 2 – Return splits for 1.8 GHz plant.....	7
Table 3 – Digital Service to Spectrum conversion .....	10
Table 4 – Downstream Spectrum Allocation to Digital Services with SDV .....	11
Table 5 – Downstream Spectrum Allocation to Digital Services without SDV .....	11

# Introduction

DOCSIS® 4.0 technology is being designed to support symmetric multigigabit services. In order to achieve these, the Hybrid fiber-coax (HFC) network will need to evolve to support both more return and more forward spectrum, all while maintaining existing services. To support more return speed, DOCSIS 4.0 technology will support expanded forward capacity on the HFC network to 1.8 GHz. Cable can already offer gigabit service in the downstream and providing upstream gigabit service will enable new symmetric multigigabit cable Internet products. However, there is still a set of legacy digital video services that need to be maintained for the foreseeable future. This paper will discuss options for maintaining legacy digital video services, while changing the HFC network to support symmetric multigigabit Internet services.

## Content

### 1. Coaxial Cable Spectrum

#### 1.1. Forward (Downstream) Spectrum

This first section explains some history of cable television, and how both the network and services have evolved over time. Readers that want to get right to multigigabit symmetric services can move to the next section.

Cable has been expanding the capacity of the coaxial network pretty much since it was invented. Originally it was a few analog TV channels. Then it was 20 channels, then 40, then 60, then 80. Then additional capacity for high definition television (HDTV). Then it was more capacity for both digital video, Internet and digital voice services. And when even more capacity was needed for expanding Internet service, systems started reclaiming analog television signals and converting them to digital television.

The cable system is generally referred to as the upper frequency supported. For example, it is common to refer to cable networks as 550 MHz, 750 MHz or even 1 GHz (1,002 MHz). Each of these upper limits has a place in cable history as shown in Table 1 below.

**Table 1 – Historical Cable Plant upper limits**

System limit	When state of the art	Why
220 MHz	1960	20 analog channels
330 MHz	1970	40 analog channels
450 MHz	1980	60 analog channels
550 MHz	1985	80 analog channels
750 MHz	1995	80 analog channels + room for digital services. Primarily intended to be digital television but soon also used for Internet service.
870 MHz	2000	Ditto, plus more high definition analog channels, digital television channels, and even more digital services.
1,002 MHz (a.k.a. 1 GHz)	2006	Ditto with room for lots of digital services including high definition television and data capacity.

There have been programs to reclaim spectrum. For example, the digital terminal adapter (DTA) was used to reclaim analog channels by converting them to digital channels, thereby reclaiming spectrum which could be used for other services, primarily Internet service. As a rule of thumb, one analog standard definition television channel would take 6 MHz of spectrum. In that same 6 MHz of spectrum it is possible to put up to 17 digital television channels using the Moving Picture Experts Group (MPEG-4) standard.

## **1.2. Return (Upstream) Spectrum**

The forward system limit is only part of the story. The coaxial cable is generally split into both forward and return spectrum; also known respectively as downstream and upstream. And whereas the forward capacity has been growing for decades, the return capacity has generally stayed the same. And this will be a key topic of this paper, to expand the upstream capacity of the cable plant.

North America generally ended the return path at 42 MHz in order to start the forward path at 54 MHz, which was the off-air location of channel 2. Europe ended the return path at 65 MHz, which allowed their forward to start at 86 MHz to carry the FM radio band.

With a nod to colleagues in the rest of the world, the rest of this paper will focus on North American cable plant. Even in the best cases in rest of world application where a 65 MHz or even an 85 MHz return path is in use, expanding the HFC network to 1.8 GHz will allow a return path to span up to 684 MHz, which is a 10x increase over the current standard return paths, and will allow for multigigabit return capacity as well as a multigigabit forward capacity.

## **1.3. Spectrum, Capacity and Speed**

The words spectrum, capacity and speed have specific meanings.

Spectrum, also sometimes referred to as bandwidth, is allocated on the coaxial cable. For example, on a North American 750 MHz system the forward spectrum runs from 54 MHz to 750 MHz, a total of 696 MHz of spectrum. The return path runs from 5 – 42 MHz for a total of 37 MHz of return spectrum. And here is a simple observation, that today's typical 750 MHz HFC network has 20 times the forward spectrum as return spectrum. This is because the early cable business was delivering television programming, and the more the better. However, with the business shifting toward both Internet service and IP service delivery, it is now a focus to grow both the forward and return to offer symmetric services. A 750 MHz HFC network can support a 1 Gbps downstream speed tier (and maintain legacy video services), however, the relatively small upstream spectrum supports a speed tier on the order of 40 – 50 mbps.

Spectrum is turned into bits by sending Sine waves over it. Coaxial cable is a fabulous medium to send sine waves over. Without going into the Physics behind it, at a certain time and amplitude and phase, a sine wave represents some number of digital bits which can be built up into IP packets and used for Internet services. Really, really fascinating stuff, but beyond the scope of this paper.

And the more spectrum there is, the more sine waves that can be sent and therefore more bits that can be carried. Simple as that. The more spectrum there is, the more Internet service capacity there can be.

Capacity is managed. Speed is an offered service. There has to be at least as much capacity as the highest speed tier. For example, in order to offer a 1 Gbps speed tier, there had better be at least 1 Gbps of

capacity. And since there are usually multiple customers sharing that capacity, there should be even more capacity available than the highest speed tier to ensure the customers are getting their advertised speed tiers. As a rule of thumb, a speed tier should be supported by twice the capacity. So, a 1 Gbps speed tier would be backed up by 2 Gbps of capacity. A 2 Gbps speed tier would be back by 4 Gbps of capacity, and so forth. This rule of thumb is generous, and actual numbers depend on the number of customers using that capacity, their usage, etc.

## 1.4. Traditional 750 MHz cable plant

A 750 MHz HFC network was state of the art in the late 1990s, and was used for plant upgrades through the mid 2000's. With digital television service, it supported a competitive triple-play offering that my home used for almost 2 decades. In North America, a 750 MHz HFC network is still the largest percentage of plant in use; decreasing for sure, but still the majority.

As shown in Figure 1, on a North American 750 MHz HFC network the return path runs from 5 – 42 MHz, and the forward from 54 – 750 MHz. Note this figure starts to introduce the new top end of 1.8 GHz (1,794 MHz), which is shown to scale for reference to allow the reader to start getting a feel for the new spectrum allocations discussed later in this paper.



**Figure 1 – North American spectrum usage; 42 MHz return split and forward to 750 MHz**

Figure 1 shows graphically how much return and forward spectrum is available; the color blue is used for return path spectrum and green is used for forward path spectrum. The issue has become the return path which is short on spectrum and hence can support only a minimal return speed and as can be seen, clearly the network is not symmetric in terms of capacity.

For decades this HFC network supported all the services that were needed including the move from analog TV to digital TV, the advent of high definition television, the introduction of Internet and digital phone services, etc. Now, it's time to look at something different.

## 1.5. 1.8 GHz, Return Spectrum and Symmetry

Why 1.8 GHz? There is nothing special about it per se, rather, about a decade ago this value was chosen when a company made HFC network equipment that could work to 1.8 GHz. For the detail-oriented reader, the actual number is 1,794 MHz and is referred to as 1.8 GHz.

With DOCSIS 3.1 technology, the top end of cable plant is 1.2 GHz (1,218 MHz), and the move to 1.8 GHz (1,794 MHz) allows an additional three downstream DOCSIS OFDM carriers, each 192 MHz wide.

The key reason for moving the top end of the spectrum to 1.8 GHz is not to increase the downstream capacity per se, rather, this change is to allow an increase in the return spectrum. In fact, a key tenet of DOCSIS 4.0 technology is to allow the HFC network to flexibly operate on any of six return path splits as listed in Table 2 below.

**Table 2 – Return splits for 1.8 GHz plant**

Return split end (MHz)	Diplex Region (MHz)	Forward spectrum start (MHz)	Forward spectrum available (MHz)	Ration of downstream to upstream capacity
204	54	258	1,536	7 : 1
300	78	378	1,416	5 : 1
396	102	498	1,296	3 : 1
492	120	612	1,182	2 : 1
588	150	738	1,056	1.7 : 1
684	174	858	936	1.3 : 1

Table 2 shows a couple of things. The first column shows the amount of return spectrum available. Since most current North American systems today support only 42 MHz of return spectrum, the move to support up to 684 MHz of spectrum is a greater than 10x increase in return capacity and is essentially the same amount of forward spectrum (696 MHz) available with a 750 MHz HFC network today.


The second column shows the amount of spectrum used by the diplex filter, which separates the return spectrum from the forward spectrum.

The fourth Column shows the amount of forward spectrum available with a 1.8 GHz HFC network, based on the return split, and should be compared with a 750 MHz HFC network that supports only 696 MHz of forward spectrum.

The fifth column shows the ratio of downstream to upstream capacity, and the trend to symmetry (1:1) as the total spectrum is partitioned for more upstream capacity. It is both the increase in spectrum and the trend toward symmetry which enables multigigabit symmetric data services and provides a new way to think about the HFC network.

A key item is that a DOCSIS 4.0 modem will be capable of all the splits listed in Table 2. This is very important because a DOCSIS 4.0 modem can be initially deployed when the network supports one return split, and later on the return split can be changed (increased to provide more return capacity and more symmetry) and that same DOCSIS 4.0 modem can stay in place and take advantage of the new return.

Figure 2 shows a comparison drawn to scale between the traditional 750 MHz HFC network (from Figure 1) and a 1.8 GHz HFC network using the highest return split from Table 2.



**Figure 2 – Comparison of 750 MHz and 1.8 GHz HFC Networks**



As can be seen in Figure 2 the 1.8 GHz HFC network has both significantly more return spectrum, and more forward spectrum. The goal is the migration from a broadcast television service, to more symmetric Internet services. The coaxial cable already out there is capable of this, though there will be operational issues to deal with when the switch occurs; however, making this switch to being predominantly an Internet services provider is the future of the business.

In summary, this next move to 1.8 GHz will allow symmetric multi-gigabit services. Internet services are the basis for the majority of communications these days, even the delivery of entertainment. What were once analog television channels are now carried over the digital technology. This conversion began in the late 1990s when cable began offering Internet Protocol (IP) services. Digital voice was originally offered as a legacy service too, before voice over IP (VoIP) was enabled. Now television services are going over IP too. Home automation, distance learning, health care, your favorite app's, etc., all these services are offered over IP.

While it is not possible to predict the services that will be used 5 years from now, it is a pretty good bet those services will be offered over IP. Hence this paper, discussing a method to increase IP capacity and move to symmetric multigigabit services, while maintaining the legacy video services.

## **1.6. DOCSIS 4.0**

To support multiple HFC network scenarios, DOCSIS 4.0 supports two types of plant operation:

- Full Duplex (FDX) operation using the full duplex spectrum defined between 108 MHz to 684 MHz.
- Traditional operation using separate forward and return spectrum separated by a diplex filter, with multiple return splits available.

In both cases, the splits are as listed in Table 2 and the intent is to add more upstream capacity for Internet service. Operating in either case depends on how the operator chooses to configure their HFC network for this future view of services.

## **2. Enabling Technologies**

The previous section discussed the increases in spectrum that are made available with a move to 1.8 GHz. This section looks at how the spectrum is currently being used for today's services including limits on those services imposed by existing network equipment and customer premise equipment (CPE). An important part of moving to 1.8 GHz will be ensuring that existing services still fit, based on how they fit into the available spectrum.

### **2.1. Legacy Service and Spectrum Limits**

When expanding the return path spectrum, the forward path spectrum has to be moved further up. That is, with a legacy 42 MHz return, the forward path started at 54 MHz. If the return path is move up to end at say 396 MHz, based on Table 2 the forward path now starts at 498 MHz. That means that all the downstream services that used to start at 54 MHz now need to start at 498 MHz, and the question becomes can the deployed modems and set-top boxes still receive those services?

The examples given in this paper assume legacy CPE can receive signals up to 1 GHz (1,002 MHz) as discussed in the following sections. Operators should check with their systems to learn the limits of deployed CPE.

### **2.1.1. Digital Video Set-top Box Spectrum**

Modern digital set-top boxes support spectrum to 1 GHz (1,002 MHz). Older set-top boxes will support lower limits, for example, both the Motorola DCT 2000 and Scientific-Atlanta Explorer 1850 set-top boxes operate to only 860 MHz. Operators should check their inventory of deployed set-top boxes and plan digital video spectrum appropriately.

If the return path is moved up, the digital video spectrum will also have to move up but cannot be moved beyond the limit of the deployed set-top boxes to receive that spectrum.

### **2.1.2. DOCSIS 2.0 and Earlier Spectrum**

DOCSIS 1.0, DOCSIS 1.1 and DOCSIS 2.0 modems only operate to 857 MHz. To keep these modems operational, a single downstream 6 MHz DOCSIS channel needs to be kept below 857 MHz.

### **2.1.3. DOCSIS 3.0 Spectrum**

DOCSIS 3.0 downstream technology in practice operates to 1 GHz (1,002 MHz) based on supplier implementations. Again, DOCSIS 3.0 spectrum will need to be kept below 1 GHz and will be seen in some cases to be a limiting factor.

### **2.1.4. DOCSIS 3.1 Spectrum**

DOCSIS 3.1 downstream technology is specified to operate to 1.2 GHz (1,218 MHz); hence, the existing DOCSIS 3.1 spectrum can be moved above 1 GHz but must stay below 1.2 GHz (1,218 MHz).

DOCSIS 3.1 modems can also use all the DOCSIS 3.0 spectrum. Because of that higher spectrum limit available for DOCSIS 3.1 modems, a program for migrating customers from DOCSIS 3.0 to DOCSIS 3.1 modems is a positive step when considering a move to 1.8 GHz.

### **2.1.5. DOCSIS 4.0 Spectrum**

DOCSIS 4.0 modems are anticipated to support all the return splits listed in Table 2 (up to 684 MHz) while also supporting downstream spectrum up to 1.8 GHz (1,794 MHz) which will make the DOCSIS 4.0 modem the first to be capable of symmetric multigigabit services.

DOCSIS 4.0 modems can also use all of both the DOCSIS 3.0 and DOCSIS 3.1 spectrum.

## **2.2. Switched Digital Video**

Switched Digital Video (SDV) provides a method to reclaim spectrum for video delivered using MPEG transport. This includes both linear television programming as well as video on demand (VOD) programming. SDV delivers programs only when and where requested by viewers, unlike broadcast digital video systems that deliver all programming all the time even if no viewer is watching.

With broadcast digital video, HFC capacity is consumed even if no one is watching. With SDV, HFC capacity is only consumed when a viewer is watching. SDV optimizes capacity and without dedicating spectrum to programming that is not being watched.

Additionally, by switching some, or all, of the broadcast video tier it is possible to significantly increase the amount of programming offered. The so-called “long tail effect” demonstrates that there is a large aggregate demand when many specialized offerings are made available.

It should be noted that both IPTV and over-the-top video providers are doing the same thing, using IP transport instead of MPEG transport. Imagine if Netflix broadcast their entire programming catalog all the time to all subscribers and allowed the subscriber to just tune to whichever “channel” they wanted. This would be a huge waste of capacity. Rather, Netflix content is only placed on the network when there is an active viewer consuming that content.

SDV and IPTV are both examples of a switched model where spectrum is only allocated when content is being watched. Several references are provided about the statistics of SDV and the evidence is clear it conserves HFC network capacity. This information will be used when providing examples of migrating existing services to HFC network with a higher return split.

### 2.3. 750 MHz spectrum example

This section will show an example of how the spectrum is used to support today’s services. A 750 MHz HFC network was chosen because this is the most prevalent in North America. These examples can be considered representative of a 750 MHz HFC network and different operators can do different things.

On a 750 MHz HFC network, the forward path starts at 54 MHz and ends at 750 MHz which provides a total of 696 MHz of spectrum to support all the forward path services offered today, including television, Internet and all the IP services that run over the DOCSIS spectrum.

This example assumes all services are digital, and that analog television has been removed. Digital television uses the MPEG-4 standard which allows compressing more television channels into a given amount of spectrum than for example the MPEG-2 standard. Table 3 lists conversion factors used to convert services into spectrum.

**Table 3 – Digital Service to Spectrum conversion**

Digital Service	Spectrum
17 standard definition digital television channels (MPEG-4)	6 MHz
9 high definition digital television channels (MPEG-4)	6 MHz
1 DOCSIS downstream channel capable of 38 mbps total capacity	6 MHz
DOCSIS 3.1	Variable

In North America the downstream channelization was traditionally built on 6 MHz of spectrum which is based on the NTSC analog television standard. (Other regions of the world use the PAL standard which uses 8 MHz of spectrum per analog channel, and these regions can also use 8 MHz DOCSIS channels.)

Digital technology allows more service to be carried in the spectrum than analog. Digital has been used to make the most of the available spectrum by compressing more services into it and is the reason that such devices as the DTA have been used to carry televisions in digital format on the cable plant.

#### 2.3.1. 750 MHz with SDV


Table 4 shows an example spectrum allocation for a 750 MHz HFC network and how the forward spectrum is used. There is a total of 115 Consumer Technology Association (CTA) channels available on this plant, each 6 MHz. With SDV in use, all the digital television services can be offered in 40 CTA

channels, using 240 MHz of spectrum. The remainder of the spectrum is allocated for DOCSIS service, including both DOCSIS 3.0 (and earlier) and DOCSIS 3.1 spectrum.

**Table 4 – Downstream Spectrum Allocation to Digital Services with SDV**

Service	CTA channels (@ 6 MHz)	Spectrum
Digital Television (SD and HD)	40	240 MHz
DOCSIS 3.0 (and earlier)	32	192 MHz
DOCSIS 3.1	32	192 MHz
Total	104	624 MHz

The spectrum listed in Table 4, with SDV in use, is shown in Figure 3.



**Figure 3 – 750 MHz plant spectrum allocation with SDV**

Note that all the services fit within the forward path that starts at 54 MHz and ends at 750 MHz. An astute reader will realize that 696 MHz of spectrum divided into 6 MHz channels works out to 116 channels, however, just below the FM band there is a discontinuity in the channelization resulting in 115 CTA channels.


### 2.3.2. 750 MHz without SDV

Table 5 shows the corresponding spectrum usage for a plant without SDV. Note that without SDV, more CTA channels are needed for the digital video services including SD channels, HD channels and VOD service. Specifically with SDV as shown in Table 4, a total of 40 CTA channels are needed for digital television whereas Table 5 shows that without SDV a total of 60 CTA channels are needed for digital television. But there are only 115 total CTA channels available on a 750 MHz HFC network, therefore with more CTA channels allocated to video there will be fewer CTA channels available for DOCSIS spectrum. Note that Table 5 shows half the DOCSIS 3.1 spectrum that is available in Table 4. On the cable plant there is only so much spectrum to be allocated between services, and this is an example of traditional broadcast digital television without SDV requiring more capacity and hence less capacity is available for DOCSIS service.

**Table 5 – Downstream Spectrum Allocation to Digital Services without SDV**

Service	CTA channels (@ 6 MHz)	Spectrum
SD digital television (425 channels)	25	96 MHz
HD Digital television (207 channels)	23	192 MHz
VOD	12	72 MHz
DOCSIS 3.0 (and earlier)	32	192 MHz
DOCSIS 3.1	16	96 MHz
Total	108	648 MHz

The spectrum listed in Table 5 is shown in Figure 4.



**Figure 4 – 750 MHz plant spectrum allocation without SDV**

In both Figure 3 and Figure 4, all the services fit within the 750 MHz limit of a traditional cable plant. A North American 750 MHz HFC network supports 115 forward CTA channels (each 6 MHz wide) and the example with SDV uses 104 CTA channels whereas the example without SDV uses 108 CTA channels (albiet with less DOCSIS capacity). No system fills every channel and it is appropriate to not show every CTA channel being used.

### 3. MultiGigabit Return Examples

A key point of this paper is that an operator can take their 750 MHz HFC network today, and without any changes to existing services, migrate the spectrum to a 1.8 GHz HFC network with more return path and provide symmetric multigigabit data services. And do this without running into legacy CPE limits as discussed in section 2.1. This is not straightforward since most legacy CPE operate to only 1 GHz (1,002 MHz), and increasing the return path means that the start and end of the forward path is also moved up.

In this section, information from the proceeding sections will be synthesized to show how this works.

This section focuses on the return path and a forward path example using a 1.8 GHz HFC network is given in section 4.


#### 3.1. 3 Gbps Return Example

A 396 MHz return will support up to 3 Gbps of upstream capacity using DOCSIS 4.0 technology.

With a 396 MHz return, the forward path starts at 498 MHz which means the legacy services which used to start 54 MHz will now start at 498 MHz, and end correspondingly higher in the spectrum too.

##### 3.1.1. 3 Gbps Return Example with SDV

Applying the information from Table 4 to a 1.8 GHz HFC network with a 396 MHz return path yields the information shown in Figure 5.




**Figure 5 – 3 Gbps return with legacy spectrum allocation with SDV**

This figure shows that legacy services on a 750 MHz HFC network with SDV will fit on a new plant with a 396 MHz return. The key is that the legacy SDV services and DOCSIS 3.0 services all fit below 1 GHz (1,002 MHz). This means that even though the services have been moved up in the spectrum that the legacy CPE should still be able to receive these services.

Note that the DOCSIS 3.1 spectrum has moved above 1 GHz (1,002 MHz) but remains below 1.2 GHz (1,218 MHz). From section 2.1.4 the existing DOCSIS 3.1 CPE will operate up to 1.218 MHz meaning they will get the same service and perhaps even better service since additional DOCSIS 3.1 spectrum could be added up to 1.2 GHz (1,218 MHz).

### 3.1.2. 3 Gbps Return Example without SDV


Applying the information from Table 5 to a 1.8 GHz HFC network with a 396 MHz return path yields the information shown in Figure 6.



**Figure 6 – 3 Gbps return with legacy spectrum allocation without SDV**

This figure shows that legacy services on a 750 MHz plant without SDV will not fit on a new plant with a 396 MHz return. The legacy video spectrum is okay, it all fits below 1 GHz (1,002 MHz) and the issue is that a portion of the DOCSIS 3.0 spectrum falls above 1 GHz (1,002 MHz). DOCSIS 3.0 CPE will not receive spectrum above 1 GHz (1,002 MHz) and the operator has a choice to make which will be discussed in the next paragraph. Also in Figure 6 note that the DOCSIS 3.1 spectrum has moved above 1 GHz, however, from 2.1.4 the existing DOCSIS 3.1 CPE will operate up to 1,218 MHz meaning those CPE will get the same service and perhaps even better service since the DOCSIS 3.1 spectrum could be extended to 1.2 GHz (1,218 MHz).

To make the DOCSIS 3.0 spectrum fit below 1,002 MHz, either the spectrum for legacy video services need to be reduced, which is likely not possible, or the DOCSIS 3.0 spectrum needs to be reduced. As shown in Figure 7, if the DOCSIS 3.0 services are reduced from 32 to 24 CTA channels then the DOCSIS 3.0 service will end right at the upper limit of 1 GHz (1,002 MHz) and the existing DOCSIS 3.0 CPE will operate with a somewhat reduced capacity. To plan for this move, the operator can begin replacing DOCSIS 3.0 CPE with DOCSIS 3.1 CPE and can expand the DOCSIS 3.1 spectrum.



**Figure 7 – 3 Gbps return with modified legacy spectrum allocation without SDV**

Also, in Figure 7 the DOCSIS 3.1 spectrum has been doubled to 32 CTA channels (192 MHz) which fits below the limit of 1.2 GHz (1,218 MHz). Whereas there was a reduction of DOCSIS 3.0 spectrum, there can be an increase in the DOCSIS 3.1 spectrum with an overall increase of spectrum available to DOCSIS 3.1 modems since DOCSIS 3.1 modems can also use the DOCSIS 3.0 spectrum. This is an example of how a proactive swap to DOCSIS 3.1 modems will position an operator for DOCSIS 4.0.

### 3.2. 4 Gbps Return Example

A 492 MHz return will support up to 4 Gbps of upstream capacity using DOCSIS 4.0 technology.

With a 492 MHz return, the forward path starts at 612 MHz which means there is only 390 MHz of spectrum up to the legacy services limit of 1 GHz (1,002 MHz).

### 3.2.1. 4 Gbps Return Example with SDV

Applying the information from Table 4 to the plant with a 492 MHz return path yields the information shown in Figure 8.




Figure 8 – 4 Gbps with legacy spectrum allocation with SDV

Much like the example from section 3.1.2, this figure shows that legacy services on a 750 MHz plant with SDV will not fit on a plant with a 492 MHz return without some tweaking, because a portion of the DOCSIS 3.0 services fall above 1 GHz (1,002 MHz). As a result, the DOCSIS 3.0 spectrum would have to be trimmed somewhat, from the original 32 CTA channels (196 MHz) down to 25 CTA channels (150 MHz) as shown in Figure 9.




Figure 9 – 4 Gbps with legacy spectrum allocation with SDV

With this reduction in the DOCSIS 3.0 spectrum, all the legacy services will fit on a plant that now supports 4 Gbps capacity in the return. Note also that all 192 MHz of spectrum for DOCSIS 3.1 modems also fits below the 1.2 GHz (1,218 MHz) limit, therefore, the DOCSIS 3.1 modems will continue to have all that capacity available.

### 3.2.2. 4 Gbps Return Example without SDV

This particular example does not work well. Applying the information from Table 5 to the plant with a 492 MHz return path yields the information shown in Figure 10 below and shows that legacy services (without SDV) on a 750 MHz HFC network will not easily fit on a new plant with a 492 MHz return.




Figure 10 – 4 Gbps return with legacy spectrum allocation without SDV

The key is that almost all of the DOCSIS 3.0 services fall above 1 GHz (1,002 MHz) which means that the legacy DOCSIS 3.0 CPE will not have enough spectrum to maintain services; only 30 MHz of downstream spectrum would be available for DOCSIS 3.0 modems, down from 192 MHz, which is not enough.

There are ways to make this work, depending on the steps an operator wants to take. Without impacting legacy video services, an operator could swap all DOCSIS 3.0 and earlier modems for DOCSIS 3.1. If the operator was OK with reducing video services to expand Internet services, they could cut back on traditional video programming. There are options and the benefit of getting to 4 Gbps of return capacity (and in fact 4 Gbps of symmetric capacity) could be a catalyst for change.

### 3.3. Other Return Examples

Table 2 lists 6 new return splits available with DOCSIS 4.0 technology including 300 MHz, 396 MHz, 492 MHz, 588 MHz and 684 MHz. The previous examples only addressed two of these: 396 MHz and 492 MHz.

What is envisioned is the operator choose one of the new DOCSIS 4.0 return splits, one that is easy to achieve while maintaining legacy services and make that change. The new HFC network should support enough capacity to get started with symmetric gigabit service. Then the operator can begin addressing legacy services with an eye of transitioning to one of the even higher return splits to offer even more upstream capacity and move to symmetric multigigabit services.


## 4. Multigigabit Forward Example

The previous sections have discussed multigigabit return services, specifically looking at fitting in existing services when the return spectrum is increased. Now let's look at the forward path.

Any of the multigigabit return examples in section 3 can also have a multigigabit forward, making for multigigabit symmetric capacity. This section builds on the example in section 3.2.1 by adding DOCSIS 4.0 downstream capacity to the HFC network which then has a downstream capacity of 6 Gbps and an upstream capacity of 4 Gbps, which would easily support a 2 Gbps symmetric service tier.

With DOCSIS 4.0 technology the forward path extends to 1.8 GHz (1,794 MHz), up from 1.2 GHz (1,218 MHz) with the DOCSIS 3.1 specifications. This increase of 0.6 GHz (600 MHz) is planned for just DOCSIS services. Couple this with the fact that DOCSIS 4.0 modems will be able to use all the DOCSIS 3.0 and DOCSIS 3.1 spectrum, it's possible to envision a multigigabit forward service with spectrum to spare, all below 1.8 GHz.

Figure 11 shows an example where 384 MHz of additional downstream is added for DOCSIS 4.0 services. Note that only up to 1.576 GHz is used, not even the full 1.8 GHz.




**Figure 11 – Symmetric Multigigabit Service Example**

A DOCSIS 4.0 modem would be able to bond all of the DOCSIS 3.0, DOCSIS 3.1 and DOCSIS 4.0 spectrum. Combined all this DOCSIS spectrum is about 6 Gbps of capacity. And the return in this



example supports up to 4 Gbps capacity. These types of numbers should easily support a 2 Gbps symmetric services, with the possibility for even higher tiers.

Figure 12 is simply Figure 11 redrawn to show the downstream DOCSIS spectrum all in same color.




**Figure 12 – Symmetric Multigigabit Service Example**

The result of the scenario in Figure 12 is over 6 Gbps of downstream capacity (726 MHz of downstream spectrum) and 4 Gbps of upstream capacity (480 MHz of upstream spectrum), all while maintaining legacy video services using SDV technology. It's a beautiful picture, all this capacity available for Internet services while gracefully maintaining legacy video services for some time into the future.

As mentioned in section 3.3, once the plant has been consolidated on an easy-to-achieve new return split that maintains both legacy video and Internet services, the operator can then consider migrating to one of the even higher return splits defined in Table 2, as well as perhaps extending the forward spectrum all the way to 1.8 GHz (1,794 MHz).

Figure 13 below shows the spectrum that will ultimately be available with DOCSIS 4.0 technology and includes a return path up to 684 MHz and a forward path consisting of 858 MHz of spectrum, ending at 1.8 GHz (1,794 MHz).



**Figure 13 – DOCSIS 4.0 Technology Spectrum Plan**

Configuring the plant in this way will support over 7 Gbps of downstream capacity and 5 Gbps of upstream capacity, which will set the stage for the existing HFC network to provide Internet services for many years to come.

## Conclusion

DOCSIS 4.0 technology is being designed to both enable symmetric multigigabit services and to provide operators with options for migrating their HFC networks to achieve this goal. Specifically, DOCSIS 4.0 technology includes adding both upstream and downstream capacity and the tools necessary to support symmetric multigigabit services.

DOCSIS 4.0 technology also includes flexible options for increasing the return capacity in steps, following either the full duplex path or by following the flexible split path. DOCSIS 4.0 will allow operators choices as the network is migrated to symmetric Internet service. DOCSIS 4.0 modems can initially be deployed on one return, and later the HFC network can be changed to another (increased) return split and that same modem can stay in service.

Depending on how the HFC network is operated today, there is the possibility to retain all existing legacy video services while the DOCSIS capacity is expanded. Scenarios were provided in section 3 that illustrate several high-runner cases, and the analysis is flexible enough to be applied to other scenarios.

Preparing the HFC network for DOCSIS 4.0 technology includes the same process the cable industry has been following from the 1960s as shown in Table 1. The capacity of the HFC network has always been expanding, and it has even more room to grow beyond what is envisioned for the DOCSIS 4.0 specifications. Coaxial cable has turned out to be a fabulous medium for broadband services.

An interesting point for discussion is upgrading the DOCSIS 3.0 modems to DOCSIS 3.1 modems. Not all DOCSIS 3.0 modems need to be replaced; the goal is to reduce dependence of the DOCSIS 3.0 spectrum such that if the amount of DOCSIS 3.0 spectrum needs to be reduced to fit below 1 GHz (1,002 MHz) that customers using DOCSIS 3.0 modems will not be adversely affected. This is an area for additional study; however, reducing dependence on the DOCSIS 3.0 spectrum will facilitate a transition to a higher return split.

DOCSIS 4.0 modems will support multiple return splits. This capability was first introduced in DOCSIS 3.1 technology where two return splits were envisioned, and this capability has turned out to be useful.

And with the availability of multiple, flexible return splits, operators have choices in how their HFC network is operated. When first moving to DOCSIS 4.0 technology, the HFC network can be moved to a first new higher split that supports legacy services, and later on the HFC network can be moved again to support an even higher split all while using the same DOCSIS 4.0 modem.

## Abbreviations

CPE	Customer Premise Equipment
CTA	Consumer Technology Association
DOCSIS	Data Over Cable Service Interface Specifications
DTA	digital terminal adapter
FDX	full duplex
FEC	forward error correction
FM	frequency modulated
GHz	Gigahertz, or billions of cycles per second
HFC	hybrid fiber-coax
HD	high definition
Hz	hertz
IP	Internet Protocol
ISBE	International Society of Broadband Experts
MHz	Megahertz, or millions of cycles per second
MPEG	Moving Picture Experts Group
NTSC	National Television System Committee
PAL	phase alternating line
SCTE	Society of Cable Telecommunications Engineers
SD	standard definition
SDV	switched digital video
VOD	video on demand
VoIP	voice over Internet Protocol

## **Bibliography & References**

Counting Channels, SCTE Technical Columns, Communications Technology, Ron Hranac, 2008.

CTA-542-D R-2018, Cable Television Channel Identification Plan, Consumer Technology Association, 2018.

DOCSIS 3.0 Physical Layer Specification, SP-CM-PHYv3.0-I13-170111, 2017.

DOCSIS 3.1 Physical Layer Specification, SP-CM-PHYv3.1-I16-190121, 2019.

Switched Unicast: It's not just about capacity, R. Oz, NCTA Technical Papers, 2006.

The Statistics of Switched Broadcast, N. Sinha and R. Oz, SCTE Conference on Emerging Technologies, 2005.

# **Predicting the Evolution of Distributed Denial of Service Attacks on Carrier Networks**

## **An Analysis of Shared DDoS Data**

A Technical Paper prepared for SCTE•ISBE by

**Kyle Haefner**

Senior Security Engineer

Cable Television Laboratories Inc.

858 Coal Creek Circle, Louisville, CO 80027

Phone:303-661-3320

k.haefner@cablelabs.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Background .....	4
1. DDoS Information Sharing Project Overview .....	4
2. DDoS Introduction and History .....	5
2.1. Cost of DDOS.....	5
3. Taxonomy of DDoS Attacks .....	5
3.1. Overview of DDoS Detection .....	6
Methods and Results .....	6
4. Global Attack Statistics.....	7
5. Attack Prediction .....	9
5.1. Long Term Short Term Memory (LSTM) Attack Prediction .....	9
6. Shodan Data.....	10
6.1. Random Forest Predictions on Shodan Data. ....	12
Conclusion .....	15
Appendix .....	16
Abbreviations.....	16
Bibliography & References .....	17

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1: DDoS Information Sharing Overview .....	4
Figure 2: Top Attacks by Type .....	7
Figure 3: Top Attacks Countries.....	8
Figure 4: Top Attacks by Bandwidth.....	8
Figure 5: Top Attacks by Packets.....	8
Figure 6: Top Attacks by Duration.....	9
Figure 7: LSTM Attack Prediction.....	10
Figure 8: Top Attackers by Service .....	11
Figure 9: Top Attackers by Protocol .....	11
Figure 10: DNS Amplification Top Features .....	12
Figure 11: NTP Amplification Top Features.....	13
Figure 12: CLDAP Amplification Top Features .....	13
Figure 13: IP Fragmentation Top Features.....	14
Figure 14: Total Traffic Top Features .....	14
Figure 15: UDP Top Features.....	14

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1: Top attackers by operating system .....	11
Table 2: Classifier results on attack data.....	12
Table 3: Attack Labels Used .....	16

# Introduction

Distributed Denial of Service (DDoS) attacks are among the preeminent threats facing the Internet today. Predicting where the next DDoS attack will emanate at an endpoint/subscriber level is a long-sought goal of the cyber-security community.

This work evaluates attack data from five contributing members of the DDoS Information Sharing (DIS) project with the intent to provide an ISP/MSO the tools to predict *at subscriber/endpoint granularity* if they will start participating in a DDoS attack. The DIS data is combined with data from the Internet search engine, Shodan, to build a detailed dataset of recent/active attackers. Statistical and machine learning analysis of this composite dataset demonstrates that by evaluating network endpoints with certain features, it can be predicted that these endpoints will participate in a specific type of DDoS attack with accuracies between 91-98%.

Finally, each feature of the attacking network endpoint that was used in the machine learning model is ranked by its predictive significance, lending insight into how ISP/MSOs might *preemptively* detect and mitigate an endpoint even before it starts participating in a DDoS attack.

## Background

### 1. DDoS Information Sharing Project Overview

The DIS project began as a pilot in early 2017 through the Messaging, Malware, Mobile Anti-Abuse Working Group (M3AAWG). At its core, its purpose is to allow large ISPs to share their attack data with the goal to help them remediate compromised and vulnerable systems running within their own networks.

This is accomplished by using a trusted third party to aggregate attack data from participants and provide API access to this aggregated data in a way that the ISP can see the attacks that are emanating from within their own AS (Autonomous System) as shown in Figure 1.




Figure 1: DDoS Information Sharing Overview

## 2. DDoS Introduction and History

Distributed Denial of Service attacks are attacks that emanate from many different and distributed sources and usually target a single entity on the Internet. The goal of the attacker is to overwhelm a service provider and deny legitimate users' access to the service.

The first notable DDoS attack that was widespread occurred in February 2000 by a 15-year-old going by the handle MafiaBoy (Michael Calce) (HERSHER, 2015). MafiaBoy enlisted several University servers to send many simultaneous requests and brought down some of the biggest names in e-commerce at the time including Amazon, CNN, Dell, eBay and Yahoo. The attack itself was small in scale and simple by today's standards, using only a handful of powerful computers to submit many legitimate requests simultaneously. The web servers of the time were overwhelmed and unresponsive for a span of several hours and in some cases several days.

In 2007 one of the first documented cyber-warfare DDoS attacks crippled the government of Estonia (Goth, 2007). The attack was relatively small in terms of bandwidth and targeted only a handful of websites however the collateral damage overwhelmed Estonia's network infrastructure effectively taking the entire nation offline. The politically motivated attack is largely blamed on Russian actors and resulted in the drafting of new international laws, notably the Tallinn manual on the international law applicable to cyber warfare (Schmitt, 2013).

The Mirai botnet attack of 2016 was not only one of the largest DDoS attacks in terms of bandwidth at 1.1 Tbps it, was also one of the most disruptive (Kolias, 2017). Composed of a botnet of over 600k compromised Internet-of-Things (IoT) devices, this attack targeted high-profile services such as the Dyn DNS (Domain Name Service) provider blocking access to many popular websites, such as Twitter, Netflix, Reddit and GitHub for many hours.

As of this writing the largest ever published DDoS attack in terms of bandwidth was reported by Imperva, where in April of 2019 they reported a SYN DDoS attack of 500 million packet-per-second attack resulting in a phenomenal 3.4 Tbps (Crane, 2019)!

### 2.1. Cost of DDOS

Denial of service attacks are so effective because they are extremely cheap for the attacker and extremely expensive for the victim. The rise of DDoS as a Service (DDoSaaS) on the dark web has commercialized and commoditized these types of attacks, drastically lowering the total cost and barrier to entry required to launch them. An analysis by Kaspersky Labs examined several DDoSaaS providers on the dark web and found that a DDoS attack lasting 300 seconds with a bandwidth of 125 Gbps will cost as low as \$6 (US). Others advertise an hourly rate of \$20 per hour for attacks in the hundreds of Gbps, and offer various plans and a simple pricing structure based on type and scope of attack (Makrushin, 2017).

For victims the cost of a DDoS is much higher, Incapsula surveyed 270 North American organizations and estimated that a targeted DDoS attack costs a victim an average of \$40,000 per hour (Mathews, 2014). B2B International research firm estimated that a DDoS attack costs enterprises an average of \$2 million per incidence (Kobialka, 2018).

## 3. Taxonomy of DDoS Attacks

DDoS attacks can generally be divided into three broad categories, volumetric attacks, protocol specific, and application specific. The DIS data have examples of attacks from all three categories.



## **Volumetric Attacks**

Volumetric attacks are designed to saturate the target network by flooding it with traffic. Some examples of this type of attack include UDP floods, and ICMP floods. Volumetric attacks often have a broad affect even if they have a narrow target. For example, a flood of spoofed UDP traffic aimed at specific server could saturate many of the network paths leading to that server causing access to other websites and services to become unreachable.

## **Protocol Attacks**

Protocol attacks take advantage of specific weakness in protocols and focus on depleting resources. Protocol attacks work against specific servers, or intermediary network equipment such as firewalls, load balancers and SDN controllers.

## **Application Attacks**

This type of attack does not produce high levels of network traffic, instead the attack is targeted at specific server applications with the goal of making the service exhaust CPU or memory resources.

### **3.1. Overview of DDoS Detection**

Detecting DDoS attacks is a challenging problem due to the heterogenous nature of how these attacks are carried out and requires an equally heterogenous set of solutions. There are three main categories that are used for DDoS detection, statistical methods, knowledge-based methods and machine learning based methods. Often combinations of these three are used simultaneously.

#### **Statistical Methods**

Statistical methods generally attempt to model the normal traffic and then test any new traffic or flows to determine if it belongs to the normal set or is an anomaly.

#### **Knowledge Based Methods**

Knowledge based approaches to detecting attacks use predefined rules and patterns to determine if the traffic is an attack or not. Some examples of knowledge-based approaches can be as simple as threshold-based systems to more complex state-transition and signature analysis.

#### **Machine Learning Methods**

Machine learning methods of detecting DDoS attacks apply machine learning enormous ability to absorb vast amounts of data and learn classifications across that data.

## **Methods and Results**


This section is divided into two parts. Part one analyzes global statistics of attacks across all participants, giving a broad picture of the scope, scale and structure of attacks. A prediction algorithm that uses sequential data was run on the global data to predict the next attack in the sequence.

Part two uses recent attack events from the DIS project to query Shodan, a search engine for connected devices. The Shodan query returns fine grained details such as port, service and operating systems of

active/recent attackers. These details are then composed into features and fed into machine learning predictive models.


## 4. Global Attack Statistics

Global statistics are generated from attack events submitted across all DIS participants. Attack events are purged by the system if they are older than 30 days; however, reports are generated on a weekly basis and sent to participants. These weekly reports store aggregated meta data such as top attack types, countries etc. The following graphs show data from the first six months of 2019.



**Figure 2: Top Attacks by Type**


Figure 2 shows the top attack types seen by the DIS system over the first six months of 2019.




**Figure 3: Top Attacks Countries**

Figure 3 show the origin of attacks by county for the first six months of 2019. These countries and their rank are consistent with other sources and analysis. (Link) (Akamai, 2019)

The following figures represent a sampling of data gathered between June 1, 2019, and June 30, 2019.




**Figure 4: Top Attacks by Bandwidth**



**Figure 5: Top Attacks by Packets**

Figure 4 above indicates that most of the attacks are not volumetric in nature with the median attack only 179Kpbs and the plurality (mode) of the attacks having a bandwidth of only 50Kbps.

Figure 5 furthers this observation where the median packets per second is only 19 pps and the plurality (mode) is only **4 pps**.



**Figure 6: Top Attacks by Duration**

Figure 6 shows that the average attack is just over 20 minutes and the most frequent attack lasts only a little more than 7 minutes. Figure 6 also shows that the longest measured attack was almost 60 *days* long (active and ongoing attacks are not purged every thirty days).

## 5. Attack Prediction

Attack data from participants comes with a label (Table 3) generated by the Netscout Arbor system running on the participant's network. This global data only contains a few usable features making it largely unsuitable for predictive analysis. However, some predictive analysis can be run using only the labels and the associated IP address.


### 5.1. Long Term Short Term Memory (LSTM) Attack Prediction

Recurrent Neural Networks (RNNs) are machine learning neural networks that are useful on time-series data. Long-Term Short-Term Memory (LSTM), are subset of RNNs that can learn long term dependencies and are particularly well suited to learning from large sets of sequential data and are used extensively on word prediction algorithms.

The hypothesis in this analysis is that a host that is compromised by one malware is likely to become compromised by another malware. Each malware forms a botnet that propagates unique types of attacks. These attacks form a timeline of attacks such that a host that propagates attack type A, later propagates attack type B. The predictive model analysis then follows that if a Host<sub>X</sub>, that propagates a series of attacks,

Host<sub>X</sub>: Attack<sub>A</sub> → Attack<sub>B</sub> → Attack<sub>C</sub>...

how likely is that host to also exhibit Attack<sub>D</sub>?



**Figure 7: LSTM Attack Prediction**

As can be seen in Figure 7, the blue line is the series of attack data from a single IP address that propagated several types of attacks (in order for the model to analyze the data, attacks names were each converted a number see Table 3: Attack Labels). The orange line represents the model being trained and the green line is the model's prediction.

Based on the graph, the LSTM prediction model did *not* accurately predict the attack type. The primary reason for an LSTM model to fail is that the data is too random in nature to be predicted. To confirm this hypothesis, the augmented Dickey-Fuller (ADF) test was run across the data (Dickey, 1979).

The ADF test shows how strongly series-based data can be defined by a trend. If the ADF test results in a positive score, then the data has a series-dependent structure to it. If the score is negative, then the data is too random to predict. The ADF score for the DIS labeled data run across several hosts that had multiple attack types resulted in a score of -5.9611. This score implies that the data is too random to perform time-series predictions.

Why is the data random? It is believed that the DHCP assignment of IP addresses to attacking hosts introduces noise into the data that prevents any predictive analysis to be run. The LSTM model could be a viable method of prediction if host data could be statically tied back to a consistent IP address. For MSO/ISPs, this could be accomplished with internal hosts using current DHCP assignment logs.

## 6. Shodan Data


Shodan is a search engine that continuously scans the Internet for open and accessible ports. When Shodan successfully connects to a port it grabs the response from the host and stores and indexes the returned data, called the banner, along with the IP address in a database. Shodan offers an application program interface (API) to query this database.

For this research the most recent 100K events representing recent/active DDoS attacks from the DIS service were used. For each event the IP was extracted and then used to query the Shodan database. Queries that returned a recent record from Shodan are considered to be recent/active attackers.


Of the 100K events in DIS, Shodan returned 2,388 recent records. By parsing the banner of each returned record, and transforming the categorical features (operating system, protocol, etc.) into numbers using a technique called one hot encoding, the data was expanded to 231,226 examples (rows) and 225 features (columns). From this new composite dataset some basic aggregation analysis was done across the dataset. Table 1 below shows a breakdown of the attackers based on operating system.

**Table 1: Top attackers by operating system**

Operating System (OS)	Attacker Count
Windows Server (various versions)	201
Windows Desktop (various versions)	52
Linux 3.x	48
Windows Embedded	6
Linux 2.6.x	4
Darwin (MacOS)	3



**Figure 8: Top Attackers by Service**



**Figure 9: Top Attackers by Protocol**

Figure 8 shows the breakdown of attackers based on service running on the host. It is notable that the MikroTik bandwidth test service is detected at a rate nearly five times more than the next nearest service, Microsoft's IIS http (web) server.

Figure 9 shows the breakdown of attackers by protocol. Each of the protocols was derived from the Shodan module that was used to detect it. Domain Name Service , a frequent source of DDoS attacks due to the high amplification potential of the protocol, takes the top two slots.

The appearance of the ikettle protocol in the top four was a notable result. The iKettle protocol is a binary protocol that runs over UDP or TCP and port 2081. This protocol is used to control a smart WIFI-connected kettle of the same name used to heat water for tea or coffee. iKettles come with a default password of "000000". Once the iKettle is joined to the user's network, it allows connections to port 23 using Telnet. An attacker can connect to the iKettle and request it to list its settings, one of which is the WIFI password that is in plain text (hughes, 2015).

## 6.1. Random Forest Predictions on Shodan Data.


The composite dataset has the label imported from the DIS data and many additional features from the Shodan data, so supervised learning classifiers can now be deployed to predict the type of attack using the weighted Random Forest classifier (Liaw, 2002). The weighted Random Forest classifier belongs to a set of classifiers that use an ensemble of decision trees to build a model of the data.

The weighted Random Forest classifier works by building several predictive models and votes on the best one. This classifier is particularly useful when applied to unbalanced data, i.e., where the labels are not evenly split. The dataset derived from top six DIS attack types as labels combined with the Shodan host data had majority sets made up of attack types DNS amplification and NTP amplification with the other four belonging to minority sets. The weighted Random Forest model assigns a higher weight and misclassification cost to the minority classes this in turn reduces the bias toward the majority classes.

The composite attack dataset was split into two parts where 80% was used for training and 20% was reserved for testing. Table 2 shows the accuracy of the classifier of predicting the correct label on the test data.


**Table 2: Classifier results on attack data**

Attack Label	Random Forest Classifier Accuracy
DNS Amplification	98.036%
NTP Amplification	98.468%
CLDAP Amplification	95.055%
IP Fragmentation	91.204%
Total Traffic	96.265%
UDP	98.234%




**Figure 10: DNS Amplification Top Features**

Figure 10 shows the top features used in predicting a DNS amplification attack, as would be expected port 53 is seen. Echo refers to the echo protocol (RFC 862) that is associated with the init.d services on Linux. Based on Figure 10, this classifier utilized the ports feature predominantly in predicting the attack classification of DNS amplification. QOTD is a quote of the day service.



**Figure 11: NTP Amplification Top Features**


Figure 11 shows the features for an NTP amplification attack. Features for the ntpd service, ntp protocol and port 123 make logical sense here. The figure shows a close ranking of feature importance, especially with the ntpd service, OpenSSH and the ports features. This implies these features are of relative equal importance to the attack classification of NTP amplification.




**Figure 12: CLDAP Amplification Top Features**

Figure 12 shows the features used in the CLDAP classification. CLDAP is associated with Microsoft Active Directory services and often serves as an application layer ‘ping’ for Active Directory. Here the open ports were the predominant feature used by the classifier.






**Figure 13: IP Fragmentation Top Features**



**Figure 14: Total Traffic Top Features**



**Figure 15: UDP Top Features**

Figure 13, Figure 14, and Figure 15 are grouped together as they have similar predominant features. Despite having similar features, the classifier was able to predict each attack classification with a greater than 90% accuracy rate.

# Conclusion

In this research, DDoS Information Sharing (DIS) data was explored showing the top attack types and top attack origins as seen by the participating members. Using the sequence of attacks from each IP address, the feasibility of using LSTM RNNs for predicting the next attack type in the sequence for a given IP address was examined. The current data is too random to pursue this method by itself, however, it is believed with the addition of DHCP lease logs, this remains a viable prediction model.

Next, the latest events seen by the DIS service was used to query the Internet search engine, Shodan. From these queries, a detailed dataset was built for recent/active attackers that showed the top operating systems, services and protocols running on each attacker.

This new dataset was analyzed using the Random Forest ensemble classifier to predict the attack type of an endpoint based on open ports and the information these hosts present when connecting to them. The model was able to correctly predict the DDoS attack type with accuracies above 90% for each type of attack. Lastly, a breakdown of each attack type and the features of the attacker that are most important to the predictive model was shown.

The research presented in this paper could be directly applied by an ISP/MSO to predict which subscribers have a node that have been or can be compromised on their network. This in turn could be used in remediation efforts and upstream DDoS prevention services, preemptively, before the compromised node has started participating in a DDoS attack.

# Appendix

**Table 3: Attack Labels Used**

#	Attack Name	#	Attack Name
1	UDP	12	NTP Amplification
2	Total Traffic	13	SSDP Amplification
3	DNS Amplification	14	chargen Amplification
4	TCP RST	15	SNMP Amplification
5	IP Fragmentation	16	MS SQL RS Amplification
6	DNS	17	rpcbind Amplification
7	TCP SYN	18	memcached Amplification
8	TCP SYN/ACK Amplification	19	RIPv1 Amplification
9	ICMP	20	mDNS Amplification
10	CLDAP Amplification	21	NetBIOS Amplification
11	TCP NULL		

## Abbreviations

API	application program interface
AS	Autonomous System
bps	Bits per second
CLDAP	Connection-less Lightweight Directory Access Protocol
DDoS	Distributed Denial of Service
DDoSaaS	DDoS as a Service
DHCP	Dynamic Host Configuration Protocol
DIS	DDoS Information Sharing
DNS	Domain Name System
Gbps	Gigabits per second
ICMP	Internet Control Message Protocol
IoT	Internet of Things
Kbps	Kilobits per second
LSTM	Long-Term Short-Term Memory
M3AAWG	Messaging, Malware, Mobile Anti-Abuse Working Group
NTP	Network Time Protocol
RNN	Recurrent Neural Network
SDN	Software-Defined Networking
Tbps	Terabits per second
TCP	Transmission Control Protocol
UDP	User Datagram Protocol

## Bibliography & References

- Akamai. (2019, July 06). *Web Attack Visualization*. (Akamai) Retrieved July 14, 2019, from <https://www.akamai.com/us/en/resources/our-thinking/state-of-the-internet-report/web-attack-visualization.jsp>
- Crane, C. (2019, May 29). *The Largest DDoS Attacks in history*. Retrieved July 1, 2019, from Casey Crane Read more at: <https://www.thesslstore.com/blog/largest-ddos-attack-in-history/>
- Dickey, D. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association*, 74(266a), 427-341.
- Goth, G. (2007). The Politics of DDoS attacks. *IEE Distributed Systems Online*, 8(8), 3-3.
- HERSHER, R. (2015, February 7). *Meet Mafiaboy, The 'Bratty Kid' Who Took Down The Internet*. Retrieved July 1, 2019, from <https://www.npr.org/sections/alltechconsidered/2015/02/07/384567322/meet-mafiaboy-the-bratty-kid-who-took-down-the-internet>
- hughes, M. (2015, October 23). *Why the iKettle HAcK Should Worry You (Even if You Don't Own One)*. Retrieved July 14, 2019, from <https://www.makeuseof.com/tag/ikettle-hack-worry-even-dont-one/>
- Kobialka, D. (2018, February 25). *Kaspersky Lab Study: Average Cost of Enterprise DDoS Attack Totals \$2M*. Retrieved July 1, 2019, from <https://www.msspalert.com/cybersecurity-research/kaspersky-lab-study-average-cost-of-enterprise-ddos-attack-totals-2m/>
- Kolias, C. (2017). DDoS in the IoT: Mirai and other botnets. *Computer*, 50(7), 80-84.
- Liaw, A. a. (2002). Classification and regression by randomForest. *R News*, 2(3), 18-22.
- Link, C. (n.d.). *CenturyLink 2018 Threat Report*. Century Link.
- Makrushin, D. (2017, March 23). *The cost of launching a DDoS attack*. Retrieved July 1, 2019, from <https://securelist.com/the-cost-of-launching-a-ddos-attack/77784/>
- Mathews, T. (2014). *Incapsula Survey: What DDoS Attacks Really Cost Businesses*. Retrieved July 1, 2019, from <https://lp.incapsula.com/rs/incapsulainc/images/eBook%20-%20DDoS%20Impact%20Survey.pdf>
- Schmitt, M. N. (2013). *Tallinn manual on the international law applicable to cyber warfare*. Cambridge University Press.

# **Preparing the Metro Core Network for Disruptive Technologies Like DAA and 5G**

A Technical Paper Prepared for SCTE•ISBE by

**Fady Masoud, M. Eng.**

Principal, Product and Technology Marketing  
Infinera

555 Legget Drive, Suite 222, Tower B, Ottawa, ON, Canada K2K 2X3  
fmasoud@infinera.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
1. Evolving the Metro Core.....	3
1.1. Network Scalability: .....	3
1.2. Network Economics.....	3
1.3. Network Automation .....	4
1.4. Network Evolution.....	4
1.5. Network Agility.....	4
1.6. Network Security .....	4
2. Enabling Technologies for Metro Core Evolution .....	5
2.1. Increase New Scalability Through Super-channels and Advanced Coherent Capabilities .....	5
2.2. Redefine Economics with Compact Modular Platforms.....	5
2.3. Evolve Legacy Services to the Next Generation.....	6
2.4. Automate Network Operations: .....	7
2.5. Enhance Network Agility with Software-defined Capacity .....	8
2.6. Protect Critical Data While on the Optical Network.....	8
Conclusion .....	9
Abbreviations.....	9
Bibliography and References.....	10

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1: Typical Optical Transport Network.....	4
Figure 2: Increasing fiber capacity with super-channels.....	5
Figure 3: New network economics with compact modular .....	6
Figure 4: Evolving the network with universal switching .....	7
Figure 5: Software-defined capacity vs. conventional methods of adding bandwidth .....	9
Figure 6: Traffic encryption .....	9

# Introduction

The metro network is one of the most dynamic areas in the entire service delivery model, and it is expected to witness first the impact of emerging technologies such as the Internet of Things (IoT), 5G, and distributed access architectures (DAA), which will drive demand for capacity and a need for increased network agility.

It is estimated that 30 billion devices<sup>1</sup> will connect to the internet in 2020, while all-fiber access networks have surpassed xDSL connections, with nearly 60 million homes able to receive fiber to the home (FTTH) and 23.8 million homes connected in North America alone.<sup>2</sup> Moreover, the world is on the verge of the massive deployment of 5G mobility, which is poised to radically transform mobile connectivity and start a new era of high-performance mobile applications and machine-to-machine real-time communication. The number of commercial 5G network deployments is expected to grow to 55 globally in 2019, up from 13 in 2018.<sup>3</sup>

The deployment of 5G networks will dictate an unprecedented level of performance from the underlying optical transport networks, including ultra-low latency, network sliceability, and scalability. Some of these requirements, such as the dramatic reduction in latency, will require the implementation of multi-access edge computing (MEC) to process and push content closer to the end user, charting a new way of architecting and operating the network as a result. MEC is also fueling the conversion of existing access network sites into mini data centers, thus creating the need for access and core networks, along with MEC resources, to be managed as a single entity from core to edge. In 5G, the concept of network slicing is used to manage this core-to-edge flow of transport and MEC resources. Slices can be created to support the differing transport, storage, and compute resources required for each individual service type.

Overall, the deployment of data centers in metro areas has been increasing at a very fast pace. As a matter of fact, there are 3,600 data centers in the top 20 cities,<sup>4</sup> with London leading the pack with 429 data centers in its metropolitan boundaries. This paper describes the evolution needs of metro core networks and the innovative technologies that will enable network operators to gracefully embrace this journey.

## Content

### 1. Evolving the Metro Core

Given its fit in the middle of the optical transport architecture (Figure 1), between access and regional/long-haul networks, the metro core network plays a key role in ensuring the successful deployment of the wireline optical infrastructure supporting 5G and DAA. Hence, its evolution must touch multiple fronts, such as:

#### 1.1. Network Scalability:

In addition to demanding significantly more bandwidth when initially deployed, FTTH and emerging technologies like 5G will continue to increase the need for higher capacity, ensuring that a scalable metro core network becomes paramount.

#### 1.2. Network Economics

Operational expenditure (OpEx) has a direct impact on bottom line (profitability). The order of magnitude and scale of connected devices enabled by 5G, as well as the bandwidth required for DAA, will certainly

increase operating costs unless new, significantly lower network economics (power consumption, footprint, total cost of ownership, etc.) are achieved.

### 1.3. Network Automation

The proliferation of IoT and cloud applications, the push to MEC to cope with 5G requirements, and the constant expansion of data centers in the metro and closer to the edge are creating new pools of points of delivery (PODs) for services and applications. They also require a new level of practical network automation and programmability within and between the PODs in order to keep up with the constant and unpredictable demand for bandwidth, streamline operations, and eliminate sources of human error.

### 1.4. Network Evolution


While most of today's services are Ethernet-based, there is still a considerable amount of legacy services (e.g., SONET/SDH) delivered over platforms that require significantly larger footprint and higher power consumption. Preparing the network for emerging technologies must include a plan for a smooth migration toward next-generation services.

### 1.5. Network Agility

A key step toward preparing the metro core network for 5G and DAA is enhancing network agility and breaking away from current methods of optical capacity planning, engineering, and activation that require numerous truck rolls, extensive manual labor, and human interaction at multiple points in the network.

### 1.6. Network Security

As more content is being pushed to the cloud, cyber-attacks and data breaches are becoming frequent occurrences. The annual damage to the U.S. economy caused by cyber-attacks is estimated to be up to \$100 billion.<sup>5</sup> Network operators must protect data traffic carried over the network from intruders and hacking tools.



**Figure 1 - Typical Optical Transport Network**




## 2. Enabling Technologies for Metro Core Evolution

As discussed in the previous section, preparing the metro core network for 5G and DAA must touch multiple fronts. Below are some of the enabling technologies that play a crucial role in elevating network performance to the levels required by 5G and DAA:

### 2.1. Increase New Scalability Through Super-channels and Advanced Coherent Capabilities

DWDM technology disrupted the telecommunication industry by enabling multiple optical carriers to travel in parallel on a fiber, thus increasing capacity and maximizing fiber utilization. However, projected growth in traffic triggered by 5G and DAA is demanding a new level of scalability and spectral efficiency (the ability to pack more capacity on the fiber). A technology called super-channels solves the challenge of increasing network capacity quickly and without operational complexity. A super-channel includes several optical carriers combined to create a composite line-side signal of the desired capacity that is provisioned in one operational cycle (Figure 2). The use of super-channels increases spectrum efficiency and thus fiber capacity by reducing spectrum waste due to guard bands, up to 2 terabits per second when using the extended C-band. Once deployed, this service-ready terabit capacity allows seamless growth without the need for truck rolls, network re-engineering, or major disruption to current operating processes. Furthermore, the latest breakthroughs in optics and DSPs have led to the introduction of advanced coherent optical capabilities, such as Nyquist subcarriers, software-programmable modulation schemes, software-decision forward error correction (SD-FEC) gain sharing, pulse shaping, and so on, that maximize capacity over any given distance.



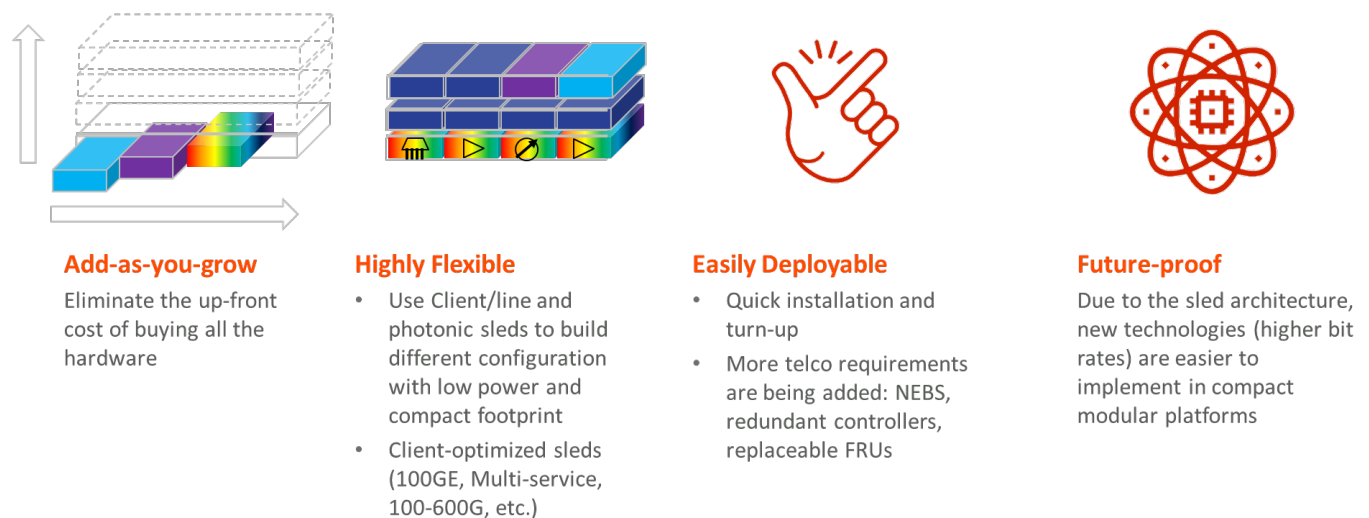
**Figure 2 - Increasing fiber capacity with super-channels**

### 2.2. Redefine Economics with Compact Modular Platforms

A new breed of optical platforms, called compact modular, has been created to set a new benchmark in scalability, low power consumption, and compact footprint. Compact modular platforms are the outcome of “disaggregating” the optical layer with design and specifications tailored to the needs of the next era of hyperconnectivity. While the first generation started around data center interconnect (DCI) applications, current and future generations are equipped with the required features and capabilities to be deployed by all optical network operators, not just internet content providers, and in a wide scope of applications

beyond DCI, including metro, regional and long-haul networks. Compact modular platforms redefine network economics (Figure 3) with


- Significantly lower power consumption, a compact footprint, add-as-you-grow architecture through a sled design that allows network operators to eliminate the up-front cost of buying all the hardware on day one
- High level of flexibility with the ability to use client/line (100 Gigabit Ethernet [GbE], multi-service, 100G-600G, etc.) and photonic (ROADM, OTDR, amps) sleds to build different configurations with low power and a compact footprint,
- Ease of deployment and turn-up so traffic can be up and running literally in minutes
- Future-proofing investment by setting a clear path for higher capacity and advanced features without forklifting, as new technologies are faster to develop and much easier to implement in a sled in compact modular platforms vs. monolithic systems.



**Figure 3 - New network economics with compact modular**

### 2.3. Evolve Legacy Services to the Next Generation

Throughout many years of service evolution, from time-division multiplexing to packet and from sub-10G to 100G+, service providers have slowly built up parallel and service-specific networks in an effort to meet service requirements and market windows. This has led to complex architectures, costly operations, and limited flexibility and scalability due to the numerous network interconnects that make the metro core network ill-suited for 5G and DAA. A new breed of packet-optical platforms supporting “universal switching” have proven to simplify network architecture and ease the evolution of legacy services by replacing many service-specific platforms and collapsing multiple networks into one flexible and highly scalable multi-service infrastructure. Deploying a universal switching platform allows service providers to smoothly and gracefully migrate legacy services to packet-optical services at their business and operational paces (Figure 4).



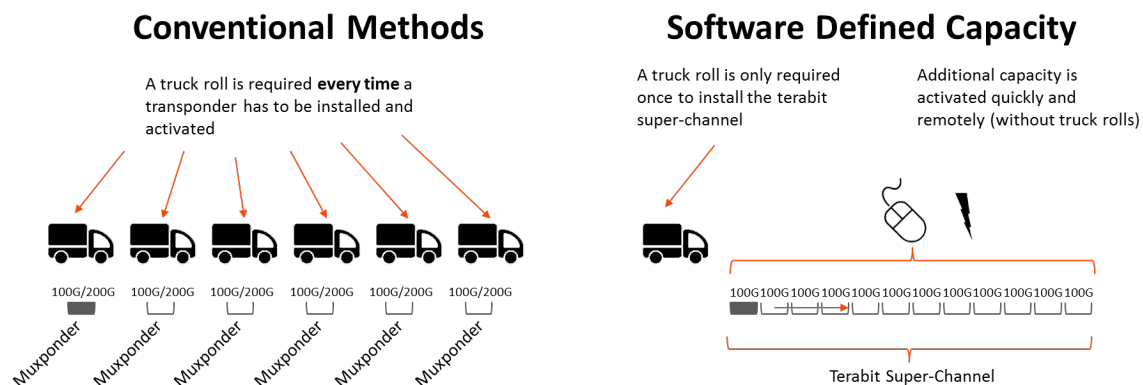
**Figure 4 - Evolving the network with universal switching**

## 2.4. Automate Network Operations:

Building practical automation throughout the network is central to preparing for 5G and DAA, as content is being processed/pushed closer to end-users with MEC, and traffic profiles and trends are becoming difficult to predict. Moreover, human error is often behind major network outages, especially in POD environments where multiple platforms (servers, routers, Layer 1/2 transport, etc.) are interconnected, hence the need to automate recurring tasks within and between the PODs for better efficiency and reliability. Network operators can use the latest developments in multi-layer, multi-domain, and multi-vendor orchestrators; SDN controllers; and virtual POD controllers. They can also take advantage of the implementation of open interfaces such as REST application programming interfaces, NETCONF/YANG, and others to simplify network management and automate recurring tasks. These multi-layer, multi-vendor, and multi-domain orchestrators and SDN and POD controllers elevate network automation to a whole new level by adding intelligence and “rule-based” capabilities for planning and traffic restoration. For example, a path computation engine (PCE) can be enhanced by adding context-oriented “rules” to automatically restore traffic while meeting certain requirements (e.g., maintain minimum latency, avoid congested links, etc.). Other automation capabilities include real-time monitoring of network parameters and optimization of networking assets like optical spectrum. Practical automation is also a stepping stone toward cognitive networking that is multi-layer, self-aware, self-organizing, and self-optimizing, and can take predictive and/or prescriptive action based on what it has gleaned from its collected data and experience.

## 2.5. Enhance Network Agility with Software-defined Capacity


A key step to preparing the metro core for 5G and DAA starts with allowing intelligent software tools to dynamically add, modify, move, and retire optical capacity based on the real-time requirements of upper-layer applications. With current and conventional methods, activating a new service or adding extra network capacity is a long, complex, and labor-intensive process that can take many weeks and requires numerous truck rolls. Breakthroughs in software and practical automation led to the creation of software-defined capacity (SDC), which provides instant software activation of additional capacity, creating a pool of bandwidth that can be dynamically allocated based on traffic demand from 5G and DAA, as depicted in Figure 5. SDC is a true game-changer from both business and operational perspectives. It enables a perfect match between the timing of capital expenditure (CapEx) and service revenue, thus accelerating time to revenue from months to minutes. SDC also reduces OpEx by streamlining operations and eliminating truck rolls. Moreover, SDC is a key enabler of automation throughout the network and across all operational levels, which is a vital element in building the foundation for cognitive networking, in which real-time network analytics, microservice-based engines, and machine-learning algorithms can dynamically increase or decrease network capacity on specific routes based on past trends, spontaneous changes in traffic demand, or an anticipated spike in capacity.



**Figure 5 - Software-defined capacity vs. conventional methods of adding bandwidth**

## 2.6. Protect Critical Data While on the Optical Network

According to a study on cybercrime<sup>6</sup>, an average organization suffered 130 cyber security breaches in 2017, up 27 percent from the previous year, with the average cost of these breaches now at \$11.7 million USD. Securing mission-critical data is more important than ever in a world where security threats are constantly on the rise. While data encryption can be performed at different layers, Layer 1 (OTN payload) and Layer 2 (MACsec) encryption provide significant advantages over upper-layer encryption like Internet Protocol Security (IPsec) encryption at Layer 3. These advantages consist of higher throughput at relatively low cost, minimized latency, and the ability to support non-IP traffic. Adding encryption requires no network re-engineering and is quite easy, as it's often already supported on compact modular platforms as well as some metro core universal switching platforms as depicted in Figure 6.



**Figure 6 - Traffic encryption**

## Conclusion

Disruptive technologies like 5G and DAA and their need to process/push content closer to end users through MEC will impact how metro core networks are built, operated, and evolved, with a new set of requirements for higher capacity, better optical performance, and a new level of automation. The latest innovations in software and hardware are evolving metro core networks to gracefully embrace this new era of machine-to-machine and human-to-machine connectivity.

## Abbreviations

5G	Fifth generation of mobile networks
API	Application programming interface
CapEx	Capital expenditure
DAA	Distributed access architectures
DCI	Data center interconnect
DSL	Digital subscriber line
DSP	Digital signal processor
DWDM	Dense wavelength-division multiplexing
FTTH	Fiber to the home
GbE	Gigabit Ethernet
gRPC	Generic Remote Procedure Call
ICP	Internet content provider
IoT	Internet of Things
IPsec	Internet Protocol Security
MPLS-TP	Multiprotocol Label Switching–Transport Profile
NETCONF	Network Configuration Protocol
OC	Optical carrier
OpEx	Operational expenditure
OTDR	Optical time-domain reflectometer
OTN	Optical Transport Network
MACsec	Media Access Control Security
MEC	Multi-access edge compute
PCE	Path computation engine
POD	Points of delivery
QAM	Quadrature amplitude modulation (8QAM, 16QAM, 64QAM)
QPSK	Quadrature phase-shift keying

REST	Representational State Transfer
ROADM	Reconfigurable optical add-drop multiplexer
SDC	Software-defined capacity
SD-FEC	Soft-decision forward error correction
SDH	Synchronous Digital Hierarchy
SDN	Software-defined networking
SE	Spectral efficiency
SLA	Service-level agreement
SONET	Synchronous Optical Networking
TCO	Total cost of ownership
YANG	Yet Another Next Generation

## Bibliography and References

- [1] <http://www.mckinsey.com/industries/high-tech/our-insights/the-internet-of-things-sizing-up-the-opportunity>
- [2] <https://www.lightwaveonline.com/business/market-research/article/16675952/ftth-now-second-to-hfc-in-north-american-broadband-homes-passed-fiber-broadband-association>
- [3] <https://www.information-age.com/the-state-of-5g-deployments-in-2019-123479261/>
- [4] Ovum, Global Data Center Analyzer 2018
- [5] <http://www.datacenterdynamics.com/security-risk/infographic-the-cost-of-cyber-attacks-in-the-us/96087.article>
- [6] Cost of Cyber Crime Study 2017 conducted by the Ponemon Institute and jointly developed by Accenture. [https://www.accenture.com/t20170926T072837Z\\_\\_w\\_/us-en/\\_acnmedia/PDF-61/Accenture-2017-CostCyberCrimeStudy.pdf](https://www.accenture.com/t20170926T072837Z__w_/us-en/_acnmedia/PDF-61/Accenture-2017-CostCyberCrimeStudy.pdf)

# Why Gaming Needs An Edge

A Technical Paper prepared for SCTE•ISBE by

**Alan Evans**

Senior Director Innovation Strategy  
EDGE GRAVITY by Ericsson  
United Kingdom  
+447826952348  
alan@edgegravity.ericsson.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
1. Foreword .....	3
2. The Classification of Gamers .....	4
3. Gaming Culture – is it possible to innovate around? .....	5
4. Delivering the game experience .....	6
5. Gaming Architectures and performance expectations .....	7
5.1. Device-Side Rendering .....	8
5.2. Server-side rendering .....	9
6. Gaming and the edge.....	10
6.1. Ping acceleration and the edge.....	11
6.2. On-demand Infrastructure at the edge .....	11
6.3. Server-side rendering (cloud gaming) and the edge .....	12
6.4. Device Offload to the edge .....	12
Conclusion .....	12
Abbreviations.....	13
Bibliography & References .....	13

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Interaction between gamer and game .....	6
Figure 2 - Potential “friction” on the interaction between gamer and game.....	7
Figure 3 - Online gaming architecture .....	8
Figure 4 - Cloud gaming architecture .....	9
Figure 5 - Latency at the broadband service provider edge .....	10

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Classifying Gamers against time, place and device.....	4



# Introduction

The gaming industry generated almost 140B USD in 2018 (Newzoo Global Games Market Report 2018/2019). It has over taken the entertainment genre of choice for many consumers both young and old, so the need for a high quality of experience has become undeniable; if the experience is poor then the game will not be successful.

The pursuit of high performance is no longer the sole domain of the PC gamer; where thousands of dollars are invested in the latest GPUs and cooling systems to deliver the greatest clock speeds and frames per second.

Rendering performance is only part of the challenge. For many gamer consumers, the performance requirement is focused on the multiplayer experience, where poor performance means poor latency, lost packets and unpredictable jitter, which translates to a laggy experience, poor game play, lost competitive matches, frustration, and abandoned game titles. Our research shows trend where blame for this poor performance is often left with internet service provider.

The concern with quality of experience (QoE) when traversing the internet is not limited to competitive experiences, PvP (Player Vs Player), it also applies to PvE (Player Vs Environment) as well. Failed collaborative experiences where a social network of gamers (aka, Group, Guild, Org, Clan or Faction) are working together on an objective for several hours is just as frustrating as a competitive match, and perhaps for some even more so.

In addition, the Quality of Experience also has an impact on the commercial aspects of the game as well. Many of the most successful titles today fall into the category of “free to play”. Some of the most popular being of a “Battle Royale” game type, are of course demanding on the gameplay performance, but also on the in-game e-commerce, analytics, business intelligence and ultimately a shift to a reoccurring revenue of the traditional one-off purchase. If these less-graphical aspects of the game are affected by poor performance, then the game will not be able to generate revenue from its audience and the game will be a failed investment for the developer.

In this paper we will explore the statement, “why gaming needs an edge”. We will consider the classification of gamers and how they lend themselves to different game architecture. We will also explore the intersection of gaming and media, eSports, and how the broadband service provider edge can play a role in delivering the experience needed. We will explore the role of a Service Provider, gaming performance metrics and our research as it pertains to gaming acceleration at the Edge.

## Content

### 1. Foreword

In 2019 no one would argue that gaming is now a mainstream form of entertainment, and access to games at anytime, anywhere, and on any device is becoming the expected norm. Just as with the TV & Media industry, the gaming industry raises similar questions of how to provide access to content on any device, provide a consistent and high-quality experience, and meet the demands of the various groups of gaming consumer; but are the answers to those question the same for both industries?

## 2. The Classification of Gamers

To be able to answer this question we should first examine the types of gamer, how the industry tends to classify them and why. NewZoo (*An excellent source of intel for the gaming industry*), and others tend to classify gamers into three categories, Casual, Competitive and Professional. These groupings are good, however do not provide a canonical view or provide any absolutes (but what does), in-fact a gamer that identifies as casual for one game, might identify as competitive in another, and so on. However, the real purpose of this classification is due to the eSports angle. Just as in more traditional sports, let's take football (soccer) as an example; a player might enjoy a casual game with friends, or they might join a competitive five-a-side team, or perhaps they have managed to make a career out of the sport and are paid and/or are sponsored for their skill.

So, let's try to describe or "profile" the three classifications, based on anytime, anywhere, and any device, i.e. time, place and device: -

**Table 1 - Classifying Gamers against time, place and device**

	Casual	Competitive	Professional
Time	Opportunistic	Session	Organised
Place		Online	Venue
Device		Expenditure	Sponsored

There is an obvious demarcation between Professional and the other two classifications, and that is that the professional gamers get paid for gaming. But if we look a little deeper we can gain a relevant insight from what is important for professional gamers, and how it relates to both the casual and competitive gamer classifications.

Professional gaming events or "Esports" are organised events, at a specific venue and are nearly always sponsored. They are organised months ahead to allow professional gamers and spectators to a) travel to a venue, b) scope and build for and manage the event, and c) get some hype going to attract as many viewers as possible, both locally and online. The venue will have a dedicated offline network to provide the best possible conditions for the gaming traffic, as well as the best devices with the best technical performance to permit the best human performance possible. Only then can the gamers know for sure that there is no advantage, or limitation based on the technology used.

To a certain extent the "Competitive" classification is very similar to the "Professional" one; games are session based, and typically are considered "online" experiences with other gamers. A session might be a quick 10 min PvP (Player vs Player), or a 5-hour Cooperative PvE (player vs environment) experience. The point here is that a) the players have come together at a specific time, sometimes across multiple time zones, and b) they are all online, and are looking for the best network performance possible to ensure their gaming experience is optimal.

The casual gamer isn't really defined by whether they play online or not, (with other people) its more of an indication of the time they have available to spend gaming, as well as being less likely to buy new hardware or watch gaming streams. One might also choose to imagine that casual gamers are less concerned with their position on a multiplayer leader board, but I have known some serious Tetris players who would disagree, and even friends who have spent hours at the arcades to finally punch their three-letter gamer tag in the number one slot on Pac Man. No, I would hypothesise that the casual gamer is one that just enjoys gaming but doesn't necessarily have the time it takes to make sure their console, PC, game, online store, software is at the latest version before spending 10min on their lunch break playing a

game, nor are they likely to spend \$120 on a season pass that provides a deep story line or extensive quests to pursue.

There is a tendency to try to segment or define the classifications by the device, specifically the amount spent on the device. But that can be misleading, most of the time the device can't be used to provide absolute classifications, it's more of a reflection of how much a gamer is prepared to spend on their device. That might be a purely income-based decision, one gamer might play several hours a day on a 200\$ console, whereas another might spend 30 min a week on a 5000\$ PC. Not to mention that many gamers will have multiple devices capable of providing a satisfying gaming experience to meet the goal of anytime, anywhere, any device.

By understanding the criteria for classification of the gamer community we gain an insight into the types of game delivery system as well as supporting architectures types and how they might be improved upon or even reimaged to meet the industry need. We will explore this in more detail later in the paper.

What I find interesting is how the amount spent on a device relates to the mindset of the gamer and their expectation on quality of experience, and the culture of gamers.

### **3. Gaming Culture – is it possible to innovate around?**

Ok, so I am not a psychologist. Neither would I describe myself as a serious student of human behaviour. But being a human being, I am inherently interested.

What was that about I hear you shouting at this text (or however you are reading this, or perhaps it made it to audio book?). Well let me explain, I believe that there is an innate desire to compare, to show off, to understand where we rank against others. I think it gives us drive, makes us try harder to improve our performance and speaks to our competitive nature no matter how hard we try to conceal it.

So, is there really any difference between Single Player vs Multiplayer games? As I hinted at earlier I have known some very serious Tetris players, and frequenters of traditional arcades, who are incredibly competitive, yes with themselves, and they exhibit a completionist focus, that probably registers somewhere on a scale of obsessive-compulsive behaviour (I may be talking about myself here a little).

So yes, of course there is a difference between single player and multiplayer games. However most of it is focussed on the architectural requirements needed to deliver the compelling gaming experience. The culture of the gaming is such that it really is about the social aspect of gaming, and our need to compare our progress, our level, our score, and often as part of a community or a social network (the traditional sense of the term). In fact I love this quote from Zach Snader of Newzoo – *“Gaming is a way for people to connect with one another on a deeply human level across a shared interest. Gaming is a form of communication as much as it is a form of entertainment”*

So, is it possible to innovate around gaming culture? I don't think it is. I believe that there isn't a single gaming delivery technology or architecture that solves all the needs of the gamer community and places it into a single paradigm (not yet anyway, maybe the holodeck?). Gamers will continue to value their social interactions above a specific technology. They will still gravitate to games, that let them demonstrate their skill, via score, level, trophies, vanity accessories, emojis, dance moves (etc etc), to strangers, friends and families, work mates.

They will continue to spend on devices, technology and services then let them achieve this, in the way they want to do that will let their skill shine through. That could mean a high-end gaming rig, with 3ms 144Hz screen, the latest low latency keyboard and mouse, that are wired as to not lose any performance

between a key stroke and the reaction of the game engine, or perhaps a service that optimises the route taken for the gaming command traffic (controller button presses) from the device to the game server, or Broadband service providers that provide the best QoS for online gaming, or game downloads and updates.

It could mean a custom controller for a console gamer that feels less awkward for extended periods of gaming, or a tether free VR headset so there is no fear of tripping up, distracting from giving the best performance. Or quite a simply a leader board that lets one “punch-in” a three-letter gamer tag when a game of Tetris has concluded during a lunch break or daily commute.

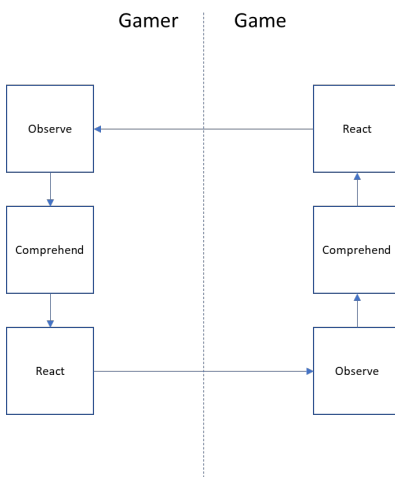
Please don’t misunderstand, this is not about one-upmanship, this is not about winning a game because the best technology or service was used or paid for. This is about removing performance impeding obstacles (or artifacts) using technology, with the objective to create an as close to a “level playing field” (a term I learnt from one of our technology partners, and fellow gamer) as possible so that high scores and victories are down to player skill, and not how much funding they have.

In our behinds-the-scenes industry of networks, cloud computing and platforms; we need to focus on how we support this culture, how can we enhance it how we can build solutions that enrich the experience of the community rather than exploiting it.

## 4. Delivering the game experience

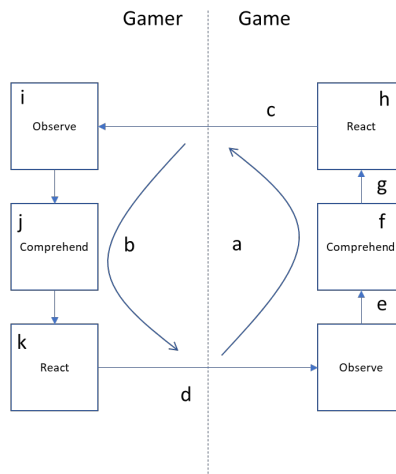
Games are traditionally hosted in the device they are being played on. i.e. the game engine runs on the device, input is taken from a control mechanism, into the game engine, the output is rendered and displayed on the screen. The gamer provides the feedback loop, and the game is experienced.

For general purposes we might draw the interaction between gamer and game like this: -



**Figure 1 - Interaction between gamer and game**

Perhaps an over simplification, but the point is that. Once the game is started the game is presented to the gamer the gamer observes, comprehends, and reacts, the game observes that reaction, decides what to do, and provides a reaction back to the gamer. The purpose for sharing this concept is to illustrate that at each stage there is potential for “friction” to be introduced as a by-product of the technology, mechanics and architecture used to deliver the game experience.



**Figure 2 - Potential “friction” on the interaction between gamer and game**

Once the propagation delay of the process (a) in figure 2 reaches a certain level, it will have a negative effect on the gaming experience. The gamers overall reaction process (b) will be so affected that the game will become unplayable and likely abandoned, in extreme cases.

If we explore this in more detail, and assume that the game has started, and the gamer is ready and about to react to the game, we can tell that delay introduced in the input (d) or information lost in the transmission would create a false observation in the game, which would affect the games comprehension (f) of what to do with the gaming input, and create an suboptimal output (g) and reaction (h) back to the gamer via (c), if further latency or poor performance affects the presentation of the game (c) to the gamers observation for the game (i), their comprehension will be frustrated (j) and the reaction (k) will attempt to compensate again potentially making the unplayable

Although the described process is quite abstract, the point here is that the interplay between gamer and game is sensitive to the potential short falls of the underlying architecture, and technology used to deliver a given gaming experience. Therefore, game developers state the minimum and optimum hardware (and software) requirements for their games to provide the experience they intended for the game and gamer, at least in the PC gaming market. The console market is slightly different, but once a console starts to “struggle” with the demands of modern game techniques, a new console is traditionally brought to market. In the case of mobile devices, there as a lot more variation of device capability so principles from both PC and Console can apply.

The interesting part about specifying device requirements to support game experience is that once the game requires connectivity for some or all its game play, it is difficult for the game studio to create a game that works with all connectivity scenarios in possible. Most state “requires broadband connection” and typically that is as much as can be done. Some will perform a network test and might suggest that the speed and ping results are not suitable to play a title. Some even might state that they are detecting poor network performance in real time and terminate a gamers connection to the online servers. Frustrating for the gamer, as this is not something that is always within their control.

## 5. Gaming Architectures and performance expectations

In this section of the paper we will talk about the gaming architectures. Not the inner workings of the game itself or game mechanics, but the gaming delivery architectures. As I mentioned earlier, to meet the

appetite of the gaming community for gaming, and gaming performance there are various delivery architectures available, and they are needed, as not all architectures work for all gamers, game types, fit all markets or network topologies etc. there are too many variables for a one-size fits all approach, today.

In my opinion we can start of with two main categories:-

1. Device-Side Rendering, and
2. Server-Side Rendering


As the options suggest, we can split the architectures by considering where the rendering occurs. i.e. where the graphics are processed, or more specifically where is the GPU (Graphics Processing Unit) if one is required for a game in question.

The following diagrams are from the Newzoo report on the gamin market. I have borrowed them and have adapted them slightly for my purpose. If you the reader would like to learn more about NewZoo and their market report. There is a reference at the back of this paper.

### 5.1. Device-Side Rendering

Traditional online gaming architectures perform the game rendering on the client device, PC, Console, Mobile Phone, Tablet, and handhelds. All these devices have varying levels of performance, with PCs clearly out in front with some gamers purchasing nVidia 2080 RTX, or Titan GPUs, that can make their gaming rig cost 10K+ USD, to custom chips made by apple in the iPhone. The purpose of these load specific devices is to reduce the load on the CPU, for a specific task. With the latest iterations of these being capable of providing ray-tracing, (lighting effects) “automatically” in the GPUs, lowering the demands on standard CPU/GPU processes.

For these solutions the game is downloaded to the device (or on physical media e.g. DVD, Blueray, which is largely a legacy approach these days) and rendered locally.




**Figure 3 - Online gaming architecture**

The challenge with this more traditional architecture; is the potential for introduced lag between the gaming device and the gaming server. Lag is the perceived delay between what the gamers interactions and what they see on screen. Its like a “glitch in the matrix” it can be very off-putting and can potentially make the difference between a positive or negative experience, or in more extreme cases a win or a loss in a competitive multiplayer experience. Therefore, professional eSports games are hosted on a private network to provide a “level playing field”. In a Virtual Reality setting, it could result in a gamer being nauseous (I speak from experience). But for player vs player, multiplayer or indeed any online game environment, Round Trip Time (RTT) is a problem. A problem for the gaming experience, and a problem for the ecommerce aspects. The more gamers are in the same online session, the more critical the RTT becomes on preserving the experience.

The great thing about this type of architecture is the game performance. The localised interaction between the gamer and the game is as good as it gets. PC providing the best performance, with graphics rendered natively and delivered to the eyeballs in flicker free 144Hz under 3ms, with all command traffic from the controller being processed locally. The only real lag being a possibility from any online requirements, and the network quality of service (QoS) required is focused on latency and jitter rather than bandwidth.

## 5.2. Server-side rendering

There is another way to deliver a game to the gamer's device, and that is streaming it as video. By streaming a game as video, the end device doesn't need to be as powerful as it is no longer performing the rendering for the game, it is simply transmitting game control input from the gamer and decoding video set to it from the Cloud. The performance required to create a compelling game experience is now the responsibility of the cloud, the game server, and the networks that deliver the video.



**Figure 4 - Cloud gaming architecture**

This technology / approach lends itself very nicely to addressing the more casual gamer who want to be able to play a game anywhere any time on any device. This doesn't mean that the more AAA titles can't be delivered in this manor, there are just some considerations that need to be acknowledged.

One benefit is that the game studio only needs to create the game once rather than a version for every platform. As I said earlier, the issue of rendering and delivering the game becomes a challenge for the cloud and the network. There are game service providers offering this type of service, with hundreds of game titles already available, in fact I think this is an important detail, this model lends itself towards the subscription gaming model where gamers pay a monthly fee in an "all-you-can-eat" model, which is a relatively new concept for the gaming community, and we will have to see how it is adopted by the free to play games, that are heavy on in-game purchases.

In addition to using this architecture paradigm for offering a catalogue of games delivered as on-demand video, some companies also offer this as a "gaming-rig-in-the sky" solution. The gamer specs their ultimate machine, and a virtual machine is created in the cloud. It doesn't come with any games, as they need to be bought separately, but the games are delivered as a video stream just the same.

The challenge this architecture faces are like the previous device-side rendering architecture. Quality of Service (QoS) parameters are important as its is critical that all the controller command traffic makes it too the cloud, and quickly, but more so that device-side rendering, because the rendering of the whole game once all input has been gathered, i.e. the player and all other players that might be in a session together, then need to be processed, then rendered, and then delivered back to the device.

The process of delivering the rendered game back to the device can't be done natively. That would require far too much bandwidth. It must be encoded as a video stream. Encoding is not a low latency process, the more time an encoder is provided to encode a picture the better the picture will look. With

some bit rates approaching 15Mbps for HD resolutions at 60fps using H.264, the picture quality, and visuals in general are a sacrifice that is made to accommodate the delivery approach.

Of course, delivering video OTT (over the top) is not a new idea. Network quality issues can be mitigated by using adaptive bit rate techniques and buffering to ensure that the video service is as uninterrupted as possible. Unfortunately, these techniques do not lend themselves well to the gaming use cases, as the gaming experience must be uninterrupted, high quality, and always available.

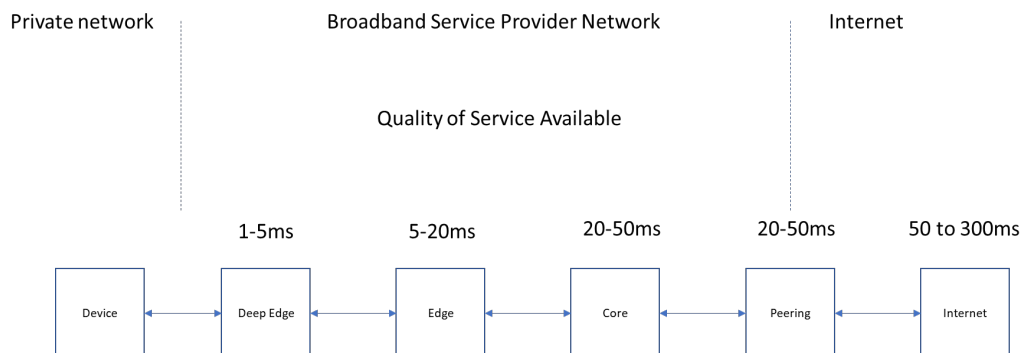
## 6. Gaming and the edge

So how can the edge help with these architectures, and does the edge offer additional opportunities for innovation around how games can be designed?

First, let's quickly qualify what we mean by the edge. For the purposes of this paper we are not going to address any single network topology, although of course in a real-world deployment the topology of different network types, 3G, 4G, 5G, HFC, DSL and Fibre (plus wifi) all have their specific considerations.

The broadband service provider edge; is the edge that we are concerned about today, and when we refer to the edge, we are referring to points-of presence (PoPs) with run-time environments that can be used to execute workloads, to gain a performance improvement by some measure.

The figures for round trip times (RTT) are purely illustrative and will vary from network to network.



**Figure 5 - Latency at the broadband service provider edge**

Earlier in this paper we introduced the concept of device-side and server-side rendering along with their pro's and con's, strengths and weaknesses, to highlight that there is potential to provide an improvement to the principle architectures that will result in an improvement of game experience, for casual, competitive and professional gamers alike.

In this paper we are not trying to suggest that there is one true gaming architecture, only that there is a place for them all, but all can benefit from optimisations the edge can provide.



## **6.1. Ping acceleration and the edge**

If a gamer has chosen device-side rendering type of environment then, as we discussed earlier the need for greater performance stems from the competitive eSports culture, and competitive online multiplayer games.

These games are typically PC based, console competitive gaming is also relevant here. Uptime and connectivity are equally relevant on any platform, PC or otherwise, and for any online-connected game. I say PC mainly because solutions in this space are provided by 3<sup>rd</sup> parties and may likely require a client to be downloaded to the device. Because latency can ruin a game, gamers turn to companies that can improve the latency, packet loss and jitter of their connection between their chosen gaming device and the gaming server.

These companies act on the command data between the game and the game server. So, this means the controller button pushes from the gamer, where they are in the game, and any information, other players in the game environment need to know about (via the game). They intercept this traffic and find the most optimal route between the gaming device and the game server.

Depending on the location of the gamer's device and the game server they are connected to, latency performance can be improved to ping levels of 10ms and lower, but perhaps just as important packet loss and jitter are reduced to zero due to the ability to perform multipath routing between the device and the game server.

Game Servers are hosted out in the internet somewhere, this can be either in a datacentre or a public cloud. By providing edge locations within the broadband service provider network, within the edge locations of the hosting datacentres and in between, the accelerator technology can make better decisions on how to route game command traffic between the device and the game server.

## **6.2. On-demand Infrastructure at the edge**

The ping accelerator technology described above acts on the command traffic between the gaming client, and the gaming server. Although this technology and approach works incredibly well and makes a significant difference for the competitive multilayer experience, the performance increase to the gamer can be variable depending on the gamer's location, the game server's location, and the route between the two.

An emerging approach to improve the gamers multiplayer experience comes from another angle, is the idea of on-demand infrastructure. Rather than accelerating the game data, move the game server closer to where the gamer is.

The challenge here is that this is that for multiplayer gaming, simply moving (instantiating) a game server to "somewhere closer" to the end user is not enough. This is largely because multiplayer gaming goes through a process called "match making".

Match making, is the process of placing a gamer into a multiplayer gaming session with other gamers of a similar skill level and gamer location. It is pretty much an internal game mechanic to make sure that the multiplayer session is a balanced one for all gamers in the session, as well as addressing language issues where communication is important to the experience.

These on-demand gaming infrastructure companies solutions complement this process. During the match making process, these solutions look at which gamers are to be added into the gaming session, and then

instantiating the game servers in the most logical point of presence to provide the best possible experience for the gamers in the multiplayer session. The placement of the server considers the performance of the PoP, both compute and network, as well as its logical and physical location to attempt a level playing field for the multiplayer session.

When the session concludes the instance is destroyed, and the resources are released. By leveraging the service provider edge, game studios would be able to provide an option for multiplayer gamers to take advantage of a multiplayer gaming infrastructure that carried with it the performance of a managed network suitable for an eSports event.

### **6.3. Server-side rendering (cloud gaming) and the edge**

Streaming video is not without its challenges. Delivery of games as video streams increases the complexity further. By bringing the cloud game servers within the broadband service provider network, the round-trip time for the command traffic and the resultant delivery of the video stream could be optimised beyond what can be delivered for this type of service from a public cloud.

The benefit of leveraging the broadband service provider edge should permit the use of higher resolution and the required bitrate of video, whilst reducing the amount of buffering in the gaming device, all of which should contribute to an improved quality of experience for the gamer.

### **6.4. Device Offload to the edge**

Until now we have mostly addressed two main architectures, but the edge might facilitate a third. With edge resources becoming available within the 1ms to 5ms RTT range. It could be possible to offload some of the processing functions of the device, towards the edge. This would allow the devices to perform more complex use cases, without draining their batteries. There would be a need to find the balance between transmit and receiving data, vs processing it on-device, but an interesting study for sure.

## **Conclusion**

Gaming technology is becoming more and more relevant to other industries, just as Sci-fi stories have predicted future technologies, the gaming industry and its technology provide insight into the consumer and enterprise applications we expect to see in the very near future. The performance requirements to create a compelling gaming experience, have even more critical impacts that can dictate the viability of a use case in various IOT scenarios.

At Edge Gravity, we are in the process testing the described architectures in conjunction with leveraging the broadband service provider edge. We are interested in collaboration in all forms, e.g. network type and topology, hardware accelerators and processors, game platforms and game developers and so on.

We expect to see test results that will demonstrate the benefits edge cloud and edge compute for these types of solutions and a show why gaming (and other industries) needs an edge.

## Abbreviations

PoPs	point(s) of presence
RTT	round trip time
AAA	triple A rated
HFC	hybrid fiber-coax
PC	personal computer
HD	high definition
QoS	quality of service
QoE	quality of experience

## Bibliography & References

Ericsson Consumer Labs 2019

Gameye CEO - Sebastiaan Heijne

Newzoo - Cloud Gaming: The Perfect Storm 2018, <https://newzoo.com/insights/trend-reports/free-report-cloud-gaming-the-perfect-storm/>

Newzoo - Global Games Market Report 2018 & 2019, <https://newzoo.com/solutions/standard/market-forecasts/global-games-market-report/>

Ericsson Consumer Labs 2019

# **Operational Impacts of Network Slicing**

## **Leveraging network slicing technologies to offer innovative business services**

A Technical Paper prepared for SCTE•ISBE by

**John Douglas**

Digital Services Sales Director, NVFi and Orchestration  
Ericsson North America  
6300 Legacy Dr, Plano, TX 75024  
(913) 241-0090  
john.douglas@ericsson.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
Conclusion .....	8
Abbreviations.....	8
Bibliography & References .....	9

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Roles related to 5G networks and network slicing.....	5
Figure 2 – Providing customer control over network slices .....	7

# Introduction

Today's networking trends create numerous opportunities for service providers, enterprises, and Mobile Virtual Network Operators (MVNOs). The migration of applications and network functions to a virtualized environment offers tremendous business value - but also requires sophisticated workflow automation across multiple domains to take full advantage and assess success criteria.

## Content

One of the most exciting opportunities created by network virtualization and workflow automation is network slicing. As a specific form of virtualization, network slicing allows multiple logical networks to be deployed on top of a shared physical network infrastructure.

5G technology is an essential business transformation enabler for communication service providers. 5G networking standards define capabilities to deliver network slicing in a new and efficient manner. Network slicing allows the same network infrastructure to deliver support for many different network use cases, each managing to different capabilities, such as, performance objectives on the same network at the same time with maximum efficiency. Network slicing takes advantage of virtualization to combine network resources and dynamically build separate logical networks for specific purposes addressing individual use cases, industries, or enterprises. Available application programming interfaces (APIs) are required to efficiently enable and operate this "Network Slice as a Service" (NSaaS) type of business model.

A key benefit of network slicing is that it is intended to provide an end-to-end virtual, distinct, and special-purpose network, including networking, compute, and storage functions. This allows service operators to deliver (slices of) networks which support specific requirements based on customized service characteristics - bandwidth, latency, volume, on-demand capacity, etc. But to ensure full advantages of this, one must re-assess certain processes to enable workflow automation.

With the dynamic nature of network slice lifecycle management, "zero touch" operations (including creation, provisioning, service enablement) and closed-loop assurance (self-healing, predictive, and proactive network/service management) are expected at the same time as operational expenditures are being reduced. This in turn creates a necessity to modernize today's operational support system (OSS) operations and maintenance procedures and evolve to process transformation, re-engineering, and automation.

Furthermore, as Virtual Machines (VMs) and containers will co-exist for the foreseeable future, a common Management and Orchestration (MANO) solution that manages both VMs and containers running on the same platform will need to be in place. With automated provisioning and insight-driven service assurance, a dynamic orchestration implementation simplifies operations, moving the service provider closer to the zero-touch experience. Additionally, dynamic orchestration enables comprehensive, modular, end-to-end solutions that support automation of all these different executions of Network Slice Instances (NSI) and Network Slice Subnet Instances (NSSIs).

- A single orchestrator for the network – managing physical, virtual and cloud native network functions
- AI-powered closed-loop assurance automatically adapts the network in real time, maintaining Service Level Agreements (SLAs)

- Automated onboarding and continuous deployment accelerating time to market in multivendor environments
- Monetization enablement of new business opportunities facilitated by an integrated converged charging network function

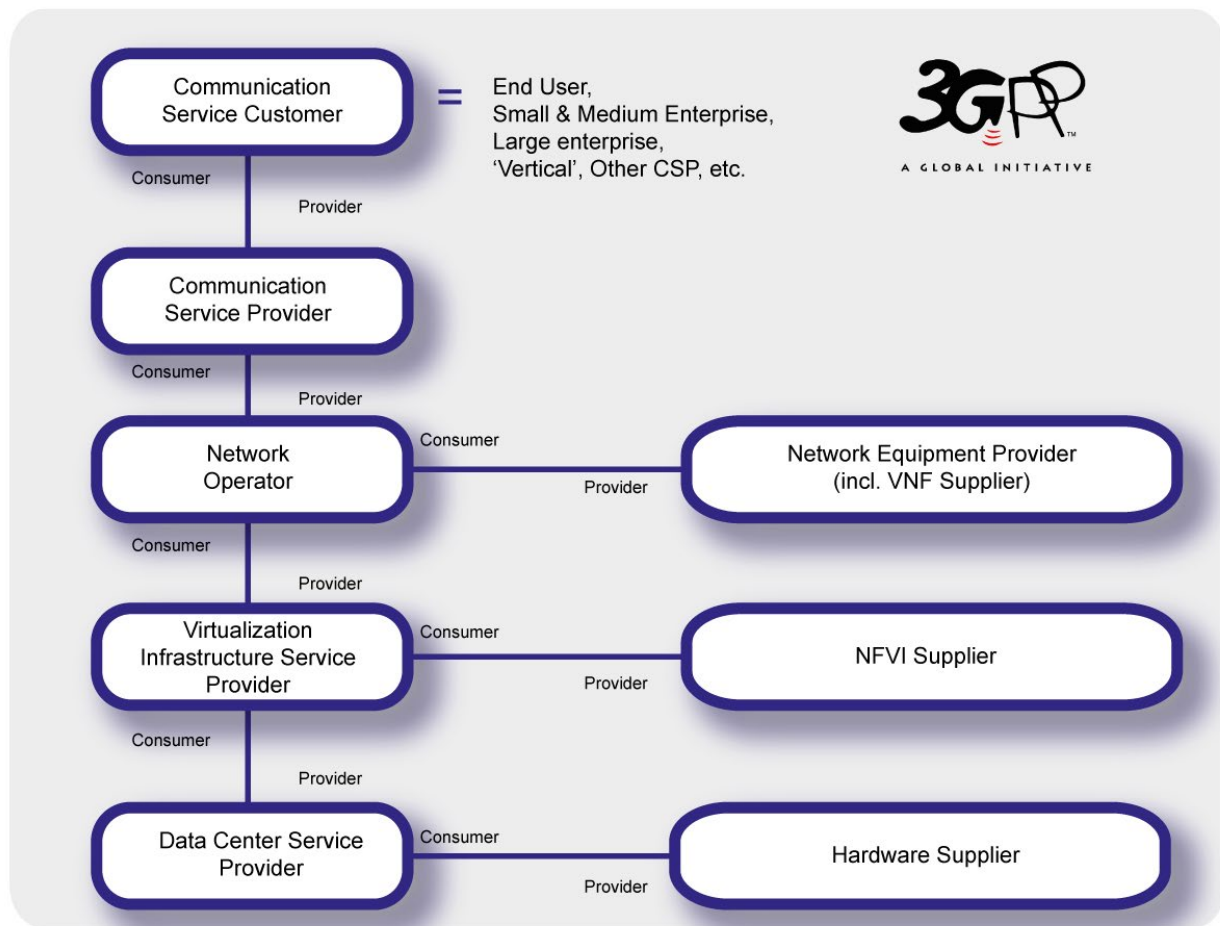
As different 5G applications and services would require different demands for latency, bandwidth, density, Quality of Service (QoS), among other requirements, the dynamic orchestration utilizes different blueprints tailored for different network slice types whether it is addressing critical Machine Type Communication (MTC) and requirements for very low latency with continuously moving devices, or massive MTC (e.g., Smart Metering) with static devices with spurts of data in the form of small messages without latency constraints.

Having dynamic orchestration and its automation in place, this also allows service providers to support and monitor these different SLAs per individual network slice. Management functions provide life cycle functions to compose, create, modify, operate, scale, heal, and monitor, as well as, removing network slices. In certain cases, certain management capabilities can be “exposed” to customers for visibility and control of the network slice instance. Furthermore, as NSIs/NSSIs are successfully implemented and executed, the respective blueprints comprising the Key Performance Indicators (KPIs), SLAs, etc. could become an eco-system library collection of baseline templates to further aid zero-touch, automated time-to-market of new innovative services.

For quick market launch of new innovative services, dynamic orchestration should be a pre-integrated solution that supports specific business use cases, enabling service providers to launch services which leverage both physical and virtual network functions. By adopting subscription and policy profiles, such solution supports unique QoS profiles for each service or tenant e.g., MVNO. To reap the most benefits of dynamic orchestration, it is essential to have a comprehensive asset management solution to track asset utilization to timely adapt to new demands and requirements in real time. Enabling automation to the operational environment (to quickly respond to traffic and capacity trending, assess end-to-end (E2E) representation of a service, service outages, better forecasting and planning) depends on accurate inventory.

From orchestration perspective, every service/slice would get created with specific characteristics. 3GPP has different roles defined for Mobile Network Operator (MNO), Mobile Virtual Network Operator (MVNO), Mobile Virtual Network Enabler (MVNE), respectively. These required service characteristics (e.g., QoS) have to get applied during the service design and provisioning/activation.

To get the required QoS, there are multiple factors (like capacity, location, etc.) which have to be considered during the design. Then for each of the subnetwork slice/service the QoS parameters like bandwidth, latency, jitter, etc. would be applied. Prioritization of the traffic for the slice would have to be classified so that the routing devices can route the traffic with the required priority over the network to get the required latency. In turn, this requires the network to have clear multi-tenant capabilities; subsequently, the system needs to make available a Self-Service Portal and Self-Ordering APIs to MNO customers to allow self-ordering of their slices with SLAs and update them dynamically when needed.



**Figure 1 – Roles related to 5G networks and network slicing<sup>1</sup>**

Enterprises, MVNOs, or any other customer consuming the services from the MNOs can be considered a Communication Service Customer (CSC). An MVNE can share the CSC and some of the Communication Service Provider (CSP) depending on the sharing of areas or concerns based on the type of services (B2B, B2C, etc.). Some of the functions of management functions CSMF, NSMF and NSSMF ought to have accessibility via multi-tenant portals. A dynamic service orchestration would need to be able to differentiate the slices created for the customers of MNOs, distinguishing between Enterprises and MVNO customers.

End-to-end service orchestration interprets and translates service definition (Service Design) into configuration of resources (physical and virtualized) needed for service establishment. The configuration of resources may be for actual amount of resources or the policy of their allocation at a later time, when the service is activated. The E2E Service orchestration further triggers the components of the management and orchestration system (ETSI NFVO, VNF Manager, VIM, SDN Controller, legacy OSS) to dynamically apply the configuration of the required resources which for some resources may result in their actual allocation.

<sup>1</sup> As defined by 3GPP



With new re-engineered workflow automation, network slices can be instantiated, and subsequently modified, upgraded, scaled, or terminated in real time. Thus, 5G with network slicing will enable communication service providers to tailor their network on-demand to the needs of different industries, enterprises, and consumers and to rapidly develop new business models in co-operations with various industry verticals. This enables traditional communication service providers to evolve beyond traffic-based business models to become Digital Service Providers (DSPs). Enterprises can choose dedicated or shared slices with dynamic orchestration assuring the respective SLAs and KPIs are met, enabling monetization of the respective slices of services.

To fully exploit the benefit of optimized, on-demand created network slices, it is required that the slicing concept allows for efficient usage of common resources such as radio resources and infrastructure, and transport links such as fronthaul and backhaul. This is particularly challenging in the Radio Access Network (RAN), which is limited by the amount of available spectrum.

Similarly to the core network, RAN can be slice-aware and apply specific policies in the RAN to different users based on the service/slice type. The 5G RAN is capable of treating different slices uniquely and/or offering different services within multiple slices. It also supports protection mechanisms for slice isolation so that events (such as congestion) within one slice do not have a negative impact on another slice.

The isolation and independence of network slices enable the service provider to offer attractive new possibilities to industry, enterprise customers, or MVNOs. Each network slice can be used as an entirely self-sufficient and autonomous ‘container’ with all the functions and resources required for independent service. With this new NSaaS business model, customers can be granted visibility of their network slice, and then modify it to suit their changing needs, or create new network slices quickly to seize a new business opportunity. Compared to today’s 4G MVNO business model, the NSaaS business model is expected to be more powerful in terms of flexibility, automation, and the ability to customize the services delivered. For on-demand capacity increase needs, a NSaaS customer would be able to perform the required service changes themselves within the network slice itself. In certain business models, the NSaaS customer may even take full operational responsibility for this service, using in-slice management functions. As NSaaS evolves, Continuous Delivery and Deployment for software updates will play a key role in expanding possibilities and automation as quickly as possible through regular, incremental functionality upgrades.

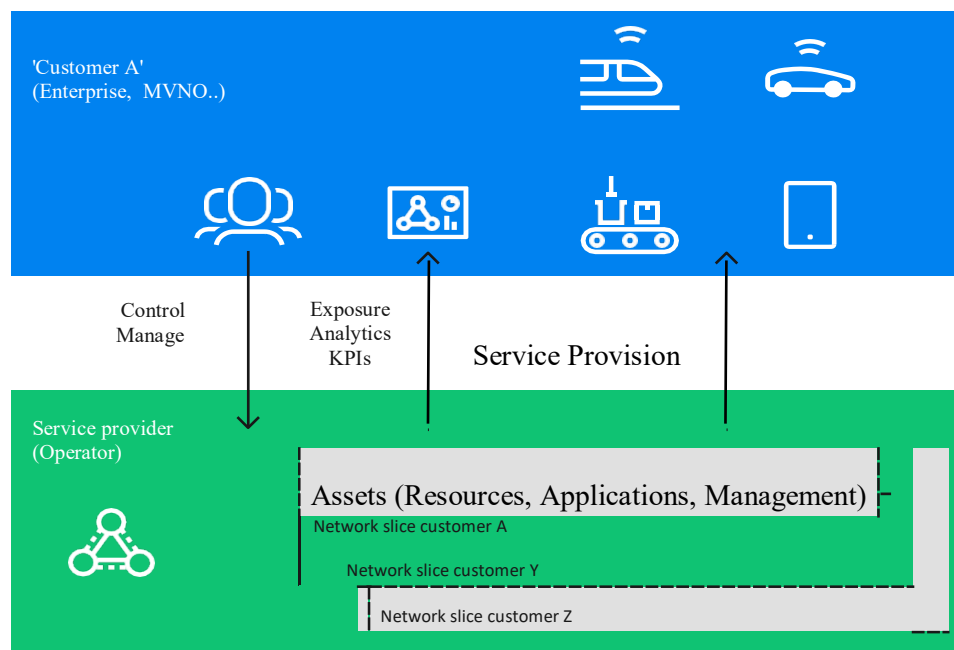
Making money from network slicing is very dependent on flexible and powerful business support systems, to enable the rapid creation of completely new services and offers. With the “Network as a Service” business model, new capabilities could mean providing capabilities to partners (e.g., MVNOs) to help them manage and monetize their relationship with their customers. Capabilities might include providing access to location information or usage analytics, or a more complete offering involving rating, charging, and billing for partners' services. All of these types of requirements can be supported by Business Support Systems (BSS) using adaptive logic to enable a platform for business innovation on a massive scale as is expected with NSaaS.

As 5G network slicing matures and evolves, it can open opportunities to embrace business models into more of an eco-system partnership where the traditional MNO and MVNO could facilitate individual aspects of services offered by eco system partners, e.g. private enterprises. This in turn will put a focus on monetization and service-based converged charging network function – where value of the offering could be a key aspect of differential pricing, rather than simple volume/usage-based pricing. The opportunity is for both the MNO and the MVNO to increase the “value” in the value chain, and smartly monetize the network with differential pricing that reflects the

“value of the bit”; consequently, a charging solution must be device-aware, network-aware, service & app-aware all at the same time to be able to properly reflect the value of the NSaaS.

Addressing charging capabilities, a convergent charging system for multiple lines of businesses, supporting network-convergence, service-convergence and subscriber-convergence should be introduced – designed to support real-time and offline charging for multiple lines of businesses of service providers’ environments. Its 5G charging capabilities would include:

- Slice-based charging
- DNN-based charging
- QoS flow-based charging
- Mobile phone location
- Connection-based charging



**Figure 2 – Providing customer control over network slices**

As requirements to distribute network resources to the edge, an Edge NFV cloud infrastructure (NFVi) solution aids the service providers with a compact platform, system-designed to run cloud native applications and virtual network functions managed by orchestration. Edge NFVi solutions shifted closer to the network edges would be used to optimize data traffic flow, addressing on-demand increased capacity needs for end-user services like HD live video streaming, Virtual Reality, etc. These solutions fit well in with the cable initiatives such as Head End Re-architected as a Data Center (HERD), Fiber Deeper, Distributed Access Architecture, and Generic Access Platform (GAP).

As it is envisaged, with more nodes distributed on the edges of the network (closer to the customer), it is important that functionality and management is software-defined with built-in flexibility to centrally adapt to new requirements and demands quickly. Edge NFVi computing provides execution resources (compute and storage) with sufficient connectivity (networking) at close proximity to the data sources,

typically within or at the boundary of access networks. It opens up opportunities for the service providers to distribute their resources to where demand exists enabling the most beneficial solutions for requirements such as low latency for VR/AR optimum user experience. To optimize such user experience, the Edge NFVi solution needs to be a distributed cloud solution – providing one execution environment for cloud application optimization across multiple sites, including required connectivity in between, managed as one solution and perceived as such by the applications. It is an end-to-end approach to edge computing allowing each individual use case to “decide” the specific edge location. Distributed cloud goes along with automated deployment of applications at just the right location in the network to optimize resource efficiency and user experience. For that reason, end-to-end orchestration for hybrid clouds is a key capability providing end-to-end management of networking, cloud infrastructure, and workload placement.

Dynamic orchestration, service-based converged charging, and Edge NFVi platforms are different ways to support multiple cable industry initiatives and leverage network function virtualization to offer more flexibility and customization of network operations. Through network slicing, these capabilities can be segmented to different customer groups to meet their specific use case requirements.

To support network slicing implementation in a multi-vendor environment with open interfaces and standardized architectures, re-engineered workflows and automation is needed to fully take advantage of new business opportunities and models. Having experience and skills in system integration and consultancy to execute and support such workflow automation initiatives and life cycle management efforts required for new service launch, software release upgrades, etc., could be a key differentiator for success.

Independent of access network (DOCSIS/HFC, 4G/5G/Private LTE, etc.), solutions and service expertise is available to support implementation and instantiation of virtual network slices for both enterprises and consumers, as well as providing unique SLAs and KPIs to manage service assurance.

## Conclusion

5G network slicing is a new paradigm, and it introduces a completely new way of specifying services and delivering them in a flexible, agile, and automated manner. For an MVNO, working with a partner that can provide a complete solution, including the orchestration, converged charging, and network edge capabilities needed to implement, streamline, optimize, and automate network slice creation, and then to enable new and emerging business opportunities quickly, is going to become crucial for future success.

## Abbreviations

3GPP	3rd Generation Partnership Project
4G	4 <sup>th</sup> Generation
5G	5 <sup>th</sup> Generation
AI	Artificial Intelligence
API	Application Programming Interface
AR	Augmented Reality
B2B	Business to Business
B2C	Business to Consumer
BSS	Business Support Systems
CSC	Communication Service Customer

CSMF	Communication Service Management Function
CSP	Communication Service Provider
DNN	Data Network Number
DOCSIS	Data Over Cable System Interface Standard
DSP	Digital Service Provider
E2E	End-To-End
eMBB	enhanced Mobile Broadband
ETSI NFVO	European Telecommunications Standards Institute NFV Orchestrator
GAP	Generic Access Platform
HERD	Head End Re-architected as a Data Center
HFC	Hybrid Fiber Coaxial
HD	High Definition
KPI	Key Performance Indicator
LTE	Long Term Evolution
MANO	Management And Network Orchestration
MNO	Mobile Network Operator
MTC	Machine Type Communication
NFV	Network Function Virtualization
NFVi	Network Function Virtualization infrastructure
NSI	Network Slice Instance
NSSI	Network Slice Subnet Instance
NSSMF	Network Slice Subnet Management Function
NSMF	Network Slice Management Function
MVNE	Mobile Virtual Network Enabler
MVNO	Mobile Virtual Network Operator
NSaaS	Network Slice as a Service
OSS	Operational Support System
RAN	Radio Access Network
SDN	Software Defined Network
SLA	Service Level Agreement
UE	User Equipment
URLLC	Ultra-Reliable and Low Latency Communications
VIM	Virtual Infrastructure Manager
VM	Virtual Machine
VNF	Virtual Network Function
VR	Virtual Reality
QoS	Quality of Service

## Bibliography & References

Management, Orchestration and Charging for 5G networks, By Thomas Tövinger, 3GPP SA5 Chairman  
[https://www.3gpp.org/news-events/1951-sa5\\_5g](https://www.3gpp.org/news-events/1951-sa5_5g)

# **Practical Considerations For Full Duplex Deployments In N+x Environments**

A Technical Paper prepared for SCTE•ISBE by

**Dr. Bill Wall**  
Principal  
Cox Communications  
6305 Peachtree-Dunwoody Rd.  
Atlanta, GA 30328  
404-269-7429  
Bill.wall@cox.com

**Michael Cooper**, Cox Communications

**David Job**, Cox Communications

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
An FDX Primer .....	4
Our Dilemma .....	5
Legacy Plant Modeling .....	5
FDX Amplifier Options .....	9
1. A traditional high split or ultra high split (UHS) amplifier with diplex filter cutoff in the FDX band; 204, 300, 396, or 492 MHz. ....	9
2. An amplifier using a triplexer design with a directionally switchable amplifier in all or part of the FDX band. ....	9
3. A true bi-directional FDX amplifier based on echo cancelation technology. ....	10
True EC FDX Amplifier Concepts.....	11
Echo Cancellation Requirements.....	12
The Case For FDD Operation .....	14
Conclusion .....	15
Abbreviations.....	15
Bibliography & References .....	16

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1. N+0 Plant with FDX .....	4
Figure 2 Upstream and Downstream Utilization over Time Based on Current CAGR.....	5
Figure 3. Amplifier Cascade used in the Analysis.....	6
Figure 4. Levels at POE Gateway .....	6
Figure 5. Tapped Feeder Used in the Analysis.....	7
Figure 6. Typical Node Leg.....	8
Figure 7. Upstream Signal Leakage Into Adjacent Amplifier Downstream.....	8
Figure 8. FDX Amplifier Concept Using Direction Switchable Amplifier.....	10
Figure 9. Basic Echo Cancelling FDX Amplifier .....	11
Figure 10. FDX Amplifier Concepts.....	12
Figure 11. Worst Case SNR vs. Echo Cancellation .....	13
Figure 12. FDX Scenario with 75 dB EC .....	14
Figure 13. FDD Scenario with 30 dB EC .....	14

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 Approximate Capacity with Traditional Diplexer .....	9

# Introduction

Last year at EXPO the concept of expanding Full-Duplex DOCSIS (FDX) beyond N+0 architectures was introduced using bi-directional echo canceling amplifiers. The expansion of FDX beyond N+0 architectures greatly expands the deployment potential of FDX and the symmetric gigabit services that it enables. As operators gather more experience in the building of N+0 plant, it has become apparent that construction time and costs are greater than initial estimates. This has resulted in more limited N+0 builds in targeted areas and a desire to pursue other methods to expand plant capacity, particularly upstream capacity. In response to this desire, the MSO community and CableLabs started two exploratory working groups, one on Extended Spectrum DOCSIS (ESD) and one on FDX amplifiers. The continuing work of these groups has led to the creation of DOCSIS 4.0, which will bring both FDX and ESD together in a single specification. This paper focuses on FDX and how it might be deployed in existing plant.

## An FDX Primer

Full-Duplex DOCSIS was designed to work in N+0 plant, that is no active amplifiers beyond the node. The node is assumed to be a Remote PHY (RPD) or Remote MACPHY (RMD) device with a single coax span of five or six taps. The FDX band of operation spans from 108 MHz to 684 MHz, divided into six upstream subbands of 96 MHz each and three downstream subbands of 192 MHz each. The FDX node transmits downstream and receives upstream on the same frequencies in the FDX band using echo cancelation techniques to remove downstream interference from the upstream receiver. The FDX modem however operates in frequency division duplex (FDD) mode, transmitting upstream and receiving downstream on different frequencies within the FDX band.

Within a given node, based on the isolation between taps, some modems can receive downstream with minimal interference on the same frequencies that other modems are transmitting upstream. Using a procedure called sounding, the CMTS core sorts modems into interference groups (IG's) such that modems in different interference groups will not interfere with each other, while those in the same interference group would. One or more interference groups are then assigned to a transmission group (TG). Each TG is then assigned to which subbands it can transmit and which it can receive through a message called a Resource Block Assignment (RBA). This effectively divides the node into two or more virtual nodes. See Figure (1) for an illustration.

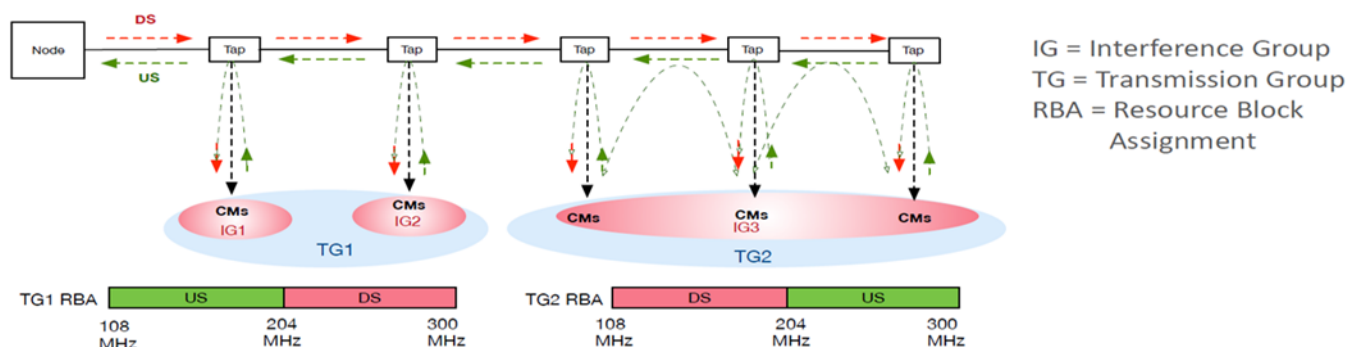



Figure 1. N+0 Plant with FDX



The primary benefit of FDX is not the re-use of frequencies, but the ability to dramatically increase upstream bandwidth as needed. Through an RBA change to a TG, it is possible to allow all of the FDX band to be used upstream; upstream capacity would be expanded to greater than 5 Gbps. RBA changes can be dynamic based on demand for capacity.

## Our Dilemma

The vast majority of our plant today is sub-split 1 GHz, typically N+5, averaging 400+ households passed (HHP). Node actions, either node splits or an N+0 conversion, are driven almost exclusively today based on upstream congestion. Modeling both upstream and downstream compounded annual growth rate (CAGR) on a node-by-node basis shows that a conversion to midsplit with a potential to offer 500 Mbps upstream virtually eliminates upstream congestion as a reason for a node action (See Figure (2)), and when coupled with a node split, pushes the next node action out 5+ years on average. Future node actions are then mostly driven by downstream congestion. This subsequent node action now could either be a node split, N+0 conversion, or ESD 1.8 GHz conversion. Midsplit however does not solve two issues, the first is the desire to widely offer symmetric gigabit services, the second would be a desire to proportionately increase upstream capacity if downstream capacity was increased through an ESD conversion.



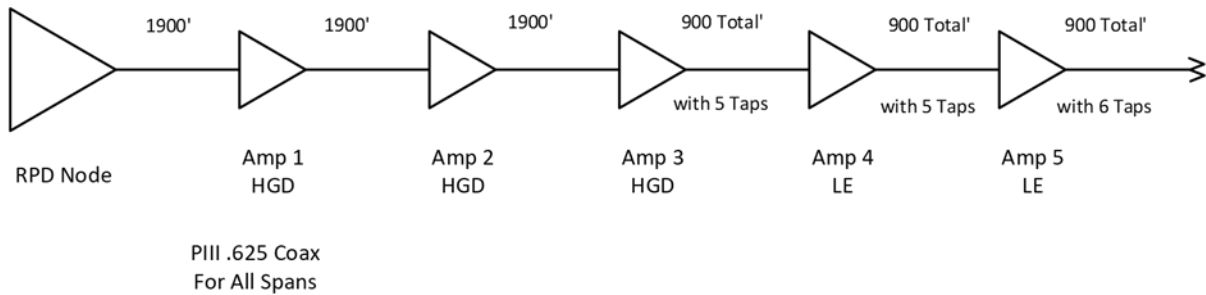
**Figure 2 Upstream and Downstream Utilization over Time Based on Current CAGR**

An ideal solution would allow FDX to operate in legacy N+x plants (N+5 or more) without changing plant topology by only replacing the node and all actives with FDX compatible products.

## Legacy Plant Modeling

We have performed network modeling associated with a “typical” legacy HFC cascade to approximate the RF performance that might be achieved with an expanded upstream (to 684 MHz) in conjunction with a 108-1000 MHz downstream. See Figure 3 for the configuration of the amplifier cascade.

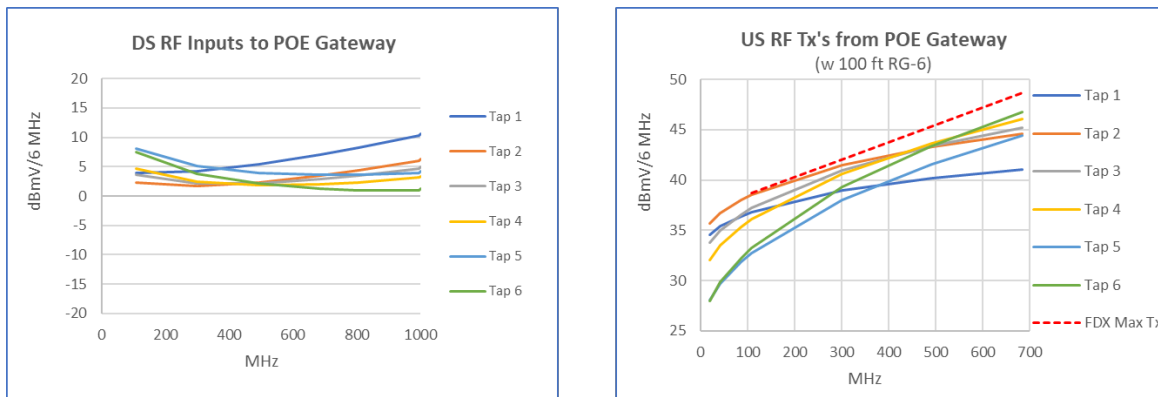
## Node + 5 Amplifier Cascade



**Figure 3. Amplifier Cascade used in the Analysis**

On our typical Node + 5 amplifier cascade, with 36 amplifiers in total fed off the node, we calculated the following performance:

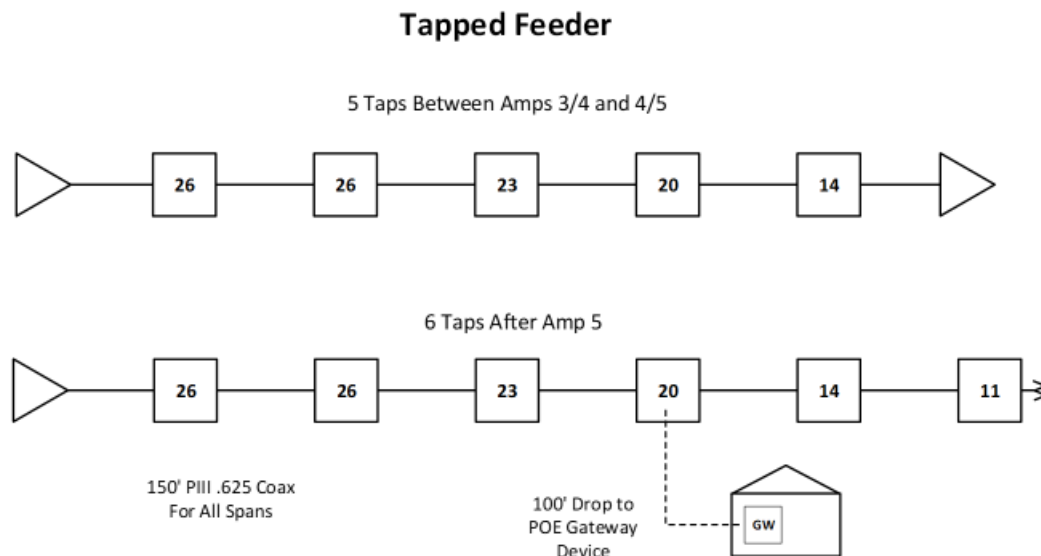
- Cumulative DS Composite Carrier to Noise (CCN) for the Node + 5 amp cascade = 41.6 dB (all ratios are relative to SC-QAM/OFDM channel power in a 6 MHz bandwidth)
- Cumulative US CCN for the 36 amplifiers plus the Node = 36.3 dB
- US Tx power for a Point of Entry Gateway device ranged from 27.9 to 35.5 dBmV/6 MHz at 20 MHz to 40.9 to 46.6 dBmV/6 MHz at 684 MHz. See Figure 4.



**Figure 4. Levels at POE Gateway**

For the modeled example the following assumptions were used:

- The node is an RPD node, with RPD module providing 55 dB composite intermodulation noise (CIN).
- Target RF output levels of the node and amplifiers are 32.8 dBmV/6 MHz at 108 MHz, and 46 dBmV/6 MHz at 999 MHz (for SC-QAM/OFDM).
- The first 3 amplifiers in cascade are “Express” multi-port amplifiers with their main output feeding directly to the next Express amplifier via non-tapped coax. The last 2 amplifiers are single port Line Extenders. Refer to the drawings of the amplifier cascade (Figure 3) and the tapped feeder line (Figure 5) for additional information.




**Figure 5. Tapped Feeder Used in the Analysis**


- Downstream amplifier output levels, internal tilts, and gains are typical of the 1 GHz types of amplifiers used in our networks.
- Upstream amplifier input/output levels and required gains were calculated assuming the losses expected at 684 MHz, with a target RF receive level at the RPD node port of 6 dBmV/6 MHz.
- Amplifier station noise figures were increased relative to legacy amplifier noise figures under the assumption that if a non-diplexed (FDD/FDX) amplifier with echo cancellation is used, additional input/output losses for internal downstream/upstream splitters/combiners would increase the noise figures.
- Downstream CIN was estimated based upon amplifier data sheet specifications for RF loading to 1 GHz, using what are typically known as high output GaN amplifiers (assuming we could use that type of downstream amplifier for this application, with downstream total composite power (TCP) of only 62.7 dBmV. Note that these are not the “super-high output” amplifiers used in 1.2 GHz N+0 applications.
- Upstream CIN was a rough approximation based upon the assumption that the amplifier stations would make use of a dual-stage upstream amplification stage, with 10 dB of tilt (via inter-stage equalization) introduced between the stages to lessen the output TCP of the 2<sup>nd</sup> stage.
- 100 ft RG-6 was used for drop loss from Tap output to POE Gateway input.
- Upstream 684 MHz EQ losses were modeled.

All of the modeling here was based on a 1 GHz plant; however extending the plant to higher frequency, either 1.2 GHz or 1.8 GHz, will not invalidate this analysis since we must keep current levels to support legacy equipment below 1GHz.

Figure 6 shows one leg of a typical N+5 node. Note the extensive use of splitters and couplers in the node. This introduces a significant issue for FDX in the formation of interference groups; the upstream output of one amplifier can couple across a splitter at a sufficient level to interfere with the downstream input to an adjacent amplifier if they are on the same frequency (See Figure 7). This basically means that modems connected through the second amplifier will be in the same interference group (IG) as modems off the first amplifier. In fact, by tracing through the potential interference paths, we find that all modems on one leg of the node will be in the same interference group.



**Figure 6. Typical Node Leg**



**Figure 7. Upstream Signal Leakage Into Adjacent Amplifier Downstream**

# FDX Amplifier Options

In order to implement FDX in an N+x environment, all actives must be replaced with FDX compatible devices. Three basic approaches for FDX amplifiers have been proposed:

## 1. A traditional high split or ultra high split (UHS) amplifier with diplex filter cutoff in the FDX band; 204, 300, 396, or 492 MHz.

This would allow FDX nodes and CPE to be used to provide extended upstream spectrum. Since all modems in a node leg in this example are in the same IG, simultaneous bi-directional use of the spectrum is not needed. We showed earlier that midsplit satisfies normal peak usage of a node, going to a higher split is needed to support higher billboard rates or to support ESD. Much of the time that additional upstream bandwidth will be unused. Further, the diplexer region grows proportionally with the upstream bandwidth, and all subtract from the downstream bandwidth. Figure 8 shows the approximate data rates available for each diplexer cutoff. The chart shows the potential upstream and downstream capacity assuming 2048 QAM D3.1 plus 32 D3.0 (256 QAM) carriers to fill the available spectrum downstream and 1024 QAM D3.1 upstream in the FDX band and four 6.4 MHz 64 QAM D3.0 carriers and one 1024 QAM D3.1 43 MHz BW block in the midsplit region. It also assumes that the plant was upgraded to the full 1.2 GHz FDX capability. Fractional D3.1 blocks indicate a block of less than 192 MHz. A 204 High Split can provide a limited 1 Gbps symmetrical service, but a 396 MHz UHS would be required to provide a 2 Gbps symmetrical service with a significant loss of DS capacity. It may be possible to have remotely switched diplexers that could switch between a lower split and a 396 UHS as demand requires.

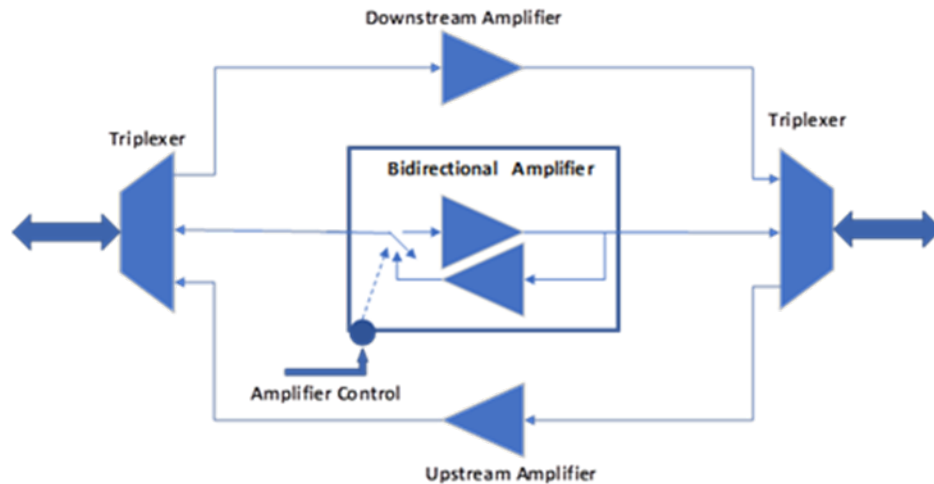
**Table 1 Approximate Capacity with Traditional Diplexer**

Upstream Freq (MHz)	FDX Blocks	US BW (Mbps)	Diplexer BW (MHz)	DS 3.1 Blocks	DS BW (Mbps)
Midsplit	0	450	23	4.7	9000
204	1	1250	54	4	7800
300	2	2000	60	3.5	7000
396	3	2800	80	2.9	6000
492	4	3600	96	2.3	5000
True FDX	6	5100	0	4.7	9000

## 2. An amplifier using a triplexer design with a directionally switchable amplifier in all or part of the FDX band.

This design proposed by CableLabs overcomes some of the limitations of the previous approach in that the selected portion of the FDX band can be remotely switched from downstream to upstream, perhaps with proper signaling to form a time division mode of operation. This potentially introduces latency in that RBA changes are not instantaneous. Early discussions in the FDX WG considered time-division duplex (TDD) as a primary solution path, but the group formed a consensus against this path in part for

this reason. However, the portion of the FDX band used is still fixed by the triplexer design and substantial bandwidth is still lost in the upper triplexer split. Figure 8 illustrates this design.

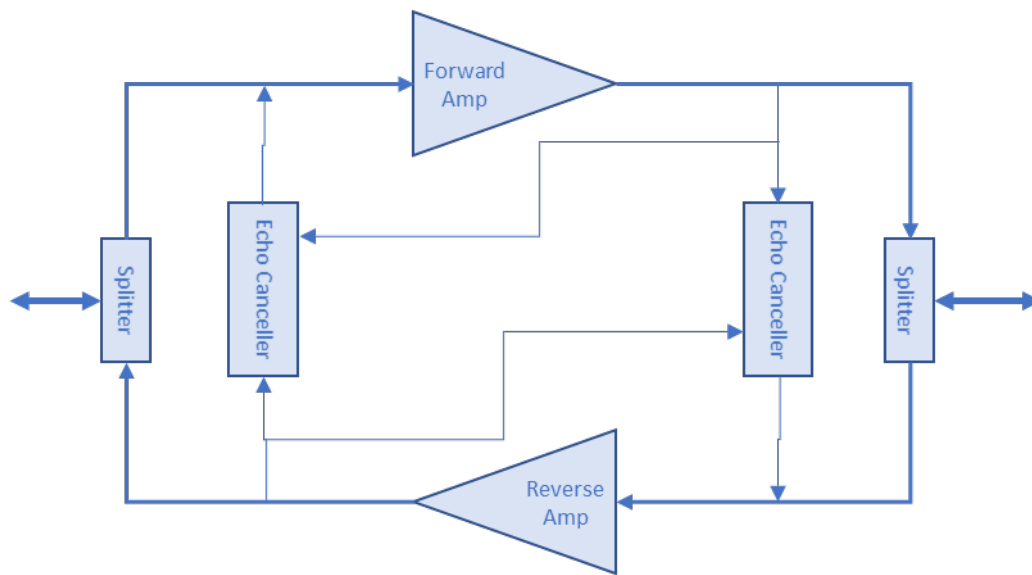


**Figure 8. FDX Amplifier Concept Using Direction Switchable Amplifier**

### **3. A true bi-directional FDX amplifier based on echo cancellation technology.**

This type of amplifier uses the same echo cancellation (EC) technology that is used in the FDX RPD, except that there are two instances of EC used, one for the forward path and one for the reverse path as illustrated in Figure 9. This amplifier offers true bi-directional amplification throughout the FDX band. There is no loss of spectrum due to a diplexer region. There are however challenges in an EC FDX amplifier. EC is not perfect, there will be residual EC “noise” that will degrade the overall performance of the amplifier. This is a particular worry in the upstream direction where noise funneling from the cascade could significantly worsen upstream MER performance. A second major concern is amplifier stability. The loop gain around the amplifier, including any echoes from upstream and downstream components, must be less than one (0 dB) or the amplifier will oscillate. In order for that not to occur, the EC’s must be trained prior to the amplifiers becoming operational.

The basic operation of EC is as follows. Consider the downstream port of the amplifier; the echo canceller samples the downstream amplifier in the FDX band to provide a reference signal and also samples the output of the upstream amplifier for training. Using a convolution process, the EC constructs a model of the leakage and echoes coming from the amplifier components and other components downstream. Using that model, it generates an out-of-phase replica of that echo that is combined with the input to the upstream amplifier, canceling the echo. The EC constantly monitors the output of the upstream amplifier and adjusts the model to minimize the resultant echo. The EC on the reverse port operates in a similar manner.



**Figure 9. Basic Echo Cancelling FDX Amplifier**

Current proponents for EC based amplifiers have proposed them to extend N+0 plant to N+1 or N+2. The focus of this paper is to understand the feasibility of using EC amplifiers in existing N+5 or higher plant.

## True EC FDX Amplifier Concepts

Two basic concepts for EC amplifiers have been explored as shown in Figure 10. The first is an analog implementation, in that the amplifier paths both forward and reverse are purely analog as in today's diplexed amplifiers. Additionally, the actual echo cancellation happens in the analog domain as well, in a directional coupler. Creation of the echo model and out-of-phase replica happens in a Digital Signal Processor (DSP) and its associated D/A and A/D converters. Since echo cancellation is only needed in the FDX band, the digital converters and DSP only need to work below 700 MHz. However, since cancellation is done only in the analog domain there will be a limit to the amount of cancellation achieved.

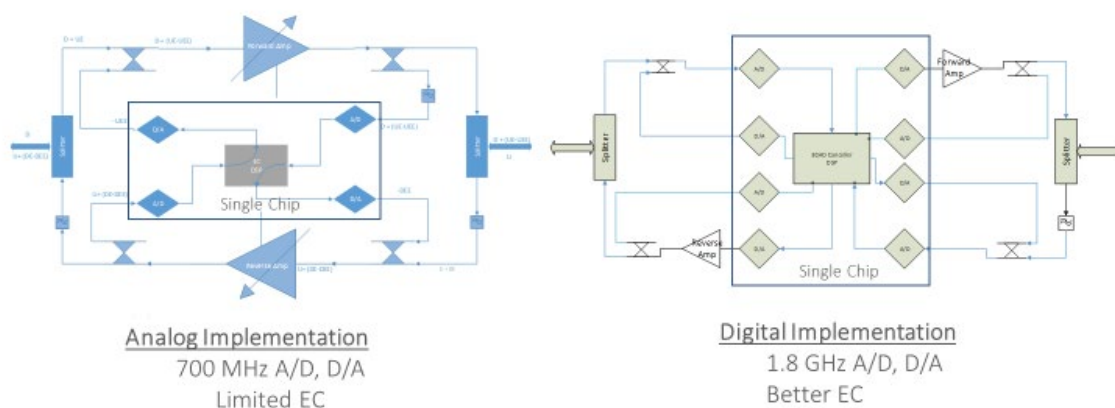
The second concept is a more digital approach; the inputs of both the forward and reverse amplifiers are digitized and then converted back to analog before amplification. Like the analog approach, there is a first stage of analog cancellation to ensure that the A/D converters are not saturated by the echo and leakage levels. A second stage of digital cancellation in the DSP follows. Significantly higher levels of EC are obtained by this two-stage approach. Current designs for FDX nodes use this two-stage approach. Once both forward and reverse paths are digitized, there is significant flexibility in the processing that can take place such as gain control and equalization, or other functions such as upstream squelch to minimize noise funneling. A disadvantage of this approach over the analog is that the A/D and D/A converters have to digitize the full band downstream. For a full DOCSIS 4.0 compliant amplifier, this could mean digitizing the full band to 1.8 GHz with the inclusion of Extended Spectrum DOCSIS (ESD). To be compliant to the current DOCSIS 4.0 specification, which includes only the existing Full Duplex specifications at this time, this could be relaxed to 1.2 GHz.

Both approaches have issues that must be solved. As mentioned earlier, the loop gain around the forward and reverse paths must be less than one or the amplifier will oscillate. We will see in the next section that in order to achieve that requirement the EC's must be trained individually prior to the loop being closed.



While there are ways to do this in both designs, the greater flexibility of the digital design makes this more straightforward. Another problem with both designs is the zero-time echo. That is the echo associated with the leakage across the output splitter (or more likely, directional coupler (DC)) and from the connector on the amplifier occurs at almost zero time, so processing through the DSP and converter chain has to occur in near zero time in order to cancel that echo. A key issue is the level of EC needed to insure the SNR desired, both upstream and downstream which will be addressed in the next section.

Finally, for an EC amplifier to be viable, it's advantages over a diplexer-based amplifier must outweigh any cost or power disadvantage that it has introduced. In either approach, for both cost and power reasons, the DSP and converter functionality needs to be integrated into a single custom ASIC. FPGA implementations most likely would not meet cost or power requirements. Here the analog approach would seem to have the advantage in both cost and power. It has half the number of converters and operates at a significantly lower frequency, and since the main path is not digitized it's converters may not need to be as accurate. However, the added functionality of the digital approach including remote gain and equalization control, proactive network management (PNM) functions such as full band capture in both direction, and knowledge of return loss profiles in both directions, may outweigh the cost advantages of the analog approach.



**Figure 10. FDX Amplifier Concepts**


## Echo Cancellation Requirements

For this exercise, High Gain Dual amplifiers have been analyzed since they contribute most heavily to both upstream and downstream noise and have higher gains. Similar requirements will apply to line extenders. The objective is to determine the level of echo cancellation required in order that the residual echo will not degrade the overall carrier to composite noise (CCN) of the node such that 2048 QAM OFDM will work downstream and 1024 QAM OFDMA will work upstream. Using the nominal input and output level of both forward and reverse paths, estimated CIN and carrier to thermal noise (CTN) performance of the amplifiers, plus a maximum reflected energy (echo) from the plant, the resultant SNR of combining thermal noise plus residual echo “noise” was calculated vs varying degrees of echo cancellation. For this calculation, the assumed value of maximum echo was -15 dB in both the forward and reverse directions from the amplifier.

Two scenarios were examined, the first was true FDX where the node is sending downstream on the same frequencies it is receiving upstream, and the second was an FDD mode where there is not simultaneous use of forward and reverse spectrum. In the first case, the full output power of the amplifier must be



canceled such that the SNR of the input is not significantly degraded. In the second case the noise plus distortion products of the amplifier must be cancelled to that level. In this second case, active transmission only needs to be canceled to the point that it doesn't significantly affect the operating point of the amplifier. Figure 11 shows the worst-case calculated SNR for both upstream and downstream directions. Both the FDX and FDD scenarios are compared. Worst case performance in both scenarios occurred at the upper end of the FDX band.




**Figure 11. Worst Case SNR vs. Echo Cancellation**


As can be seen, there is a dramatic difference in the degree of echo cancellation required in the two scenarios. The “knee” in the curve in the FDX scenario for both upstream and downstream EC’s is about 75 dB. It is not known if this level of EC can be reliably achieved in an amplifier or if this level of isolation could be achieved in a single chip implementation. For the FDD case, both the upstream and downstream ports only require about 30 dB of echo cancellation. This value seems much more achievable.

A more detailed look at the full FDX case is shown in figure 12. Shown for a 75 dB EC, the upstream SNR and downstream SNR are plotted. The slight decrease in SNR versus frequency is due to the increasing power output with frequency and it’s residual after EC. In the FDD scenario Figure 13 shows a similar result. Here the EC is 30 dB in each direction. An additional parameter is plotted here, the transmit leakage into the input side in of both the forward and reverse amplifiers. With significantly less EC than in the FDX scenario, the leakage of the active downstream and upstream paths into the opposite input path could contribute to the TCP of that amplifier. However, this result shows that at any given frequency that leakage is lower than the normal input signal and will not significantly contribute to the TCP of the amplifier.

In both cases the loop gain is well below 0 dB. The highest loop gain is in the FDD case where the worst-case loop gain is -28 dB. This will result in a low level ringing of the loop that will appear as low-level echo, easily handled by OFDM or the SC-QAM equalizer.



**Figure 12. FDX Scenario with 75 dB EC**



**Figure 13. FDD Scenario with 30 dB EC**

## The Case For FDD Operation

As shown earlier, each leg of a node in Figure 6 style N+x forms a single interference group, meaning all modems in that leg must be assigned the same FDX subbands for both upstream and downstream, operating in an FDD mode. If the RPD also operates in FDD mode, that is it will not transmit downstream on the FDX subbands assigned for upstream, then amplifiers with reduced EC capability can be used. There is no loss of capability in that leg since it is a single interference group. Such a mode, “static” FDX, is supported in DOCSIS 4.0.

It is possible to operate a node in full FDX, with each of up to four legs of a node being independent IG’s and TG’s. Modems on each leg now will be sharing the non-FDX spectrum so overall capacity of a leg will be reduced compared to each leg operating independently in FDD mode. Further, to operate in full FDX mode, each amplifier will have to support the much higher EC requirements. The cost of having separate RPD modules for each leg should be compared to the total cost of upgrading all actives, the potential of higher cost for the higher level of EC required to operate in full FDX mode, and the loss of capacity of having a single RPD in the node.

It is potentially possible to operate in a “dynamic” FDD mode where the upstream/downstream capacities are changed through an RBA message. Such changes could happen at millisecond time scales and react to demand requests from modems. This could allow for optimization of capacity utilization where both upstream and downstream capacities are scaled for normal busy hour with reserve bandwidth to support “Tmax” (maximum speed offered) billboard rates now shared between upstream and downstream, available on demand.

Dynamic FDD will likely require a specification enhancement. Current FDX specifications anticipate that the RPD, once provisioned, will transmit continuously downstream on the active FDX subbands. However, to operate in FDD mode, the RPD will need to mute downstream transmissions on subbands that are upstream. This will require an RBA-like message be sent to the RPD from the CCAP core and timed to be coincident with the RBA message to the modems.

There are a number of advantages of operating an FDX-based system in FDD mode with EC based amplifiers.

- First, it allows the operator to adjust his upstream/downstream splits as needed without touching the plant, only a configuration change at the CCAP core for static operation, or an RBA change for dynamic operation.
- There is no bandwidth lost in a diplexer region; this could amount to as much as a gigabit/sec of throughput for the higher splits.
- Legacy OOB operation will just pass through and is not affected.
- FDD operation reduces the EC requirements for both node and amplifiers. The EC requirements for full FDX operation in N+5 may not be achievable.
- The use of FDD rather full FDX simplifies operation in that sounding to establish IG's is not required.

## Conclusion

A potential path to the use of FDX in N+X plant has been described. Though both technical and cost challenges are present in the development of EC based amplifiers, they offer a number of advantages over fixed diplexer solutions. The use of dynamic or static FDD is shown to be a preferable solution over full FDX in an N+x plant, both from reduced EC requirements on the amplifiers as well as from higher overall capacity of the node. As the industry moves forward with DOCSIS 4.0 and extended spectrum, we urge the industry to push forward the development and the use of bi-directional EC based amplifier technology.

## Abbreviations

A/D	Analog to digital
BW	Bandwidth
CAGR	Compounded annual growth rate
CCN	Carrier to composite noise ratio
CIN	Carrier to intermodulation noise ratio
CMTS	Cable modem termination system
CPE	Customer premis equipment
CTN	Carrier to thermal noise ratio
D/A	Digital to analog
dB	Decibel
dBmV	Decibel relative to one millivolt
DC	Directional Coupler
DOCSIS	Data over Cable system interface specification
DSP	Digital signal processor
EC	Echo cancellation

ESD	Extended spectrum DOCSIS
FDD	Frequency division duplex
FDX	Full duplex (DOCSIS)
GaN	Galium Nitride
GHz	Giga Hertz ( $10^9$ Hz)
HFC	Hybrid fiber-coax
HHP	Households passed
Hz	Hertz
IG	Intrference group(s)
MAC	Media access control
MSO	Multi system operator
N+0	Node plus zero actives
N+x	Node plus “x” actives
OFDM	Orthogonal frequency division modulation
PHY	Physical layer interface
RBA	Resource block assignment
RMD	Remote MAC PHY device
RPD	Remote PHY device
SC-QAM	Single carrier quadrature amplitude modulation
SCTE	Society of Cable Telecommunications Engineers
TCP	Total composite power
TG	Tansmission group
UHS	Ultra High Split

## Bibliography & References

*FDX DOCSIS Line Extender: Deploying FDX DOCSIS Beyond N+0*, John T Chapman & Hang Jin, Cisco Systems; SCTE ISBE and NCTA 2018 Fall technical Forum

*DOCSIS 4.0 Physical Layer Specification*, CM-SP-PHYv4.0-I01-190815, CableLabs, August 15, 2019, <https://apps.cablelabs.com/specification/>

*DOCSIS 4.0 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv4.0-I01-190815 CableLabs, August 15, 2019, <https://apps.cablelabs.com/specification/>

# Telecom Argentina

## Transport Network Evolution For Future Services

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

**Esteban Poggio**

Cloud Infrastructure Strategy Architect  
Telecom Argentina  
Aguero 2392 - C1425EHZ , CABA - Argentina  
+541155304756  
epoggio@teco.com.ar

**Miguel Masache Ojeda**

Cloud Infrastructure Strategy Architect  
Telecom Argentina  
Aguero 2392 - C1425EHZ , CABA - Argentina  
+541140015669  
mmoqueda@teco.com.ar

**Furquan Ansari**

Partner  
Bell Labs Consulting  
600 Mountain Avenue, Murray Hill, NJ 07974  
+1-732-768-5876  
furquan.ansari@bell-labs-consulting.com

**Ben Tang**

Principal Consultant  
Bell Labs Consulting  
4573 Kentucky Drive, Plano, TX 75024  
+1-469-910-3698  
ben.tang@bell-labs-consulting.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Abstract.....	3
Content .....	3
1. Overview Of Telecom Argentina.....	3
2. Strategic Objectives And Drivers .....	4
2.1. Strategic Objectives.....	4
2.2. Drivers For Network Architecture Change .....	5
3. Backbone Network Challenges & Architecture Evolution.....	5
3.1. Target Topology Selection .....	6
3.2. IP Network.....	8
3.3. Optical Network .....	12
3.4. CDN Strategy .....	14
3.5. End-To-End Architecture (2023+) .....	15
3.6. SDN Use Cases .....	18
4. Summary .....	19
Abbreviations.....	20
Bibliography & References .....	21

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Presentation of the new Telecom Argentina. ....	4
Figure 2 - Former companies existing topologies. ....	6
Figure 3 - Candidate topologies. ....	7
Figure 4 - FMO selection process. ....	8
Figure 5 – FMO IP Topology.....	10
Figure 6 - New topology by layers and main nodes geographical locations.....	11
Figure 7 - Options to solve G.653 fiber capacity limitations.....	13
Figure 8 – Failures supported considering optical and IP layers behavior.....	14
Figure 9 - Traffic distribution in the Telecom Argentina Network. ....	14
Figure 10 - CDN strategy for content distribution.....	15
Figure 11 - End-to-end architecture overview.....	16
Figure 12: FMO vs. PMO comparison. ....	17
Figure 13 - Optimization of the links use to Tier I.....	18
Figure 14 - Traffic engineering with SDN controller. ....	19

# Abstract

CSPs<sup>1</sup> all over the world are evolving their networks to be future-ready as the march towards 5G<sup>2</sup> continues. The new Telecom Argentina, formed through the merger of a cable and a telco companies, is not only undertaking the complex merger and consolidation of the two disparate and diverse networks, but at the same time also transforming to a unified and converged next-generation architecture that can support all future needs.

This paper addresses the transformation of the IP<sup>3</sup> and optical Core/Backbone architecture and technology, considering not only an optimized and flexible topology from a TCO<sup>4</sup> and current/future services point of view, but also features and functionalities this network has to support to meet the aforementioned objectives for the next 5 years and beyond. It addresses some of the key challenges, drivers, and specific characteristics such as fiber plant locations, diversity and resiliency requirements, traffic patterns, and densities, etc. that must be taken into consideration to develop a candidate set of network topologies for the given geography and demography. Both qualitative and quantitative criteria are then applied to select the right future architecture. It also provides rules that can be applied to determine optimal conditions to adapt this architecture to future demands with minimal impact, i.e. to create a dynamically adaptive network that can support any future, on-demand services.

The work will further elaborate on some of the salient points of this new architecture - including traffic and capacity optimization, content distribution strategies, network function distribution, DC<sup>5</sup>/edge cloud distribution, IGP<sup>6</sup> strategies, optical features, etc. to ensure an economically and operationally efficient backbone. It will discuss the incorporation and use cases of new technologies such as SR<sup>7</sup>, SDN<sup>8</sup>, NFV<sup>9</sup>, etc. as key enablers to truly achieve this target flexible network able to support next-generation services.

# Content

## 1. Overview Of Telecom Argentina

The new Telecom Argentina is the result of the merger between two companies, Cablevision Argentina, which was a cable company, and Telecom Argentina itself, a traditional telco and one of the most important telecommunications companies in the country.

The new company has approximately 30 million customers in Argentina, distributed as follows.

- 18.6 million mobile subscribers
- 4.1 million fixed broadband subscribers
- 3.5 million TV subscribers

---

<sup>1</sup> Communications Service Providers.

<sup>2</sup> 5th Generation Networks.

<sup>3</sup> Internet Protocol.

<sup>4</sup> Total Cost of Ownership.

<sup>5</sup> Datacenter.

<sup>6</sup> Interior Gateway Protocol.

<sup>7</sup> Segment Routing.

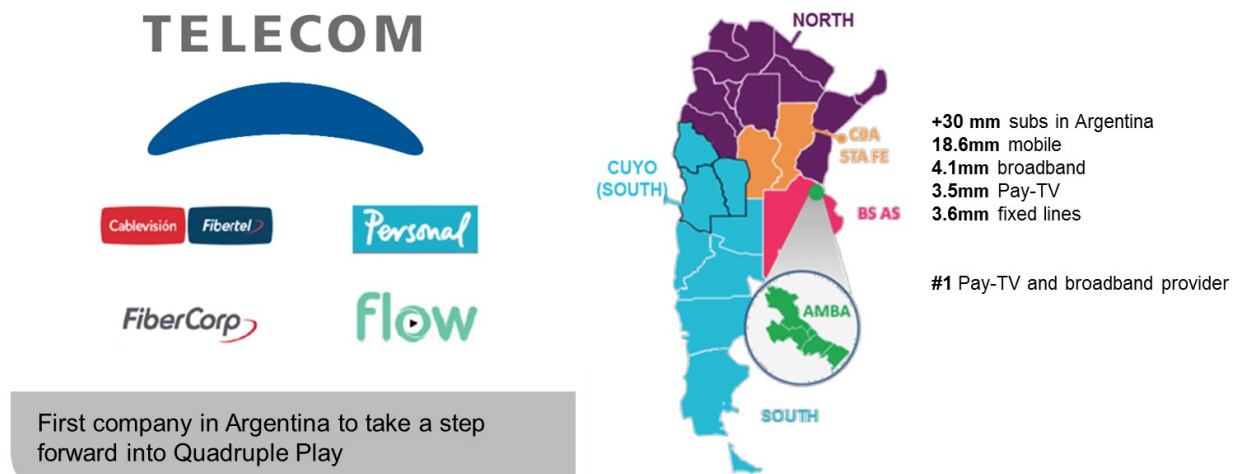
<sup>8</sup> Software Defined Networks.

<sup>9</sup> Network Functions Virtualization.

- 3.6 million fixed telephone lines

These numbers clearly demonstrate the importance and leadership position of new Telecom Argentina and sets itself as the most important company to offer 4Play<sup>10</sup> services (see Figure 1)

## Telecom Argentina Overview



**Figure 1 - Presentation of the new Telecom Argentina.**

## 2. Strategic Objectives And Drivers

### 2.1. Strategic Objectives

As part of the merger process, the new Telecom Argentina defined strategic objectives at both business and technical levels across the combined company, including for the Core/Backbone network. The strategic technical objectives were defined to meet the strategic business objectives, in such a way that they were coherent with and complimentary to each other.

The goal of the merger is to make (and retain) the new Telecom Argentina as the leading CSP in Argentina. To achieve this, the following strategic business objectives are defined:

- **Market Leadership:** Maintain the leadership in the fixed broadband service, increase the market share of the mobile service and corporate services with diversified solutions.
- **User Experience:** Enhance this aspect to increase the ARPU<sup>11</sup>, to evolve the FLOW<sup>12</sup> offer, leveraging the granularity of fixed and mobile access to offer bundled service packages.
- **Enhance Synergies:** Optimize investments (avoiding overlaps) through integration of the Core/Backbone networks. The execution of the integration could last a period of approximately 3 years.

<sup>10</sup> 4Play - Quad Play Products - Offers products that package Internet, Video, Mobility and Fixed Telephony.

<sup>11</sup> Average Revenue Per User.

<sup>12</sup> Telecom Argentina IPTV Services and Video On-Demand.



These strategic business objectives are translated into strategic technical objectives for the Core/Backbone network in the form of a consolidated, evolved, future-proof Core/Backbone network that has a flexible and scalable architecture that allows for growth of current services, deployment of futures services and has ability to adopt new technologies on the 5G path. It should also be simple to plan and operate, distributing content and traffic efficiently to improve the user experience and optimize the deployment of infrastructure, capacity and resources.

## 2.2. Drivers For Network Architecture Change

In order to meet the key business objectives defined in the previous section, the need arose to develop a single, consolidated Core/Backbone network for the merged company that would be sufficiently different from either of the existing former company networks (known as PMO<sup>13</sup>). The consolidation was meant to retain and leverage the strengths of each of the former networks. The high-level strategic objectives were further evolved into the following requirements for the Core/Backbone network:

- Minimize the TCO: In order to maintain the market leadership position, it is imperative that the network is run as cost efficiently as possible. Network TCO is a major challenge mainly due to the aggressive traffic CAGR<sup>14</sup> values that are being observed and the demand that the future new services could generate.
- Improve Time to Market (TTM<sup>15</sup>): Create a flexible and highly adaptable network architecture that will increase revenues by accelerating the service creation and instantiation with just-in-time delivery.
- Minimize latency: Based on service specific requirements, especially those driven by 5G (for example, IoT<sup>16</sup>).
- Assure resilience: Support fault conditions minimizing possible impact on services and customers.

In order to cover these requirements, e.g. to adapt services and product offers dynamically (a classic example is being able to generate bandwidth on demand), the flexibility, agility and simplicity of the future architecture are fundamental properties that should be taken into account.. The aforementioned architecture must adapt to the traffic matrix and its possible changes in a scalable manner. It should allow for designing optimal routing, fiber layout, connection schemes, operations and multilayer redundancies with coordinated work between the IP and optical layers. It should avoid lock-in by suppliers/vendors and improve innovation cycles of architectural components.

## 3. Backbone Network Challenges & Architecture Evolution

As mentioned, the architecture design of the new Telecom Argentina network should not only consider unifying the networks of both former companies but also evolve them by taking into account the combined traffic growth and profiles. This new architecture will be called the FMO<sup>17</sup>.

On one hand, from the point of view of the IP layer there are a series of challenges to consider, including resolving possible overlap of IP addressing (mainly private) when unifying the IGP, consolidating into a

---

<sup>13</sup> Present Mode of Opeation.

<sup>14</sup> Compound Annual Growth Rate.

<sup>15</sup> Time to Market.

<sup>16</sup> Internet of Things.

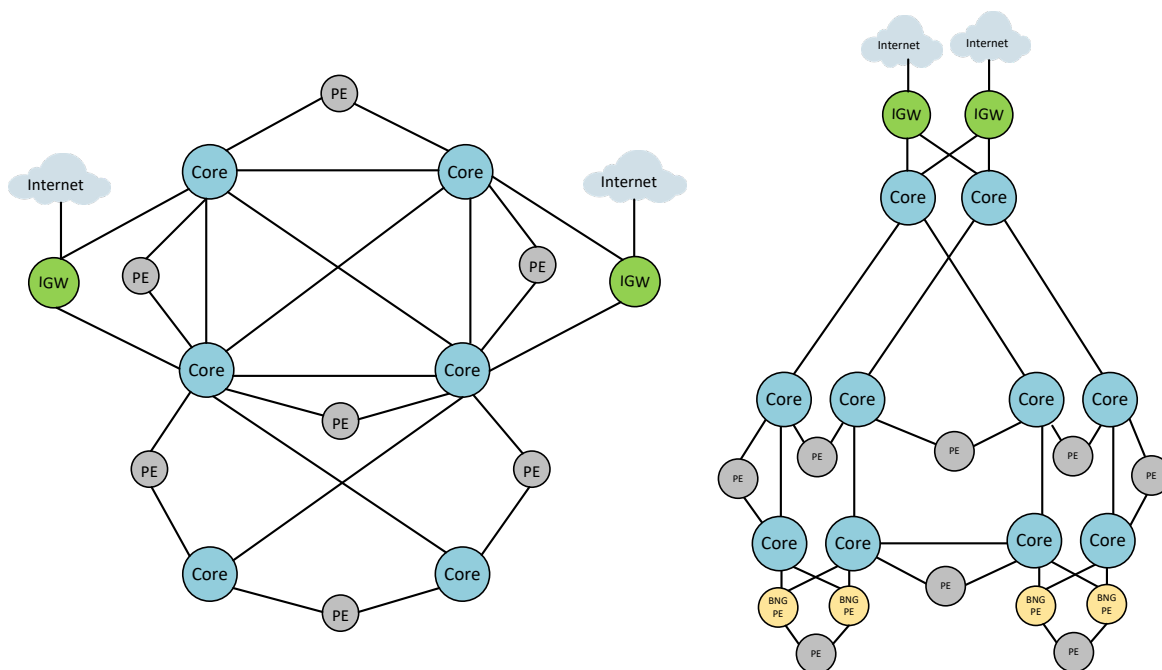
<sup>17</sup> Future Mode of Operation.

single ASN<sup>18</sup>, connections to Internet providers and route reflectors, unifying QoS<sup>19</sup>, provisioning and monitoring systems, etc.

On the other hand, from the optical layer point of view, there are different challenges to consider, such as diversity of optical platforms, partitioning of the optical network into vendor-specific networks, type and quality of the optical fiber deployed by the two companies (G.653, G.655 and G.652), etc. It is therefore necessary to analyze (not in this first stage but at a future point) the optical domains individually to allow for their optimization and also identify a strategy for the deployed optical fiber that will solve possible limitations of bandwidth, either due to the conditions inherent to certain types of fiber or to the distances of the links based on the transmission rates required.

### 3.1. Target Topology Selection

This section describes the procedure for selecting the FMO. As a starting point, the existing topologies of each former company prior to the merger should be considered (see Figure 2).



**Figure 2 - Former companies existing topologies.**

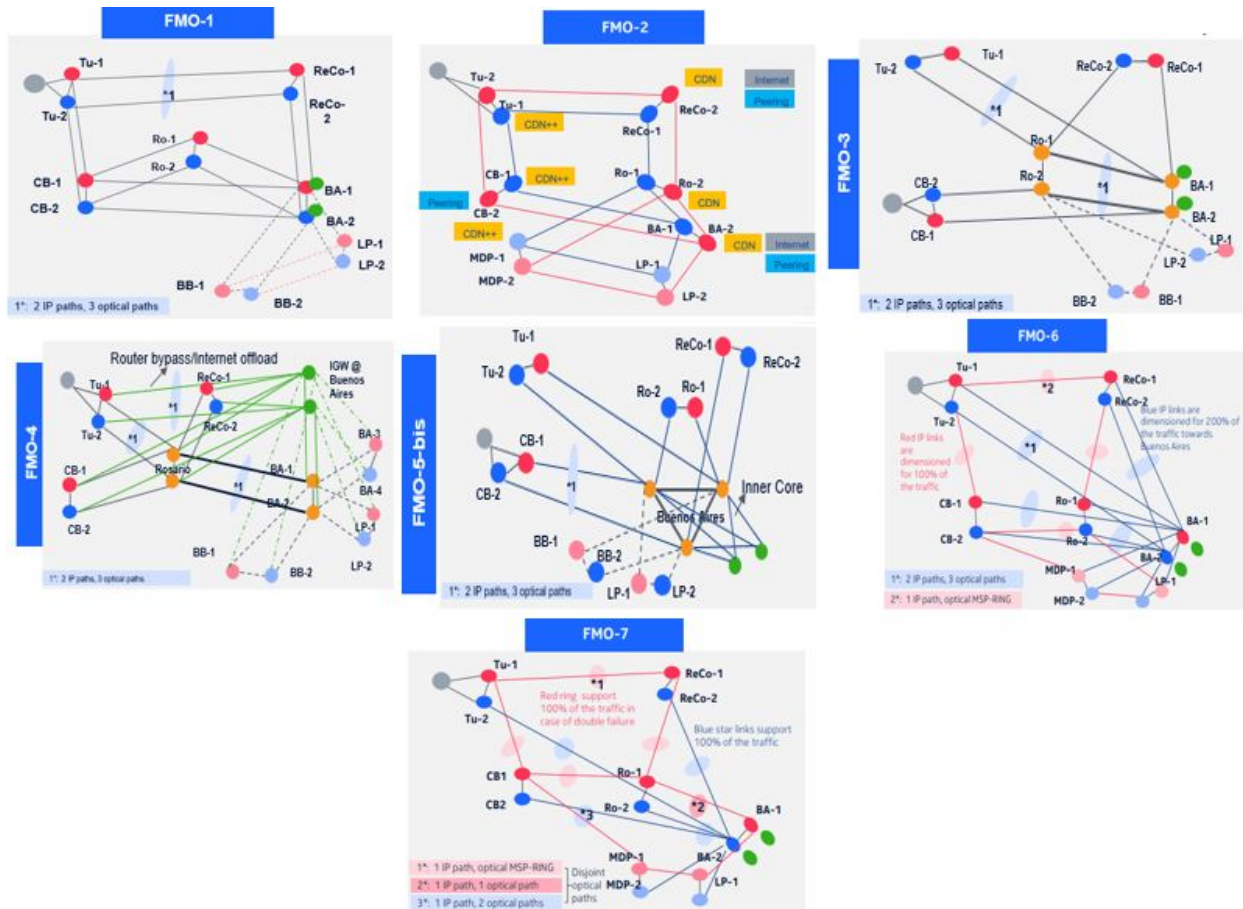
As part of the process, different topologies for the FMO selection were evaluated. Various candidate topologies were proposed, taking into account the subscriber base, traffic profiles and future traffic growth. From them, 7 candidates were selected (see Figure 3)

Figure 3 These candidate topologies, called FMO1 to FMO7, covered all the alternatives proposed by the working group for further analysis. These topologies represent different configurations of polygons with diverse geographic locations of different devices and numbers of Inner-core and Outer-core, to star

<sup>18</sup> Autonomous System Number.

<sup>19</sup> Quality of Service.

topologies, ring and combinations thereof. It also includes potential optimization using optical bypass, where possible.



**Figure 3 - Candidate topologies.**

These candidate FMOs were analyzed and scored based on various evaluation criteria, each with a specific metric and weight.

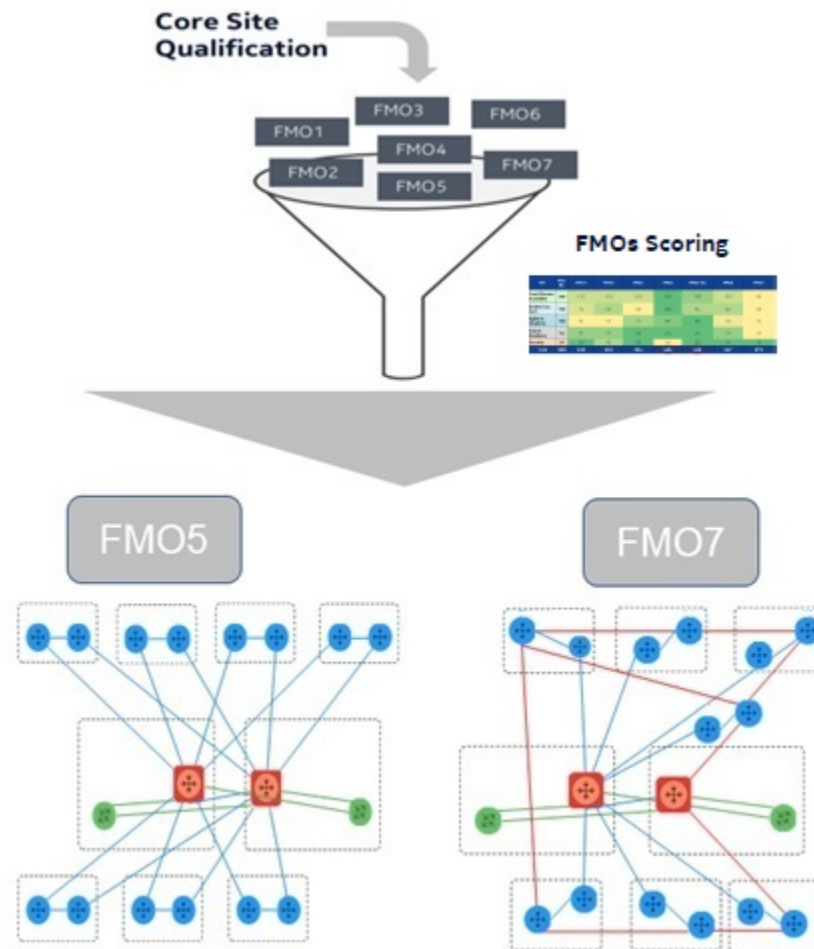
- Cost Efficiency
- E2E<sup>20</sup> Latency
- Resiliency
- Architecture Flexibility
- Operational Simplicity
- Simplicity of Migration
- Future Evolution Capacity
- Security

A qualitative analysis was accomplished based on the proposed weighting of all criteria for all the candidate FMOs. This qualitative analysis resulted in two finalist FMO architectures (at this point, in

<sup>20</sup> End-To-End.

addition to the topologies, different options of functionalities, technologies, etc. were also considered and analyzed). This procedure is outlined in Figure 4.

Finally, for the two final candidate FMO architectures, a detailed quantitative analysis was performed that included traffic routing and failure simulations using network modeling tools and took into account other technical aspects, including flexibility and change impacts, etc. The quantitative analysis resulted in a 5-year TCO calculation for both candidate FMO architectures. Based on these results, the definitive one was selected. The selected architecture is the one identified as FMO5.



**Figure 4 - FMO selection process.**

### 3.2. IP Network

As a result of the process described in section 3.1, the target architecture for the FMO, in the IP layer, is shown in Figure 5.

This architecture defines a star configuration where there are two main nodes called Inner Cores (as star centers) deployed in the AMBA<sup>21</sup> region that connect the co-located Internet nodes (IGW<sup>22</sup>).

In Argentina, connections to Internet providers are only available in AMBA. This leads to the conclusion that any content that can not be served from within our network will arrive to AMBA (that is, there is a star-oriented traffic matrix with traffic flows emanating/terminating from center of star located in AMBA). This traffic distribution characteristic and profile was critical in developing the appropriate topology.

On the other hand, there are “n” pairs of nodes called Outer Cores connected to each other and both connected in a dual-home configuration (figure of a “U”) to the Inner Cores. These Outer Cores aggregate the Edge PE<sup>23</sup> nodes from the different defined regions.

According to the previously described traffic patterns and the demographic and economic particularities of Argentina, the need for 2 pairs of Outer Cores in the AMBA region is justified by volume/capacity (traffic, number of connected PEs, customers, products and services, etc.), where one of these pairs is in the same physical site of the Inner Cores.

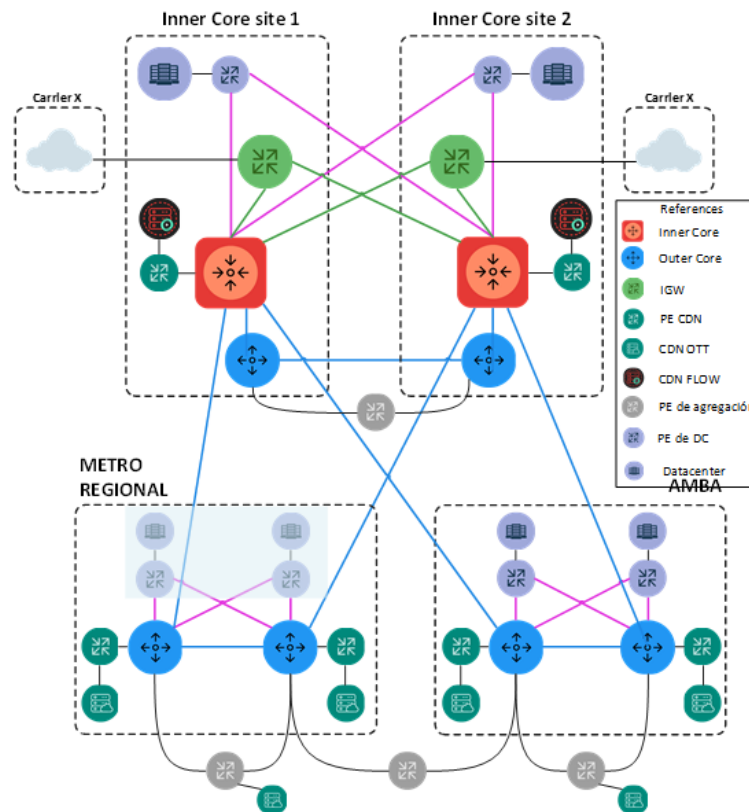
In the current architecture, there are PE nodes that connect to the national (linked to Inner Cores) and regional DC (linked to the Outer Cores), PEs that aggregate connections from final clients, as well as self-owned and/or third-party CDNs, etc. In the future, an Edge DC level hierarchy with multiple sub-levels will be added and the PE nodes and their associated infrastructure (i.e. NFVI) will be sized based on the needs of the areas and customers they serve, volume of traffic, services and products, etc.

---

<sup>21</sup> AMBA (Buenos Aires Metropolitan Area) is the common urban area determined by the Autonomous City of Buenos Aires and other 40 municipalities of the Province of Buenos Aires. According to the 2010 census, this area has 14,800,000 inhabitants, representing 37% of Argentina total, approximately 45% of the country's total GDP and more than 40% of our network traffic.

<sup>22</sup> Internet Gateway.

<sup>23</sup> Provider Edge Router (Label Edge Router).



**Figure 5 – FMO IP Topology.**

Transitioning the Core/Backbone to a full IPv6 network was evaluated. However, given the complexity, cost and technical maturity involved, it was considered best to offer IPv6 services from the edge.

The new Core/Backbone will continue to be based on MPLS<sup>24</sup> and IP4 and it is defined that the Inner and Outer Core nodes should only have P-node<sup>25</sup> functionality within the MPLS domain. This approach could eventually reduce hardware costs because it reduces the functionalities required in these devices.

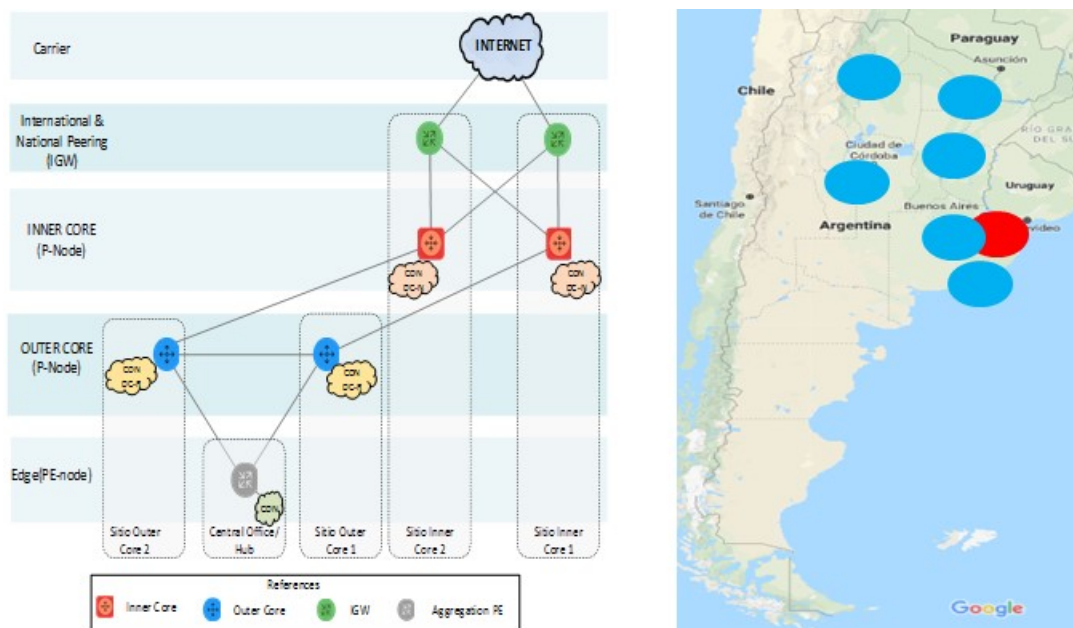
IPv6 services will be provided through the 6PE and 6vPE functionality.

Complementing the previous figure and as a synthesis of what is being developed, Figure 6 shows the new architecture divided by functional layers (figure on the left) and is shown together with the initial geographic distribution (on the right) of Inner Core and Outer Core nodes.

With a view to the future, rules for opening/adding Outer Core sites are defined according to strategic business and technical parameters, such as: number of subscribers, traffic (% increase or threshold-based), number of PE Edge nodes aggregated, latency due to distances, etc.

<sup>24</sup> Multi-Protocol Label Switching.

<sup>25</sup> Provider Router (Label Switching Router).



**Figure 6 - New topology by layers and main nodes geographical locations.**

Given that both former networks were running OSPF protocol, a link state IGP, it was determined to be operationally simpler to unify the IGP's with a single OSPF to begin with. However, in the longer term, a decision has been made to transition from OSPF to IS-IS as the IGP protocol of choice due to a number of reasons – including a future migration to full IPv6.

The definitions for the BGP<sup>26</sup> control plane will follow the usual practices for CSP networks. That is, a full-mesh iBGP<sup>27</sup> of the PE nodes must be achieved through RR<sup>28</sup> dedicated to each particular AFI<sup>29</sup>. The FRT<sup>30</sup> should only be present in the IGW nodes. The PE Edge, DC Gateway and CDN Gateway nodes should include in their routing table only the default route and the networks marked as internal routes for both IPv4 and IPv6 (criteria on which the forwarding plane is based). Thus, starting the MPLS forwarding in PE edge nodes, DC Gateway and CDN Gateway for all services, the Inner and Outer Core nodes must limit their functionalities to switch labels within the MPLS domain.

Although the current RR architecture does not present scalability limits, as the number of RR clients can increase considerably due to the PE nodes number growth in the new network, a flat centralized distribution differentiated by services is defined, connecting the RR in the Inner Core layer. In this way, each RR associated with a service will have its own scalability. Regarding the performance of BGP, this process could be optimized due to the lower number of clients to update.

The IGP (IS-IS in the future, OSPF currently) contains in its topological database only the infrastructure prefixes.

<sup>26</sup> Border Gateway Protocol.

<sup>27</sup> Internal Border Gateway Protocol.

<sup>28</sup> Route Reflectors.

<sup>29</sup> Address Family Identifier.

<sup>30</sup> Full Routing Table – Complete Internet Routing Table.



Currently, the MPLS control plane of both companies is based on LDP<sup>31</sup>. It is proposed to maintain this protocol while the unification process lasts. Then, once the integration of the networks and the new architecture is achieved, it is necessary to evolve towards a control plane based on SPRING<sup>32</sup> (also known as Segment Routing) in order to provide value-added services such as Traffic Engineering.

### 3.3. Optical Network

In the design of the optical network lies the key to the resilience expected for the FMO. It is defined that the optical link between Inner and Outer Core nodes must use an optical DWDM<sup>33</sup> network based on ROADM<sup>34</sup> CD (Colorless-Directionless) nodes, upgrade to FlexiGrid, a control plane that allows the restoration of links and connecting through three, disjoint optical paths.

The new Telecom Argentina has an important fiber optic plant at the country level comprised by the combination of the fiber assets of both former companies. Due to the geographic distribution of the main nodes, in the future, the definition of the optical domains and the vendors co-existence will be reviewed. These definitions and the vendors strategy are subjects of future analyses<sup>35</sup>. Other options such as open platforms will also be part of these future analyses.

It is also necessary to consider that the access/entrance of the three fiber connections to the different sites and / or buildings must be done through at least two disjunct routes. However, it would be desirable for resiliency purposes that three routes be available. Additionally, the fiber to be used must be one dedicated to core services exclusively. Access or clients fiber must not be used to connect to the core in order to avoid unwanted interventions that can generate an impact on it (for example, connecting new clients).

With focus on operational simplicity, the deployment and use of alien lambdas between optical networks of different suppliers is discouraged in this first stage (again, the analysis of open platforms will be done in the near future).

Otherwise, there are some options to overcome the limitations associated with G.653 fiber, such as higher transmission capacity on the channel (100G to 200G, etc.), increase in the amount of lambdas supported by single fiber or replacement by G.652 fiber (incurring in additional costs for fiber deployment). Each of these options has an associated cost and technical implications that should be considered to choose the most appropriate option.

In the case of the rest of the fibers, the capacity upgrade of the lambdas is also considered based on the traffic evolution and considering the TCO to define appropriate capacity.

As an example, Figure 7 shows a normalized cost analysis for the main option of increasing transmission capacity on G.653 fiber. This figure also evaluate different distances that correspond to some of the links currently most compromised in terms of capacity.

---

<sup>31</sup> Label Distribution Protocol.


<sup>32</sup> Source Packet Routing In NetworkG.

<sup>33</sup> Dense Wavelength Division Multiplexing.

<sup>34</sup> Reconfigurable Optical Add-Drop Multiplexer.

<sup>35</sup> The current situation is clearly conditioned by the migration scenarios of the current architectures to the FMO in order to achieve the reduction of the proposed TCO as soon as possible.





**Figure 7 - Options to solve G.653 fiber capacity limitations.**

There will be three disjoint optical paths between Inner and Outer core nodes that were previously mentioned. This requirement is justified by analyzing the current failure rates of the various components, mainly fibers and considering independent failure probabilities. This determines failure cases to tolerate, based on the desired availability.

In the target FMO, it is necessary to be able to support two simultaneous failure scenarios, considering the joint behavior of the IP and optical layers. These failure scenarios are the following (see Figure 8):

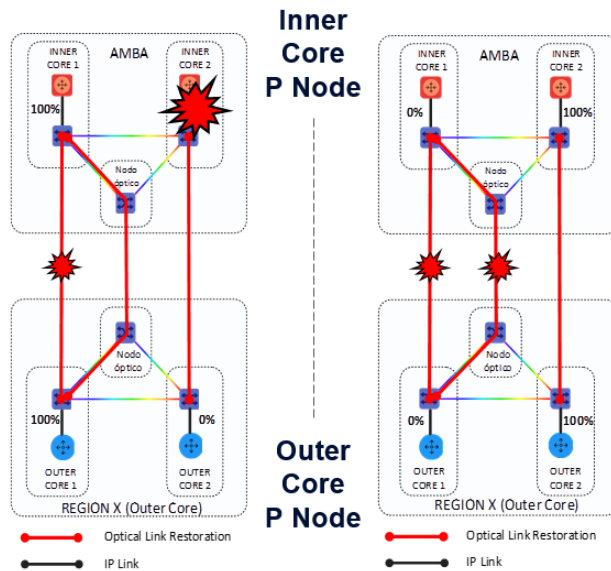
- Outer core site (or Inner Core) failure, simultaneously with a fiber cut
- Double fiber cutting in the Inner and Outer core links

When the first fiber cut of the optical link between Inner core 1 and Outer core 1 occurs, the optical control plane starts the process of restoring the link through the alternative optical node (optical switch). Temporarily, during the restoration (several seconds), there will be a re-convergence in the IGP (with the BFD<sup>36</sup> assistance to accelerate the process) and 100% of the traffic will be transported through the Inner-core 2 and Outer-core 2 links. Once the path is restored, the traffic between the IP nodes is rebalanced back and 50% is established for each link.

As the picture on the left side of Figure 8 shows, if the fiber cut described previously is added with an Inner or Outer core node going down (be it at the DWDM or IP level), the network will continue to be available. This is because there will still be an optical link between Inner-core 1 and Outer-core 1, through the alternative optical node (optical switch), with the IP interface running 100% of the traffic (again relying on BFD to accelerate re-convergence and minimizing the impact on packet loss).

Finally, the failure case of simultaneous fiber cuts is shown in the picture on the right side of Figure 8. In this case, a second concurrent fiber cut is added, where the link between Inner-core 1 and Outer-core 1 going by the alternative optical node (optical switch), also becomes unavailable. In this condition, the IGP re-converges and 100% of the traffic will begin flowing between Inner-core 2 and Outer-core 2 (once again with the BFD help to speed-up the re-convergence). Once the traffic is established in this way, the optical link between Inner-core 1 and Outer-core 1 must not be restored, that is, there only has to be one restoration of each link. Traffic should be maintained between Inner-core 2 and Outer-core 2 in this contingency situation.

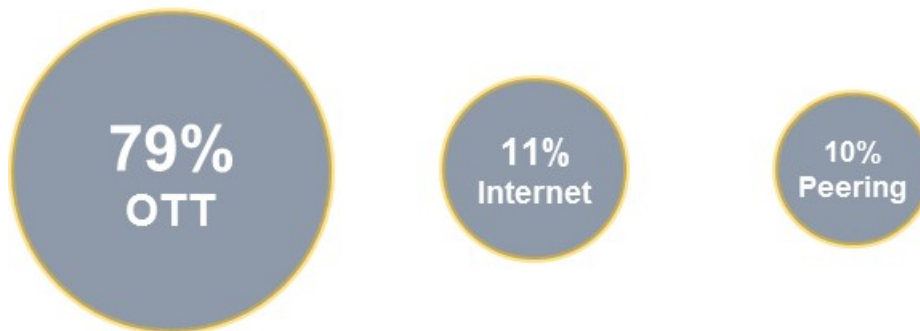
<sup>36</sup> Bidirectional Forwarding Detection.



**Figure 8 – Failures supported considering optical and IP layers behavior.**

### 3.4. CDN Strategy

One of the fundamental premises of the new architecture is the distribution of content, services and functionalities, whenever possible, in such a way that allows minimizing the deployment of infrastructure and transport capacity and to improving QoE<sup>37</sup> (as a consequence of latency and RTT<sup>38</sup> optimization). This distribution scheme serves the needs of future services on the path to 5G. Considering the above, one of the cases to be considered is an optimal distribution of CDN, mainly due to the volume represented by OTT<sup>39</sup> services over the total traffic of our network (see Figure 9).



**Figure 9 - Traffic distribution in the Telecom Argentina Network.**

To accomplish this, two CDN cache hierarchies are proposed, one at the Outer core level and next, even closer to the subscribers, at the Edge site level. These categories of the CDN are based on measured


<sup>37</sup> Quality of Experience.

<sup>38</sup> Round Trip Time.

<sup>39</sup> Over The Top.

traffic patterns previously mentioned and shown in Figure 9 (at peak time, 79% of the traffic currently corresponds to OTT traffic and the remaining 21% is divided into Internet and Peering).

The traffic thresholds for the decentralization of the CDNs towards the Edge sites can be variable according to the CSP and there is also a dependence on the OTT. These PE Edge node level caches are from multiple content providers, therefore, the links between core and Edge nodes must be sized to support the OTT server failure with the highest traffic volume. In order to maintain consistency in the savings related to the sizing of the links, the redundancy of these CDNs should be within the same region and located on the corresponding Outer Cores to avoid over-dimensioning of long haul links.




**Figure 10 - CDN strategy for content distribution.**

Figure 10 shows a CDN distribution scheme, where all the Outer core sites will host CDNs and connect to the Outer-core routers through dedicated PE CDN gateways. To further help with content distribution and network efficiency, CDNs will be placed at select Edge PE site locations as well. Not all Edge PE node sites justify having a CDN – criteria for selecting an Edge PE site to host CDN is based on traffic volumes. Two Edge sites (grey colored PE nodes) corresponding to different regions with their respective CDNs are shown as a reference in the figure. This PE node not only connects customers and access technologies, but also acts as an aggregator of the CDN, without a need for an additional PE CDN gateway as an aggregator, as was the case in the afore-mentioned Outer-core sites.

### 3.5. End-To-End Architecture (2023+)

Considering the integration of both former companies and based on the evolution of the current services and the new services that will emerge on the path to 5G, a unified and distributed architecture of a telco cloud is defined. This architecture is related to the DC hierarchies that are mentioned in the previous sections and in the disaggregation of services and components, including aspects as hardware, software,

CUPS<sup>40</sup>, segmentation of the network, etc. This is shown in Figure 11. In this architecture, only three-level DC hierarchy is shown: two National DCs (co-located to Inner cores), dozens of Regional DCs (colocated to Outer cores) and hundreds of DC Edge (on PE Edge nodes). However, the hierarchy may vary and the edge DC may have different levels according to the volume of customers, traffic, demand for services, infrastructure to support these demands and so on. The common aspect is that in all of them we will have a general purpose generic infrastructure that will allow us to deploy services in a unified manner regardless of their location<sup>41</sup>.



**Figure 11 - End-to-end architecture overview.**

Within these hierarchies, the applications, most of the time, will not be the same, so the scalability requirements of these three types of sites (and the different sub-types within Edge's DCs) must also be different. The NFVI<sup>42</sup> cloud for each level will vary at least the number of instances of each application, the constitution/structure of the components and the dimensions of each one of them. In this way, functionalities now supported on routers and specialized hardware, will be disaggregated, moving some of them to VNF<sup>43</sup> form over computing servers. Temporally, there will be hybrid scenarios requiring coexistence of services mounted on legacy PNF<sup>44</sup> and VNF, until complete transition to the disaggregated model, including PNF just where it makes sense. The fundamental change of this model implies decoupling the innovation cycles of the modules and functional blocks that constitute the infrastructure and services. In turn, this disaggregation should lower entry barriers for new players. Both aspects should make it possible to reduce the TTM of new services and products and lower their TCO and the infrastructure that support them.

An additional aspect to consider is related to the physical limitations that different types of sites could have, both from an operational and TCO point of view. This includes accounting for the criticality of the

<sup>40</sup> Control and User Plane Separation.

<sup>41</sup> This generic definition has a series of non trivial technical aspects to be solved. An example can be the deployment of the control nodes of the VIM and the volume of infrastructure to be deployed in the smaller DC sites.

<sup>42</sup> Network Function Virtualization Infrastructure.


<sup>43</sup> Virtual Network Function.

<sup>44</sup> Physical Network Function.

information to be protected. The characteristics of a National DC will not be the same as those of a DC Edge. Generally, this last type of site is restricted in aspects related to building infrastructure, such as, for example, possibilities to increase physical space and energy capacity (considering refrigeration as a contributor to energy consumption). Again, the volume of customers, traffic, demand for services and products, etc. is what will determine the needs in each case.

Based on the proposal, there will be different NFVI scales for each type of DC hierarchy, but all of them will have a common infrastructure that will respond to a consolidated framework for the administration and management of resources and automation of specific tasks that will, in turn, allow the consistency in the orchestration of end-to-end services and products. As shown in Figure 11, there will be several SDN controllers covering various technological domains (SDN for access, SDN for DC, SDN for transport network and its sub-domains –i.e. optical and IP), but all of them under the aforementioned framework of reference and defining a hierarchical vision that facilitates the orchestration.

Finally, in order to determine the impact of all the proposed optimizations on the final target architecture, an analysis was performed comparing the PMO with the FMO. The results of this analysis is shown in Figure 12.



**Figure 12: FMO vs. PMO comparison.**

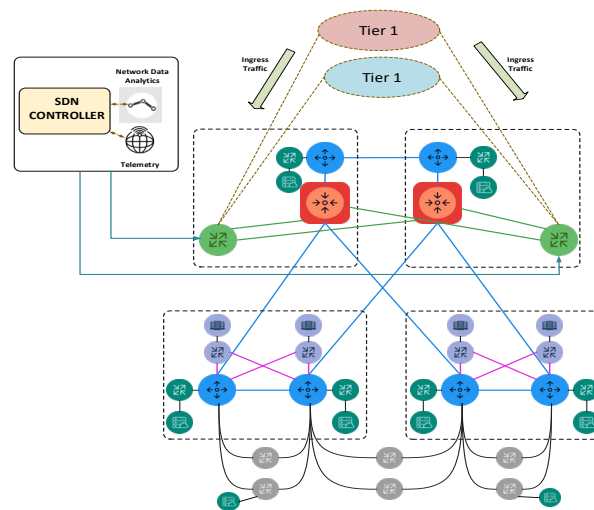
The figure shows a qualitative comparison of the PMO vs. the FMO, with the focus on comparing against the key characteristics/attributes listed in section 2.2: Cost Efficiency, Scalability, Resilience, Security, Operational Simplicity, Future Readiness, Performance and Agility. It can be seen that the future target

architecture (FMO) will have superior characteristics across all measured dimensions compared to the present one (PMO).

### 3.6. SDN Use Cases

In this section, two use cases of SDN application in the transport network are presented. These use cases are presented to illustrate the evolution process until achieving the complete orchestration highlighted by the FMO.

On the one hand, in Figure 13, the new architecture is shown optimizing the use of links to Internet providers (also called Tier I). The challenge in this case is to be able to manipulate the downstream traffic to our network based on a certain condition (congestion of some of these links, a certain level of desired balance, optimization of OPEX and/or the service based on a preferred provider or some other condition). For this purpose, an automated solution based on the analysis of the customer's arriving traffic and then the SDN controller selectively diverting the traffic by manipulating network prefixes on the IGW devices is proposed.




**Figure 13 - Optimization of the links use to Tier I.**

On the other hand, Figure 14, also based on the FMO, shows a potential use case of SDN in combination with traffic engineering. In this case, in the event of a network failure, the delivery of traffic from the CDN results in the circulation of traffic through a sub-optimal path, with higher TCO due to the need to deploy more infrastructure and transport capacity and with higher latency degrading, eventually, the user experience (QoE). Given this scenario, the solution is based on a combination of traffic engineering to configure a explicit route that passes through certain devices<sup>45</sup> (others than those defined by the IGP SPF<sup>46</sup>) complemented by the automatic manipulation of network prefixes. Everything done through the SDN controller.

<sup>45</sup> Basically through Segment Routing Explicit Route Object (SR-ERO) according to the FMO definitions.

<sup>46</sup> Shortest Path First algorithm.



**Figure 14 - Traffic engineering with SDN controller.**

## 4. Summary

The paper can be summarized as follows:

- An architecture is defined, including its topology. The architecture and topology are optimized from the point of view of the traffic matrix and with sufficient flexibility to adapt to eventual changes to it (driven by multiple factors, i.e. current and new services and products evolution, traffic capacity, etc.).
- This architecture is also optimized from the TCO point of view.
- At the IP layer, the FMO will work with IS-IS as a link state IGP with a control plane based on Segment Routing for MPLS, which will enable traffic engineering support.
- In the meantime, the trends and deployments observed in the industry will be analyzed to determine the convenience to migrate to IPv6 and SRv6 without affecting the services and products currently being offered to our clients (based on our evaluation, these technologies are not mature enough for deployment).
- PEs are connected through iBGP via RR receiving only the default route and the networks marked as internal routes for both IPv4 and IPv6 (6PE and 6vPE). The IGP contains in its topological database only the infrastructure prefixes.
- The FRT should only be present in the IGW nodes.
- In the new architecture, both the Inner and Outer core nodes only have the function of switching MPLS labels for packet forwarding. Therefore, they do not run BGP (practically they do not need a IP data plane, only the IGP control plane as support for MPLS).
- On the optical layer, the links between Inner and Outer core nodes will use a DWDM network, with ROADM CD (Colorless-Directionless) nodes, upgrade to FlexiGrid, a control plane that allows the restoration of links and connections through three, disjoint optical paths.
- ONLY core fiber will be used for these connections.
- These three disjoint optical paths between Inner and Outer core nodes, provide high availability. This will allow the support of two simultaneous failure scenarios, considering the joint behavior of the IP and optical layers.



- The distribution of the optical domains that allow the scalability of the platforms and the co-existence of vendors should be analyzed. This definition and vendor strategy will be evaluated in future work.
- The limitations of the G.653 fiber should also be overcome as long as they reach their capacity limits. This fiber type presents particular challenges.
- For the rest of the fibers, capacity limitations will be solved with greater capacity in the channels or lambdas, being the same determined by the traffic evolution and the TCO.
- IP and optical layers should work together in an incremental process of integration that allows the multilayer concept for restoration, capacity and network use optimization, etc.
- As a strategy to improve performance and TCO, content and services should be distributed as close as possible to the end customer, in order to optimize latency, improve customer QoE and minimize transport costs.
- Finally, the framework that allows end-to-end orchestration of the services and products must be defined. It will be leveraged in technologies such as SDN, NFV, SR, Network Slicing and trying to define and deploy the appropriate level of components disaggregation for a CSP Tier 2. This has to be done to decouple this components and functional blocks and improve their innovation cycles, facilitate competition, achieve greater levels of flexibility, agility and automation to allow the ultimate objectives of significantly improving the TTM and reduce the TCO of the network.

## Abbreviations

4Play	Quad Play Products
5G	5th Generation Networks
AFI	Address Family Identifier
ARPU	Average Revenue Per User
ASN	Autonomous System Number
BFD	Bidirectional Forwarding Detection
CDN	Content Delivery Network
CAPEX	CAPital Expenditures
CSP	Communications Service Provider
DC	Datacenter
E2E	End To End
FMO	Future Mode of Operation
IGP	Interior Gateway Protocol
IGW	Internet Gateway
IoT	Internet of Things
IP	Internet Protocol
IP IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
LER	Label Edge Router
LDP	Label Distribution Protocol
LSR	Label Switching Router
MPLS	Multiprotocol Label Switching
NFV	Network Function Virtualization
NFVI	Network Function Virtualization Infrastructure



OPEX	OPerational EXpenditures
OTT	Over the Top
PMO	Present Mode of Operation
P Router	Provider Router (LSR)
PE Router	Provider Edge Router (LER)
QoE	Quality of Experience
QoS	Quality of Service
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RTT	Round Trip Time
SDN	Software Defined Network
SPRING	Source Packet Routing In NetworkinG
SR	Segment Routing
SR-ERO	Segment Routing Explicit Route Object
SRv6	Segment Routing IPv6
TCO	Total Cost of Ownership
TTM	Time To Market
VNF	Virtual Network Function
WAN	Wide Area Network
DWDM	Dense Wavelength Division Multiplexing

## Bibliography & References

- ITU-T G.652. *Characteristics of a single-mode optical fibre and cable. November 2009.*
- ITU-T G.652. *Characteristics of a single-mode optical fibre and cable. November 2016.*
- ITU-T G.653. *Characteristics of a dispersion-shifted, single-mode optical fibre and cable. July 2010.*
- ITU-T G.654. *Characteristics of a cut-off shifted single-mode optical fibre and cable. November 2016.*
- ITU-T G.655. *Characteristics of a non-zero dispersion-shifted single-mode optical fibre and cable. November 2009.*
- ITU-T Y.3300. *Framework of SDN.*
- IETF RFC 3031. *Multiprotocol Label Switching Architecture.*
- IETF RFC 1142. *OSI IS-IS Intra-domain Routing Protocol.*
- IETF RFC 7142. *Reclassification of RFC 1142 to Historic.*
- IRTF 7246. *SDN: Layers and Architecture Terminology. Internet Research Task Force.*
- ONF TR-502. *SDN Architecture 1.0 – Open Networking Foundation.*

IETF. *Source Packet Routing in Networking*.

ETSI GS NFV-MAN 001 V1.1.1. *Network Functions Virtualisation (NFV) ; Management and Orchestration*.

SCTE - Journal of Network Operations - Volume 1, Number 2. *Framework for The Evaluation of SDN WAN Controllers*.

# Changing the World: IoT Chaos as a Ladder to Improving Security

## The ISO/OCF Standard and Implementation for Security

A Technical Paper prepared for SCTE•ISBE by

**Brian A. Scriber**

Distinguished Technologist and VP of Security Technologies  
CableLabs  
858 Coal Creek Circle, Louisville, CO 80027  
@brianscriber  
b.scriber@cablelabs.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Changing the World .....	3
1. IoT Security Misconceptions.....	3
2. How Can We Fix This?.....	4
3. Entrepreneurial Incentives.....	4
The Open Connectivity Foundation Security Specification.....	5
4. The OCF Network .....	5
4.1. Network Provisioning.....	5
5. Onboarding Detail .....	5
5.1. Discovery .....	5
5.2. Ownership Transfer .....	6
5.3. Provisioning.....	6
5.4. Normal Operation .....	6
6. State Transitions for OCF Devices .....	6
6.1. RESET .....	7
6.2. Ready For Ownership Transfer Method (RFOTM).....	7
6.3. Ready For Provisioning (RFPRO).....	7
6.4. Ready for Normal Operation (RFNOP).....	8
6.5. Soft Reset (SRESET) .....	8
7. Access Control.....	8
7.1. ACEs and ACLs .....	8
7.2. Access Control Types.....	9
7.3. Access Management Service.....	9
8. Message Integrity and Confidentiality .....	9
8.1. Credential Management Service.....	10
Conclusion .....	10
Abbreviations.....	10
Bibliography & References .....	11

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Device Provisioning States .....	7
Figure 2 - Access Control Example.....	9

# Introduction

With each device on our networks vulnerable to attack with processing/memory/storage and network credentials, how can we help to protect consumers, the network and other participants? How do we raise the security tide for all boats? How do we attack the economic problem here while we work the technical issues? If an entrepreneurial CTO is ever asked by his CFO or CEO “How much is security going to cost me?” we’ve already lost. The Open Connectivity Foundation (OCF) is an organization of over 450 companies, with several SCTE member companies actively participating and a couple serving on the Board of Directors. OCF is working together to not only write the specifications for interoperability and security, but also build out an entire open-source (as in free) implementation of this specification (IoTivity) to help even the most cash-strapped start-up build secure IoT software. Join us to learn the current state of OCF, IoTivity, certification and how you can get started changing the world of IoT security.

## Changing the World

IoT Security has an economics problem. There is a misalignment between incentive structures with respect to the externalities of botnet attacks. It takes an investment of time and money for manufacturers to secure devices and for manufacturers to update a device they have already sold, it pulls directly from their bottom line; this is a tangible disincentive for manufacturers to address problems caused by insecure devices. The victims of botnet activity, Distributed Denial of Service (DDoS) attacks, and other malicious traffic are not those manufacturers; those costs are born by downstream inhabitants of the IoT ecosystem.

Those effects impact consumers, governments, network operators, and any target of a DDoS attack. The economic burden of such an attack on a retail site averages between \$20,000 and \$40,000 per hour<sup>i</sup>, however the cost to hire a 100Mbps attack for an hour is roughly \$5<sup>ii</sup>. The ability to rent botnets stems from the raw number of devices available for malicious actors to compromise as well as the ease of compromise of connected devices with little or no security (botnets Torii, Demonbot, Mirai/clones, and Chalubo all heavily leverage IoT). There are roughly 8 connected devices per person on the planet in 2019 with a 50% growth expected over the next three years. That disparity between the ease of attack and damage from an attack, times the 793,377 attacks through cable networks in 2018<sup>iii,iv</sup> is the economic burden of insecure devices.

### 1. IoT Security Misconceptions

The major parts of this problem come from the realities of the niche smart home electronics manufacturers occupy and a few erroneous assumptions about IoT include the following:

- **Misconception #1: “Device pricing is the only mechanism available for differentiation against competition.”** However, as Apple has shown in 2018 and 2019<sup>v</sup>, device security can be a differentiator and buyers will make decisions based on how well devices protect their investments.
- **Misconception #2: “Security isn’t important for end users.”** However, security is cited within the top three buying concerns<sup>vi</sup> and has an influence on 75% of IoT consumers<sup>vii</sup>.
- **Misconception #3: “Security is about the protection of the device”** Concerns raised in discussions around IoT Security in relation to the NIST direction on IoT<sup>viii,ix</sup> continue to return to classes of device with a desire to create different security levels for different device types. The argument is that a lightbulb or a connected Barbie® doll don’t require the same security as other devices. The reality is that those devices have a processor, memory, power, a network stack, and

networking credentials that the owner provided which make each of them an ideal launching point for further attacks<sup>x</sup> and make them excellent participants in the botnets that have garnered legislative attention. Regulators have no motivation to care about the device being compromised if that's where it ended, they care about this sector of the economy because botnets using those devices are the weapons that can take down strategic infrastructure. It isn't about the device, it's about the network; it's always been about the network.

## 2. How Can We Fix This?

The quick answer, the one that falls into Mencken's "neat, plausible, and wrong"<sup>xi</sup> category is that we can just create a specification for IoT devices, make it a standard and everyone will use it. The trick here is that IoT isn't as simple as the three letter acronym makes it appear, these devices are actually computers, some with more or less resources than others, but they have operating systems, drivers, radios, networking stacks, layers of code libraries, processors, memory, power needs, some have cryptographic coprocessors and Trusted Platform Modules (TPMs), others are general purpose computers with small form factors. Writing a specification that doesn't take this into consideration is a recipe for failure.

Even with an agreed-upon specification, adoption of that specification by the global manufacturing community is another cognitive leap that takes a measure of suspension of disbelief. The real trick to this, however, is the entrepreneurial company that has decided to build a connected device. They may not have a cryptographer on staff, they may not have a network security expert or even a networking expert; this company may outsource the entirety of their development to other lowest-bidder firms where simple hardware and old software are combined to provide the Minimum Viable Product (MVP) for launch of the idea. This MVP likely has no consideration for security, the company releasing it may not know all the code libraries used, the versions, the vulnerabilities for each version, or even have applied it to the correct hardware. Permissions are likely granted through hard-coded credentials and as the mirai botnet leverages, those passwords are easy to try in series because no effort to limit unsuccessful attempts is ever engaged. This entrepreneurial company has no interest in reading a 200 page security specification let alone comply with that for what they consider a novel product with perceived market pressures and investors who follow Eric Reis's Lean Startup methodology driving them to release as quickly and as early as possible. The prevalence of botnets shows that addressing security isn't always treated as a priority in this environment.

## 3. Entrepreneurial Incentives

Even if security is brought up at this hypothetical entrepreneurial company, if it were to negatively impact a schedule or if it led to additional costs, a conversation between the CFO or CEO would ensue about those impacts. If this conversation happens, without an answer of "security is free" or "it would cost us extra time or money to \*remove\* security from the software", the result is a less secure product. When it comes to product security, it is for the perceived cost of security, that specifications alone cannot satisfy the market; specifications must be accompanied by implementation software that is free-to-use, modifiable, reviewed, and preferably open-source.

The OCF is such a combination; it includes a detailed security specification and has IoTivity, a separate, Linux Foundation managed, open source implementation of that specification available for use without cost. The added incentive from the OCF is that in using this specification/software, the product created will be interoperable with the IoT product lines from other members of the OCF which include over 450 of the world's most influential and prolific manufacturers including Intel, Shaw, Samsung, LG, Microsoft, Electrolux, Cox, Haier, Comcast, Qualcomm, Charter, Cisco, and others. Now, an entrepreneurial firm has the incentive to use the IoTivity code because it solves many of their non-application layer concerns. They are incented because they know that IoTivity is built with security which includes device identity, confidentiality of messages, access control, and strong authentication. The are

also incented because they now know that they have the door open for interoperability with product lines from the hundreds of OCF member companies. This is the argument that wins over not just the large players in the IoT space, but provides cover for all of the entrepreneurial companies as well. This is how security becomes part of the foundation for an ecosystem.

# The Open Connectivity Foundation Security Specification

## 4. The OCF Network

OCF is a Layer-4 and Layer-5 technology, it is transport agnostic and can run on WiFi, Bluetooth, Thread and efforts are underway for other transports. As new OCF Devices are introduced to this network, the Device Ownership Transfer Service (DOTS – aka “onboarding tool”) broadcasts a discovery message to which the OCF Device responds. The DOTS then interrogates device credentials and, through the Credential Management Service (CMS) issues new local OCF-network credentials to the device. Through the Access Management Service (AMS) the DOTS then establishes and enforces access control to the Resources within the Device. Resources are collection of related Attributes that follow the OCF Data Model. Access to a Resource is limited through Create, Read, Update, Delete, and Notify (CRUDN) operations stored in Access Control Entries (ACE) within the Access Control List (ACL) aspect of the AMS. This allows Devices to interact with other proximal devices and engage in a Scenes or Rules to enable advanced functions like “vacation mode” where lights turn on at sporadic intervals and curtains raise and lower accordingly.

### 4.1. Network Provisioning

Because OCF is a Layer-4/Layer-5 technology, there is a presumption that the Device has been onboarded into the network where it can receive multicast messages from the DOTS. With WiFi, the WiFi Alliance (WFA) has a Device Provisioning Protocol (DPP) for securely introducing the Device to the WiFi network and provisioning it with credentials to receive additional instructions from OCF. The OCF, the WFA, and CableLabs, are engaged in making this a seamless transition and effortless for the user.

## 5. Onboarding Detail

Onboarding has three distinct phases: Discovery, Ownership Transfer, Provisioning, and Normal Operation.

### 5.1. Discovery

In the Discovery phase of onboarding, the DOTS sends a multicast discovery message which is not encrypted. Devices receiving this message, which have not been “owned” reply directly to the DOTS. The reply message includes the list of methods by which the Device can communicate through the ownership transfer – these methods are called Ownership Transfer Methods (OTM). The currently supported OTMs include the “Just Works” Authenticated Diffie Hellman exchange, the Pre-Shared Key symmetric cryptographic exchange, and Certificate Based Onboarding. The Certificate-Based Onboarding OTM has a “Baseline” which is equivalent to a self-signed certificate and then additional optional “profiles” for enhanced Device identity protection.

- **Purple:** The “purple” profile makes (unverified, but attested to) claims about the capabilities of the device (e.g. that there is a Secure Execution Environment (SEE)).

- **Blue:** The “blue” profile relies upon the DOTS to optionally check on the current certification status for the Device based on the make and model of the Device through an online certification verification system.
- **Black:** The “black” profile requires that Devices have passed the OCF testing protocol at one of the Approved Test Labs (ATL) before it is issued a digital certificate which is part of the official OCF Public Key Infrastructure (PKI). The certificate for the Device chains to one of the approved Certificate Authorities who can only issue certificates in accordance with the Certificate Policy defined by OCF for compliant devices.

The supported OTMs might include multiple methods to acceptably onboard the Device. The decision about which OTM to select is the prerogative of the owner of the network through the DOTS.

## 5.2. Ownership Transfer

There are three steps in the Ownership Transfer process:

1. The DOTS, having received the OTMs supported, selects the OTM and configures that OTM to the Device – this interaction is also over an unsecured channel.
2. DOTS and the unowned Device perform the OTM using the credentials associated with the OTM. Part of this exchange is the establishment of the TLS handshake.
3. DOTS configures the Security Virtual Resources (SVRs) of the Device, the CMS and the AMS to authorize itself for further provisioning of the Device. These secure interactions occur over TLS.

At this point, the Device is “owned” by the DOTS and the Ownership Transfer is complete.

## 5.3. Provisioning

The two steps involved in provisioning for the Device are the creation of credentials and the establishment of access control properties.

The CMS issues credentials to the Device, currently both symmetric and asymmetric credentials are supported, but the preference is toward asymmetric to support Role Based Access Control (RBAC) which is discussed later. Interactions with the CMS are under secure communications using TLS.

The AMS provisions the access control policies for the Device, creating ACEs for each relationship with other Devices that are to be granted access for reading, notification, or setting of attributes using the CRUDN ACE for the ACL. All interactions with the AMS require secure communications using TLS.

At this point, the Device is provisioned with credentials for the OCF Network and has been granted the access it needs to be effective, it is now ready for Normal Operation.

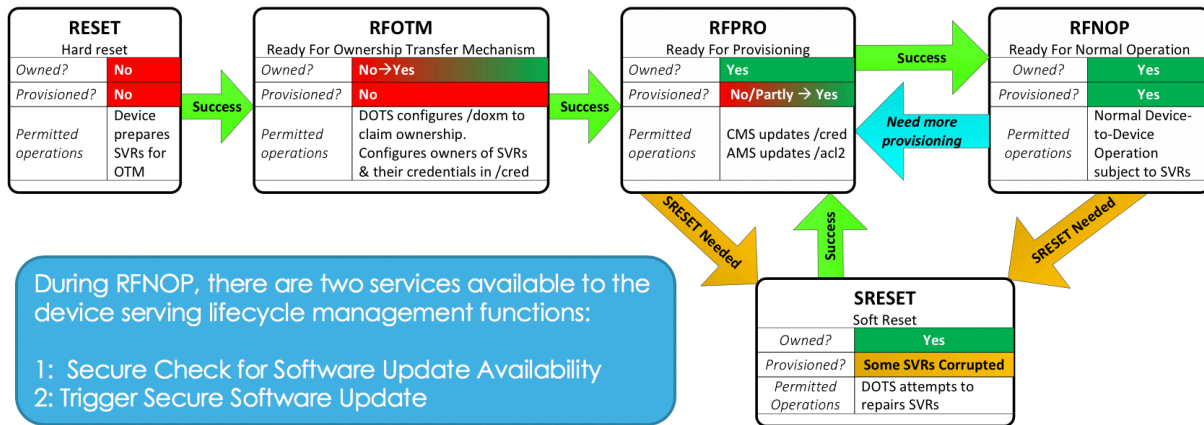
## 5.4. Normal Operation

During Normal Operation, the Device acts within the access control granted through the AMS using the credentials it was issued by the CMS. Should either of these require changes, the Device returns to Provisioning.

# 6. State Transitions for OCF Devices

In the Onboarding Detail section, some implied state transitions for the Device occur; this section explores those at a greater level of detail. Figure 1 - Device Provisioning States explores these state transitions.





Device can transition to **RESET** from any state (these transitions are not shown)

**Figure 1 - Device Provisioning States**

### 6.1. RESET

The Reset state is the original state the device came from out of the box, or as close to that state as possible if it arrives back at that point through a factory-reset or “hard reset”. Initial identity credentials such as the digital certificate issued to the device by the manufacturer are unchanged in a transition back to this state, but all local credentials are removed from the Device.

### 6.2. Ready For Ownership Transfer Method (RFOTM)

RFOTM can be transitioned to only from the RESET state. The action taken during this transition is the preparation for the security resources (SVRs) to be established which is likely already the case in RESET.

Transitioning out of RFOTM is the Ownership Transfer described in section 5.2. If it is successful, the Device transitions to RFPRO, but if the Ownership Transfer is unsuccessful, the Device sets its “Owned” status to false, deletes any credentials other than the initial identity credentials, resets any other Resources to their defaults and transitions back to RESET.

### 6.3. Ready For Provisioning (RFPRO)

RFPRO can be transitioned to from three different states: the RFOTM during initial onboarding, from SRESET during a partial update, and from RFNOP if additional credentialing or access management need to be provisioned for the Device. In all of these cases the “Owned” status remains true.

Transitioning out of RFPRO also can go to three different states. In the nominal case, after the provisioning of credentials with the CMS and provisioning of access management with the AMS occur (as described in section 5.3), the Device transitions successfully to the RFNOP state. In the case where there are errors with the SVRs, and if those errors are potentially recoverable, the Device transitions to the SRESET state to hopefully remedy any issues. If the errors persist or are irrecoverable, the Device transitions to the RESET state with the accompanying loss of local credentials and a reset to default Resource values.

## 6.4. Ready for Normal Operation (RFNOP)

RFNOP is the nominal state for a Device after it has been through onboarding. This is the state for normal operation and, in a perfect world, where the Device will spend the majority of its time. Transitioning to RFNOP can **only** be accomplished through successful provisioning and a transition directly from RFPRO.

Transitions from RFNOP are to SRESET in the case of corrupted SVRs that may be recoverable, and also transition back to RESET if the corruption requires a hard reset. It is also possible to transition from RFNOP to RFPRO when either new credentials or new access control is required as those can only be granted to Devices in the RFPRO state.

## 6.5. Soft Reset (SRESET)

The SRESET state exists to prevent complete reworking of all of the credentialing and access management assigned to the Device. It carries some inherent risk because there is no way to determine exactly how the corruption may have occurred. The SRESET state is managed by the DOTS, which is controlled by the owner of the OCF Network. If the owner prohibits the SRESET state, any transition to this state transitions directly to RESET instead, including all of the requisite resetting of credentials and access control. It is possible to transition to SRESET from either RFPRO or RFNOP.

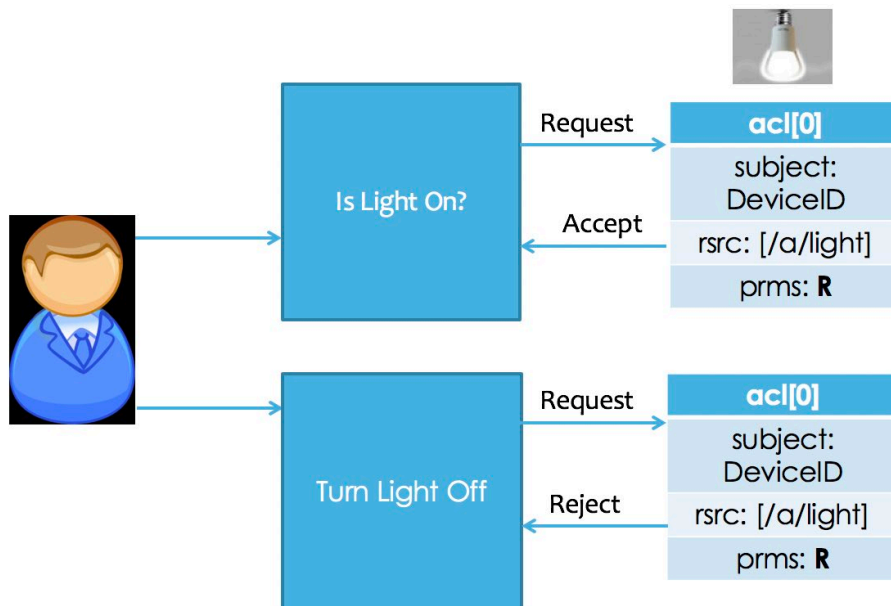
Transitioning out of SRESET to RESET can happen either through the policy as described above, or through a failed attempt to restore SVRs through the DOTS. If the SVRs can be restored successfully through the DOTS the Device can transition to only the RFPRO for verification and any refinements to either the credentials installed in the CMS or to the access control granted through the AMS.

# 7. Access Control

Access Control limits the Smart Toaster from opening the Smart Lock on the front door. It is critical in an ecosystem and the entire network can be compromised if the access control isn't part of the design from the very start.

## 7.1. ACEs and ACLs

Devices have Resources, and each of those Resources are protected by an ACE. The ACE allows another Device the ability to act on the Resource with different granted permissions: Create, Read, Update, Delete, and Notify. Most of these are self-explanatory, but Notify can be thought of as a perpetual Read where the reader is advised of changes made to the Resource. **All** requests to any of a Device's Resources are subject to the ACL policy verification where the appropriate ACE is evaluated. A simple overview of a successful Read on a Device and an unsuccessful Update is shown in Figure 2 - Access Control Example.



**Figure 2 - Access Control Example**

## 7.2. Access Control Types

There are two types of access control, Subject-Based Access Control (SBAC) and Role-Based Access Control (RBAC).

In the SBAC approach a single Device is given the ability to perform actions specified by the CRUDN options and recorded in the ACE. The Device is named by the network identity of that Device assigned at onboarding time.

In the RBAC approach the ACE specifies the Role permitted to perform the CRUDN operations. The Role can then be assigned to any of the appropriate Devices. An example of this might be Temperature Sensors that allow anything with the Thermostat Role to read different temperature Resources.

## 7.3. Access Management Service

The Access Manager Service (AMS) is the only mechanism for updating ACE/ACLs. To ease deployment options and to make some communications methods easier, wildcarding of permissions was added to the model.

It is important to note that permissions are checked inbound to Resources – OCF does not currently have outbound ACLs restricting a Device from attempting a specific communication, it relies on the ACE of the Resource being requested to check permissions.

## 8. Message Integrity and Confidentiality

Within the onboarded OCF network, all unicast messages are secured using TLS or DTLS. Multicast messages are not secured, but also do not have a requirement that Devices respond to, or consume, the multicast messaging. All unicast messages are signed, ordered, and encrypted. This protects against eavesdropping on message contents, tampering with messages, and replay attacks.

## 8.1. Credential Management Service

In order for Devices to have this level of communication protection, the Devices must have usable credentials. Those credentials are assigned in the RFPRO state before the Device is in the RFNOP state associated with normal operation. When the credential expire, or if new credentials need to be provisioned, the Credential Management Service (CMS) installs or renews as necessary and in compliance with the policy defined by the network owner.

## Conclusion

The Kaizen<sup>xiii</sup> approach of asking the five whys helps to drive the understanding of IoT Security and how to change it. **Why** are legislators looking at regulating IoT? Because our networks are being attacked. **Why** is that? Networks are being attacked because of the ease and availability of attack platforms. **Why** are these platforms so inexpensive? Attack platforms are comprised of the prolific deployment of low-security IoT devices. **Why** are these prone to attack? Because there is an economic disincentive to manufacturers to design these devices with security in mind and to update to defend against new threats. **Why** can't we unwind this disincentive? The economics of adding security at the point of design adds cost, the consumer cannot adequately differentiate between the security provided by different products, and different devices are challenged by not "speaking the same language" let alone share a common security paradigm.

The answer to those three primary drivers are the primary forces that brought together 450+ major manufacturers to create the Open Connectivity Foundation. The three problems each have a pillar in the structure of OCF:

- 1) **Interoperability:** OCF provides a common data model for Devices and a set of rules for how device security needs to interact. Devices must have a unique, attestable, immutable identifier used to onboard into the network; communication is secured via TLS to guarantee message integrity and confidentiality; access is controlled at the Resource level and Device integrity is managed through clearly defined state transitions.
- 2) **Cost of Security:** Through building not just a specification, but an actual implementation of that specification, and then providing that implementation without cost through the Linux Foundation's open source project, IoTivity, the OCF has created an architecture with security as a primary design consideration. This creates a world where it actually costs more if an entrepreneurial company wanted to build a product without security considerations.
- 3) **Consumer Education:** The OCF and their Approved Test Labs which certify compliance of the Devices against a rigorous framework to check alignment with the specification. This logo enables the consumer to know that, if the OCF label appears on the box, the Device inside has passed all of the security compliance and interoperability tests approved by the OCF.

This is the beginning of the real value and climb out from the IoT insecurity morass. The work isn't over, there will still be improvements made and there are still millions of insecure devices currently deployed, but these are steps in the right direction and the hope is that we will see the tide rise for all boats as the tools have now been made available.

## Abbreviations

ACE	access control entry
ACL	access control list
AMS	access management service
ATL	Approved Test Laboratory

CMS	Credential Management Service
CRUDN	Create/Read/Update/Delete/Notify access control permissions
DOTS	Device Ownership Transfer Service
DPP	Device Provisioning Protocol (WiFi Alliance)
OBT	Onboarding Tool
OCF	Open Connectivity Foundation
OTM	Ownership Transfer Methods
PKI	Public Key Infrastructure
RBAC	Role Based Access Control
SEE	Secure Execution Environment
SVR	Security Virtual Resources
TPM	Trusted Platform Module
WFA	WiFi Alliance

## Bibliography & References

ISO/IEC 301 18-1:2018 Information Technology – Open Connectivity Foundation (OCF) Specification – Part 1: Core specification <https://www.iso.org/standard/53238.html>

<sup>i</sup> <https://www.coxblue.com/12-ddos-statistics-that-should-concern-business-leaders/>

<sup>ii</sup> CableLabs: Survey of Dark Web Activity, most recently verified, March 2019

<sup>iii</sup> <https://securelist.com/ddos-report-q1-2019/90792/>

<sup>iv</sup> [NETSCOUT Threat Intelligence Report 1H 2018](#)

<sup>v</sup> <https://www.techtimes.com/articles/239693/20190314/apples-new-iphone-ad-shows-how-much-privacy-matters.htm>

<sup>vi</sup> Open Connectivity Foundation primary research, 2017

<sup>vii</sup> <https://www.cs.cmu.edu/~pemamina/publication/CHI'19/CHI19.pdf>

<sup>viii</sup> <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8228.pdf>

<sup>ix</sup> <https://csrc.nist.gov/CSRC/media//Publications/sp/800-53/rev-5/draft/documents/sp800-53r5-draft.pdf>

<sup>x</sup> <https://www.cablelabs.com/just-lightbulb-need-security>

<sup>xi</sup> “Explanations exist; they have existed for all time; there is always a well-known solution to every human problem — neat, plausible, and wrong.” – H.L. Mencken, “The Divine Afflatus” in *New York Evening Mail* (16 November 1917); later published in *Prejudices: Second Series* (1920) and *A Mencken Chrestomathy* (1949)

<sup>xii</sup> <https://en.wikipedia.org/wiki/Kaizen>

# **The Pivotal Role of Cable Gateways in the Internet of Things**

A Technical Paper prepared for SCTE•ISBE by

**Ryan Michael Cunningham**  
Hardware Design Engineer  
Comcast Corporation  
1800 Arch Street  
Philadelphia, PA 19013  
267-254-4232  
Ryan\_Cunningham@comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Overview of Cable Gateway.....	4
1. IEEE 802.11.....	4
2. IoT .....	5
2.1. IEEE 802.15.4 .....	5
2.2. Bluetooth Low Energy.....	6
Physical Layer Coexistence.....	6
3. PSK & QAM .....	7
4. FSK .....	8
5. OFDM & OFDMA.....	9
6. IEEE 802.11.....	11
7. IEEE 802.15.4.....	12
8. Bluetooth Low Energy .....	13
9. Theoretical Calculations.....	14
Network Coexistence Techniques.....	15
10. Network Management.....	16
10.1. Scheduling .....	16
10.2. Channel Selection .....	16
11. Prioritization.....	17
12. Spatial Mapping .....	17
Conclusion .....	18
Abbreviations.....	18
Bibliography & References .....	19

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Typical North American Wi-Fi channels in the 2.4 GHz ISM Band .....	5
Figure 2 - Typical Wi-Fi channels and IEEE 802.15.4 channels.....	6
Figure 3 - Bluetooth Low Energy operating channels .....	6
Figure 4 - Inphase and Quadrature QPSK data.....	7
Figure 5 - Power spectral density plot of PSK and QAM signal.....	8
Figure 6 - FSK IQ data set example.....	9
Figure 7 - OFDM Modulated Carrier as the addition of multiple subcarriers .....	10
Figure 8 - OFDMA signal separated into individual subcarriers over frequency and time .....	11
Figure 9 - Transmit spectral density mask of an IEEE 802.11 DSSS 20 MHz signal .....	11
Figure 10 - Transmit spectral density mask of an IEEE 802.11 OFDM 20 MHz Signal .....	12
Figure 11 - IEEE 802.15.4 Transmit spectral density mask.....	13
Figure 12 - Bluetooth Transmit Spectral density mask.....	13

Figure 13 - IEEE 802.11, IEEE 802.15.4, IEEE 802.15.1, Bluetooth, Bit error Rate vs Signal to Noise Ratio.....	15
Figure 14 - IEEE 802.11 ax 20 MHz channel resource units.....	16
Figure 15 - Angle of Arrival calculation using multiple antennae .....	18

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Typical North American IEEE 802.15.4 channels in the 2.4 GHz ISM Band .....	5
Table 2 - Criteria for calculating IEEE 802.11 spectral density.....	12
Table 3 - Criteria for calculating IEEE 802.15.4 spectral density.....	13
Table 4 - Criteria for calculating IEEE Bluetooth Low Energy spectral density .....	13



# Introduction


The term the Internet of Things, IoT, was coined in 1999, the same year IEEE 802.11 b standard was released, and the beginning of the age of the internet. At the time the IoT was strictly conceptual, but as the connected home continues to explode through improvements in cellular, Wi-Fi, cable standards, IoT continues to become more relevant. Despite the growing interest in IoT, the lack of conformity throughout the industry is resulting in a kludgy user experience, in most cases, requiring each device to be connected to individual hubs, which are subsequently connected to the home gateway. Previously IoT implementations created awkward user experiences discouraging the average consumer from fully immersing themselves in the technology, and restricted the customer base to technology-savvy individuals. Driving conformity in IoT by consolidating the hubs into the gateway, gives the end user to the ability to easily commission and control their video, Wi-Fi, and IoT devices through one interface. More importantly, the gateway provides insight into the entire connected home ecosystem through a constant flow of data provided by connected clients, allowing different subsystems in the home to coexist, and improve performance. In summary, consolidating IoT infrastructure and services into the gateway can significantly improve customer experience by driving coexistence performance between individual subsystems through the gateway.

## Overview of Cable Gateway

To increase performance, and improve user experience the cable gateway continues to add functionality. A typical gateway, consists of Wi-Fi for high throughput wireless applications such as streaming video, Docsis serves as the data pipeline in and out of the home, and IoT provides home automation and home security capabilities. Despite having different functions, all of these different standards and protocols together create a whole home experience, providing the end user with an entire ecosystem that cannot be replicated with individual sub systems. The key differentiation from an individual subsystem implementation is the coexistence benefits, particularly for the the wireless systems. The most common IoT protocols share the same ISM band as Wi-Fi creating many challenges for IoT devices the can result in truck rolls for services providers.

### 1. IEEE 802.11

Wi-Fi is the primary protocol utilizing the IEEE 802.11 standard in gateways. The IEEE 802.11 standard has had multiple iterations supporting a large range of data rates from 1 Mbps to 1201 Mbps, as well as DSSS, OFDM, and OFDMA modulations techniques. Each iteration in the standard has provided additional techniques to improve throughput, and optimize channel utilization through different modulation, coding schemes, and bandwidths. The current standard supports 2.4 GHz and 5 GHz, with additional frequency bands in the horizon. The end goal is to maximize data rates, and is ideal for streaming services, and high throughput applications. Operating channels occupy either 20 MHz, 40 MHz, 80 MHz, or 160 MHz bandwidth in the 5 GHz frequency band and 20 MHz or 40 MHz in the 2.4 GHz frequency band. The typical operating channels in the 2.4 GHz frequency band are illustrated in Figure 1



**Figure 1 - Typical North American Wi-Fi channels in the 2.4 GHz ISM Band**

## 2. IoT

### 2.1. IEEE 802.15.4

The two primary IoT protocols utilizing IEEE 802.15.4 are Zigbee and Thread. The primary applications of Zigbee and Thread devices are low power sensors for home automation and home security. IEEE 802.15.4 supports O-QPSK, GFSK modulated signals with data rates of 250kbps, 40 kbps, and 20kbps at 2.4GHz, 915 MHz, and 868 MHz subsequently. The protocols are designed to be as robust and as low power as possible. Devices are typically sleepy end devices, waking up on user interactions, and periodically providing device diagnostics to report device status. The 2.4 GHz frequency band supports 16 channels with 2 MHz bandwidth and 5 MHz spacing. The primary operating channels, listed in Table 1, have the least spectral overlap with 2.4 GHz frequency band IEEE 802.11 channels, and are illustrated in Figure 2. Zigbee and Thread have ability to dynamically change channels, however due to protocol restrictions it is not trivial. Devices can potentially get stranded, as well as impact battery life.

**Table 1 - Typical North American IEEE 802.15.4 channels in the 2.4 GHz ISM Band**

Channel	Center Frequency(MHz)
15	2425
19	2445
20	2450
25	2475

## 2.4 GHz ZigBee Channels




Figure 2 - Typical Wi-Fi channels and IEEE 802.15.4 channels

## 2.2. Bluetooth Low Energy

The primary application of Bluetooth Low Energy is low powered streaming devices, and low power sensors including; door locks, medical devices, wearables, and audio streaming devices. Bluetooth Low Energy channels are illustrated in Figure 3. Bluetooth Low Energy supports GFSK modulation with 125kbps, 500kbps, 1 Mbps, and 2 Mbps data rates which are achieved by varying the spreading factor. Bluetooth Low Energy channels occupy the 2.4 GHz frequency band consisting of 40 2 MHz channels. Bluetooth Low Energy takes advantage of the tightly spaced channels through frequency agility changing channels at predetermined intervals, and channel masks can be implemented to avoid channels with high interference. Three of those channels are designated as advertising channels, and cannot be avoided with a channel mask and are illustrated as the green channels in Figure 3.




Figure 3 - Bluetooth Low Energy operating channels

## Physical Layer Coexistence

I/Q based communications is a form of digital modulation where in-phase, and quadrature data are mixed together to form a modulated carrier. The in-phase data is represented by Equation 1 and quadrature data is represented by Equation 2. The in-phase and quadrature data can be manipulated by adjusting the amplitude represented by  $A_c$ , the frequency represented by  $f_c$ , or phase represented by  $\phi$ . Different implementations of I/Q modulated carriers produce different power spectral densities. These power spectral densities can be used to estimate the Signal-to-Interference Ratio, SIR, which can be used to


calculate chip error rate, bit error rate, and packet error rate. The CER, BER, and PER determines whether or not a signal of interest can be received. Individual wireless subsystems, can determine SIR with respect to adjacent networks. Essentially, a Zigbee network can easily determine how well it can coexist with other Zigbee networks, but not necessarily a Wi-Fi or Bluetooth network, and vice-versa. Consolidating IoT device traffic to the cable gateway, provides the opportunity to calculate if devices in other networks can concurrently communicate on the physical layer.

$$\text{Equation 1 } x_i(t) = A_c \cos(2\pi f_c t + \phi)$$

$$\text{Equation 2 } x_q(t) = A_c \sin(2\pi f_c t + \phi)$$

### 3. PSK & QAM


Phase shift keying is a form of digital modulation where IQ, In-phase and Quadrature, data is mixed together at different phase offsets to represent bit(s). BPSK represents 1 bit per symbol where in-phase data remains constant, but quadrature data phase,  $\phi$ , is shifted. QPSK, illustrated in Figure 4 - Inphase and Quadrature QPSK data, represents 2 bits per symbol where both in-phase data, and quadrature data phase,  $\phi$ , is shifted. Quadrature amplitude modulation consists of both phase,  $\phi$ , and amplitude,  $A_c$ , is shifted for both in phase and quadrature data. QAM can support 2, 4, 6, 8, and 10 bits per symbol by increasing the amplitude variations. The number of bits represented by the modulated carrier as an interferer is directly correlated to the spectral desensity and, as a result QAM, BPSK, and QPSK will have different impact on devices as an interferer.



**Figure 4 - Inphase and Quadrature QPSK data**

The power spectral density of these modulations can be calculated using Equation 3, where  $f_c$  is the center frequency,  $bw$  is the signal bandwidth, and  $M$  is number of constellation symbols. An example of the modulated carrier over frequency is illustrated in Figure 5. QAM and PSK PSD can both be calculated using Equation 5 as both carriers are typically modulated similarly. This equation provides the ability to estimate the interference seen by a receiver over any given bandwidth at any center frequency. For

instance, if Wi-Fi is currently operating on channel 1 the power of the side lobes interfering with a IEEE 802.15.4 signal on channel 19 can be estimated. This determines if Wi-Fi can transmit and IEEE 802.15.4 can receive concurrently, and is the basis of coexistence on the physical layer.



**Figure 5 - Power spectral density plot of PSK and QAM signal**

Equation 3

$$p_{psk}(f) = \left( \frac{\sin(T_s \pi (f - f_c))}{T_s \pi (f - f_c)} \right)^2$$

Equation 4


$$T_s = \frac{\log_2 M}{b_w}$$

Equation 5

$$p_{psk}(f) = p_{qam}(f)$$

## 4. FSK

Frequency shifted key is a form of digital modulation where in-phase and quadrature frequency  $f_c$  is adjusted to represent a bit, and is illustrated in Figure 6. GFSK is adaptation of FSK that is passed through a Guassian filter to smooth signal allowing for a much more narrow bandwidth.



**Figure 6 - FSK IQ data set example**


The spectral density of a frequency shifting key can be calculated using Equation 6 where  $f_c$  is the center frequency of signal,  $D$  is  $\frac{1}{2}$  distance between mark and space frequencies, and  $b_w$  is the bandwidth of the signal. The distance between the mark and space is directly correlated to the spectral, and as a result different GFSK modulations will have different impact as an interferer, and provide different challenges with coexistence.

Equation 6

$$p_{fsk}(f) = \left[ \frac{1}{D^2 - (f - f_c)^2} \right]^2 \frac{\left[ \cos\left(\frac{2\pi D}{b_w}\right) - \cos\left(\frac{2\pi(f - f_c)}{b_w}\right) \right]^2}{1 - 2 \cos\left(\frac{2\pi D}{b_w}\right) \cos\left(\frac{2\pi(f - f_c)}{b_w}\right) + \cos^2\left(\frac{2\pi D}{b_w}\right)}$$

## 5. OFDM & OFDMA

Orthogonal frequency-division multiplexing is a digital modulation technique where multiple subcarriers are multiplexed over a single bandwidth, and is illustrated in Figure 7 - OFDM Modulated Carrier as the addition of multiple subcarriers. In IEEE 802.11 OFDM implementation consists of PSK, and QAM subcarriers.



**Figure 7 - OFDM Modulated Carrier as the addition of multiple subcarriers**

The power spectral density of the OFDM signal can be calculated by summing the individual subcarriers power spectral density, and is represented by Equation 7 & Equation 8, where  $\Delta f$  is the subcarrier bandwidth, and  $N$  is the number of subcarriers. OFDM will exhibit a different spectral mask than a standalone modulated carrier. An OFDM QPSK signal occupying 20 MHz of bandwidth versus a DSSS QPSK occupying 20 MHz will have different side lobes, that can have coexistence advantages or disadvantages dependent on the relative position spectrally.


Equation 7

$$p_{OFDM}(f) = \sum_{k=-\frac{N}{2}}^{\frac{N}{2}} P_{psk}(f - f_{sc}(k))$$

Equation 8

$$f_{sc}(k) = \frac{k}{|k|} \frac{2|k| - 1}{2} \Delta f$$


Orthogonal frequency-division multiple access is an implementation of OFDM breaking up the subcarriers, and assigning them to individual clients over time, and is illustrated in Figure 8 - OFDMA signal separated into individual subcarriers over frequency and time. The power spectral density of an OFDMA signal can be calculated by determine the PSD of the active sub carriers using equation, and dynamical changes with time depending on the active sub carriers. OFDMA provides the ability to coexist spectrally by utilizing the time domain. The benefit of the small subcarriers, and the ability to dynamical disable them is the additional granularity provided on the side lobes to increase the ability to coexist. Whereas OFDM will be restricted strictly to the channel spacing provided by the standard.



**Figure 8 - OFDMA signal separated into individual subcarriers over frequency and time**


## 6. IEEE 802.11

The IEEE 802.11 standard take advantage of multiple modulation techniques focusing on DSSS and OFDM, and PSK and QAM. Each modulation technique has different requirement for spectral density mask requirements to follow FCC regulations. The spectral density mask for IEEE 802.11 b DSSS signals is illustrated in Figure 9, and the spectral density mask for IEEE 802.11 OFDM signals is illustrated in Figure 10. Depending on the operating modulation, and bandwidth spectral density can be calculated. This calculation is not effected by coding rates or spreading factors. The criteria for calculating spectral density in the 2.4 GHz and 5 GHz is found in Table 2 & . IEEE 802.11 ax spectral density can be calculated using the same criteria, by determining the active subcarriers in the OFDMA modulated carrier. The large bandwidth requirements pose challenges to the smaller bandwidth protocols such as IEEE 802.15.4 and Bluetooth, but supports a large amount of modulations techniques providing flexibility to coexist on the physical layer.



**Figure 9 - Transmit spectral density mask of an IEEE 802.11 DSSS 20 MHz signal**





**Figure 10 - Transmit spectral density mask of an IEEE 802.11 OFDM 20 MHz Signal**

**Table 2 - Criteria for calculating IEEE 802.11 spectral density**

M(bits)	Bw (MHz)	Modulation
2	10, 20, 40, 80, 160	BPSK
4	20, 40, 80, 160	QPSK
16	20, 40, 80, 160	QAM16
64	20, 40, 80, 160	QAM64
256	20, 40, 80, 160	QAM256
1024	20, 40, 80, 160	QAM1024

## 7. IEEE 802.15.4

IEEE 802.15.4 specifies the spectral mask for the DSSS O-QPSK signal, and is illustrated in Figure 11 - IEEE 802.15.4 Transmit spectral density mask. Spectral density for IEEE 802.15.4 in the 2.4 GHz band, and be calculated using Equation 3 with the criteria in Table 3. IEEE 802.15.4 only supports one bandwidth and one modulation in the 2.4 GHz, thus power spectral density can only be calculated one way. Although IEEE 802.15.4 offers sub GHz channels with GFSK modulation. Those frequency bands are not as widely used as the 2.4 GHz frequency band, and are minimally impacted by Bluetooth Low Energy or IEEE 802.11 due to sufficient spectral separation. The 2.4 GHz band does not have flexibility on different modulation techniques, however due to low bandwidth occupancy with many channels interference is minimal in comparison to Wi-Fi.




Figure 11 - IEEE 802.15.4 Transmit spectral density mask

Table 3 - Criteria for calculating IEEE 802.15.4 spectral density

M(bits)	Bw (MHz)	Modulation
4	2	O-QPSK

## 8. Bluetooth Low Energy

Bluetooth Low Energy specifies the spectral mask for FHSS GFSK signal, and is illustrated in fig Spectral density, and be calculated using equation with the criteria in table . Bluetooth offers two PHY's 4-GFSK, for higher throughput, and 2-GFSK. 4 GFSK is only available from BLE 5.0 on. As the data rate is increased the modulated carrier become more spectrally dense. Spreading codes introduced in the BLE 5.0 standard improve BER vs SNR, however have no impact on the spectral density of the modulated carrier. Due to the flexibility of both modulation techniques and low channel bandwidth Bluetooth has multiple advantages in physical layer coexistence.




Figure 12 - Bluetooth Transmit Spectral density mask

Table 4 - Criteria for calculating IEEE Bluetooth Low Energy spectral density

M(bits)	Bw (MHz)	Modulation
2	1	2-GFSK
4	1	4-GFSK

## 9. Theoretical Calculations

The spectral density of a signal is the effective power of a signal over a given bandwidth. Understanding the bandwidth and center frequency of the signal of interest, and the modulation, bandwidth, and center frequency of the interfering signal the SIR can be calculated. IEEE 802.15.4 on Zigbee channel 25, 2475 MHz, SIR measured with respect to a Wi-Fi DSSS signal on channel 1, 2412 MHz, is demonstrated in Equation 9. The results of the SIR calculation can be applied to the SNR curve illustrated in Figure 13 - IEEE 802.11, IEEE 802.15.4, IEEE 802.15.1, Bluetooth, Bit error Rate vs Signal to Noise Ratio to estimate bit error rate. This logic can be applied to any combination of channels, bandwidths, and modulations between 802.11, 802.15.4, and Bluetooth low energy to determine the impact of other in band signals. Having multiple radios consolidated gateway enables a large scale understanding of the entire ecosystem from a physical layer perspective. It provides insight into what operating channels, bandwidth, and modulation techniques are used by every device, unveiling what devices can communicate concurrently. If devices cannot communicate concurrently the gateway acting as the network coordinator has the ability to adjust the modulation, channels, and bandwidth to improve physical layer coexistence. Understanding the SIR between each device on the network is not possible with individual subsystems, and eliminates the ability to effectively manage multiple networks within one home, or determine the impact of bandwidth, modulation, or channel selection on other subsystems.


*Equation 9*

$$\int_{2475\text{MHz}}^{2476\text{MHz}} P_{psk}(f)$$

$$f_c = 2412 \text{ MHz}, 2437 \text{ MHz}, 2462 \text{ MHz}$$

$$b_w = 11 \text{ MHz}, 22 \text{ MHz}$$

$$M = 2, 4$$



**Figure 13 - IEEE 802.11, IEEE 802.15.4, IEEE 802.15.1, Bluetooth, Bit error Rate vs Signal to Noise Ratio**

## Network Coexistence Techniques

Coexistence is a balance of maintaining high throughput for Wi-Fi devices based on optimal links and channel utilization, optimizing battery life for IoT devices by reducing retransmissions, network rejoins, or other battery draining network interactions. In a cable gateway decisions to optimize performance on one subsystem can be influenced by the other sub systems in the gateway. The cable gateway can make decisions not only based on connected devices but also based on other connected devices effecting the ecosystem, such as other Wi-Fi or IoT devices in a multi-family home. Although the gateway can't directly impact the non-connected devices it pass on that data to the other sub-systems. For instance, if, the gateway's Wi-Fi is occupying channel 1 in 2.4 GHz band, but has detected other Wi-Fi networks on channel 6, Bluetooth, and IEEE 802.15.4 can adjust channel selection accordingly. Similar decisions can be made when scheduling free air time, adjusting modulation or bandwidth, and these decisions can be made by understanding the SIR between the networks.

## 10. Network Management

### 10.1. Scheduling

Bluetooth Low Energy, Zigbee, and Thread end devices typically are sleepy to preserve battery life. These sleepy end devices send status updates to the network coordinator on an interval basis that free air time can be scheduled to mitigate interference with Wi-Fi. Bluetooth Low Energy operates on an interval basis, transferring data through that interval, and then switching channels. Dynamically scheduling freeing RU's in 802.11 ax based on BLE current operating channels opens opportunities to have free air time in the network. An example of the RU's available in a 20 MHz IEEE 802.11 ax channel is illustrated in Figure 14. Increasing BLE throughput and increasing battery life. Using the power density calculations it can be determine how many RU's need to be inactive so the current IEEE 802.15.4 and BLE operating channels will have sufficient SIR to get the packet to be recieved. Older 802.11 standards are more restrictive requiring that the free airtime be scheduled, having an negative impact on throughput. Essentially, scheduling is determined by understanding what devices cannot communicate concurrently based on SIR, ensuring that those devices are given separate air time, and that all protocol timing restrictions are met. An individual subsystem not only does not have access to the the ecosystem SIR, but also is strictly limited to the scheduling and timing restrictions of the devices that it manages. Due to the inability to consider scheduling and SIR constraints of other individual sub systems, consolidated gateways have scheduling performance advantages.




Figure 14 - IEEE 802.11 ax 20 MHz channel resource units

### 10.2. Channel Selection

Ideal channel selection for IEEE 802.11 and IEEE 802.15.4 would be spectrally separating as much as possible. Bluetooth low energy ideal channel mask would include a list of all channels that spectrally separated enough that the device has sufficient SIR to receive a packet. In a single family with no neighbors nearby, coexistence may be as simple as that. However, in congested environments with networks competing for air time, channel selection may not be as straight forward. If the current operating IEEE 802.11 channel is to congested, it may be desirable to move to a less congested channel. However, the channel will adjust the SIR of the other devices connected to gateway, and as result should be part of the overall decision making process on channel changes, including adjusting channel bandwidth. For instance, an IEEE 802.11 network with no competing networks operating on channel 1 2412 MHz at 20 MHz with sufficient SNR for a 40 MHz. If the IEEE 802.15.4 and BLE networks are operating on

channel 17, 2435 MHz, and channel 15, 2436 MHz they would then be co-channel with the IEEE 802.11 network, and as a result will most likely not have sufficient SIR resulting retransmission, or loss of connectivity to the network resulting in negative impacts on battery life of those devices. Channel selection is made by understanding the channel selection of all networks being managed by the gateway, as well as competing unmanaged networks. A combination of SIR, and free air time is used to determine the channel that will be most effective for the networks being managed by the gateway.

## 11. Prioritization


Determining which subsystem has overall priority can be simple or complex depending on the ecosystem that the gateway is placed in. For example, if a cableway is supporting a Home Security network while that network is armed, those devices may have priority, or while streaming video content the Wi-Fi system may have priority. A complex dynamic approach to priority provides the opportunity to improve performance on a case by case basis, but requires insight into other networks, and their priorities. Priority can be applied to channel selection as well scheduling, and typically is determined by importance. Optimal channels can be selected for devices with the highest priority, but this decision can only be made by understanding the current state of the entire ecosystem, and each network managed by the gateway. Individual subsystem's dynamic priority capabilities are restricted in comparison to a consolidated gateway solution simply due to the lack access or control other individual subsystems.

## 12. Spatial Mapping

Understanding relative distance between connected devices and the gateway as well relative distance between connected devices is key for optimizing channel selection as well as scheduling. Assume a single IEEE 802.15.4, and IEEE 802.11 device are connected to the gateway line of sight. The IEEE 802.15.4 has an RSSI of -37 dBm, and IEEE 802.11 has an RSSI of -40 dBm. Assume both devices have a perfectly omnidirectional TRP, total radiated power, of +20dBm, implying the IEEE 802.15.4 is experiencing 57dB of free space path loss, 7 meters using equation Equation 10, and the IEEE 802.11 is experiencing 60 dB free space pathloss, 10 meters assuming equation. Worst case scenario both devices are directly in the same line of sight with distance of 50 dB, 3 meters, between each other, however best case scenario the devices are 64 dB, 17 meters, apart on the same line of sight in the opposite axis spatially. In the first scenario where the devices are 3 meters away they are spatially very close, and therefore will need to be separated spectrally. However, in the second scenario the devices are 17 meters away with 64 dB of pathloss between them, and as a result may not need to be as separated spectrally. Using phase differences between antenna's on different IEEE 802.11, IEEE 802.15.4, and BLE devices the AoA, angle of arrival, of other devices be detected, and is illustrated in Figure 15. By mapping that information to devices addresses and RSSI data, each network can create a map of devices spread throughout an ecosystem. This map of all the devices is key to the gateway determining best operating channels, modulation, and bandwidth for each network as it can calculate the SIR between every device in the ecosystem. Individual subsystems can apply a similar technique however, the subsystem will be limited to mapping devices of the same protocol impacting the the ability to coexist.

*Equation 10*

$$FSPL_{dB} = 20 \log_{10} \frac{4\pi df}{c}$$



**Figure 15 - Angle of Arrival calculation using multiple antennae**

## Conclusion

Consolidating IoT traffic through the cable gateway has significant impact on system performance. Reducing IoT device retransmissions with minimal impact on IEEE 802.11 improves device battery life, reduces latency, resulting in an overall improved user experience. These improvements are achieved by optimizing network performance by understanding physical coexistence and SIR of the devices in network. SIR can be collected on individual IoT subsystems, however it cannot necessarily be correlated to other systems SIR, or spatially correlation. More importantly changes cannot be made to those out of network devices systems from that network. The lack of the correlation and control between other networks precludes optimal performance. Continued consolidation of IoT traffic into the gateway will continue to optimize the entire whole home experience.

## Abbreviations

AP	access point
BER	Bit Error Rate
BLE	bluetooth low energy
bps	bits per second
BPSK	binary phase shift keying
CER	chipping error rate
DSSS	Direct Sequence Spread Spectrum
FCC	Federal Commuication Commission
FHSS	frequency hopping spread spectrum
FSK	frequency shift keying
GFSK	gaussian frequency shift keying
IEEE	Institute of Electrical Electronics Engineers
ISM	industrial, scientific and medical
IoT	Internet of Things
IQ	inphase and quadrature

Hz	hertz
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multi access
O-QPSK	offset quadrature phase shift keying
PSD	power spectral density
PSK	phase shift keying
QAM	Quadature Amplitude Modulation
QPSK	quadature phase shift keying
RU	resource unit

## Bibliography & References

Anderson, David, and David Cohen. *NTIA 93-298, Digital Emission Spectral Model*. US Department of Commerce, 1993, pp. 1–63, *NTIA 93-298, Digital Emission Spectral Model*.



# **Practical Lessons of a DAA Deployment with a Virtualized CMTS**

A Technical Paper prepared for SCTE•ISBE by

**Asaf Matatyaou**

Vice President, Solutions and Product Management, Cable Access Business  
Harmonic, Inc.  
4300 North First Street, San Jose, CA 95134  
1-408-490-6834  
Asaf.Matatyaou@harmonicinc.com

**Brian Bendt**

Engineering Director, P.E.  
Comporium  
471 Lakeshore Parkway, Rock Hill, SC 29730  
1-803-326-7173  
Brian.Bendt@Comporium.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Architecture Description.....	3
Deployment Details .....	5
Real-World Considerations .....	7
1. Services.....	7
2. Compute Location Options.....	8
3. Converged Interconnect Network .....	9
3.1. CIN Switch Location Options .....	10
3.2. Traffic Prioritization and Capacity Management .....	10
4. Timing.....	11
5. Access Network Maintenance and Monitoring .....	12
6. Operations.....	13
Conclusion .....	14
7. Lessons Learned .....	14
8. Summary .....	14
Abbreviations.....	15
Bibliography & References .....	15

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Traditional HFC, Centralized and DAA Deployment Comparison.....	4
Figure 2 - vCMTS Deployed in a Distributed Access Architecture.....	6
Figure 3 - A Traditional HFC Deployment Transmits and Receives RF to Analog Optical Nodes .....	8
Figure 4 - Upstream Spectrum Analysis and Sweep with DAA .....	13

# Introduction

The promise and potential of virtualizing a cable hub has been discussed over the past few years. There are many opportunities when discussing virtualization. In 2018, Harmonic presented one possible starting point, by virtualizing a CMTS in a centralized deployment model in a headend or hub. The focus of this paper will be specific to a Distributed Access Architecture (DAA) deployment, leveraging the same approach of virtualizing the CMTS.

While the benefits and general considerations of a virtualized CMTS (vCMTS) have been described in past technical papers, this paper will focus on real-world experiences and lessons learned deploying a vCMTS in a DAA. The combination of virtualization and the use of Remote PHY as a protocol between a vCMTS Core in a centralized facility and distributed Remote PHY Devices (RPDs), when migrating from a hardware-based CMTS deployed on a traditional HFC infrastructure, will be described. Important areas of consideration when deploying DAA will include usage of IEEE-1588 over aware and unaware networks, as well as the Converged Interconnect Network (CIN) connecting the vCMTS with the RPDs. Supporting legacy services, such as video and out-of-band, and operational aspects necessary to field-deploy DAA, such as leakage and upstream spectrum, will be described. The paper will also examine the differences between DAA and a traditional HFC deployment, where the RF is generated in the headend or hub in comparison to the Remote PHY node. Lastly, the usage of streaming telemetry will be explored in comparison to legacy monitoring techniques, paying particular attention to important metrics that are gathered by the RPDs.

## Architecture Description

The architecture described in this paper is made possible by advancements in standards and technology, specifically the Remote PHY specifications published by CableLabs and the virtualization of the CMTS into commercial off-the-shelf (COTS) x86-based servers. CableLabs issued the first set of Remote PHY specifications in June 2015 and most recently updated the specifications in March 2019 with the twelfth revision of the specifications. [8]

“In a Remote PHY Architecture, the classic integrated CCAP (I-CCAP) is separated into two distinct components. The first component is the CCAP Core and the second component is the Remote PHY Device (RPD).”<sup>1</sup> The RPD consists of the physical layer functionality defined for an I-CCAP, with the remainder of the I-CCAP functionality residing in the CCAP Core. The CCAP Core is logically the combination of a CMTS Core and EQAM Core, and is connected to the RPD via IP transported over digital fiber. In this paper, the CCAP Core is implemented as a vCMTS with a separate legacy EQAM pre-existing in the operator’s network and will be referred to as “vCMTS.”

Since 2015, SCTE technical papers, such as “Transforming the HFC Access Network with a Software-Based CCAP” [4] and “Real-World Deployment of a Virtual Cable Hub” [3], have defined virtualization and the benefits of a vCMTS. The benefits and general considerations for vCMTS are beyond the scope of this paper.

“Remote PHY” is an implementation of a DAA, but not necessarily restricted to DAA deployments. In fact, the term “remote” doesn’t restrict the RPD from physically being co-located with the vCMTS, and the CableLabs specifications describe examples where the RPD and RF may be located in the network, in

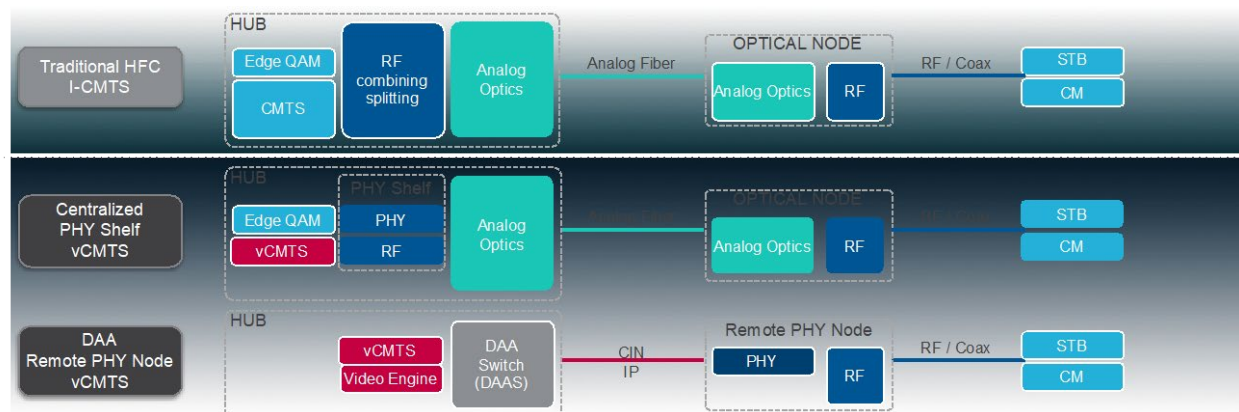
---

<sup>1</sup> Remote PHY Specifications, CM-SP-R-PHY-I10-180509, pg. 10

the headend/hub or in an optical node. The “remote” aspect is that the physical layer or PHY is separate or remote from the CMTS Core.

Figure 1 shows an I-CCAP deployment architecture, as well as Remote PHY-based centralized and DAA deployment architectures, both of which use the Remote PHY standard signaling to communicate between the vCMTS and the RPD (existing in the PHY shelf and Remote PHY Node (RPN)). Remote PHY signaling includes the Downstream External PHY Interface (DEPI), Upstream External PHY Interface (UEPI) and the Generic Control Plane (GCP).

The benefits of Remote PHY and a detailed description of the specifications, including the signaling, are beyond the scope of this paper and are well documented in the industry over the past few years.



**Figure 1 - Traditional HFC, Centralized and DAA Deployment Comparison**

This paper focuses on a DAA deployment, whereby the RPDs are deployed in outdoor Remote PHY Nodes. The terms “Distributed” or “DAA” are used to describe this deployment architecture as the RPDs are deployed in a distributed fashion, deeper in the HFC infrastructure. Per the CableLabs DAA overview, this involves “moving functions into the network reduces the amount of hardware the headend (hub) needs to house, thus creating efficiencies in speed, reliability, latency and security in support of 10G.”<sup>2</sup> A traditional deployment with I-CMTS is based on chassis deployed in each hub location within the operator’s footprint. This is the common architecture today. In Figure 1, this means that the top drawing is replicated over and over again. This was a necessary approach at the dawn of DOCSIS services, but technology has evolved with DAA and vCMTS, enabling much more efficient implementations, while simultaneously being more scalable and capable.

This traditional I-CMTS deployment was the starting point before the transformation to deploy vCMTS and Remote PHY Nodes. Cable operators continue to increase their bandwidth and service offerings to keep up with subscriber usage demands and telco competition. Continuous bandwidth growth has historically created challenges when deploying in a legacy I-CMTS deployment, such as:

- Equipment requirements driven by repeated node splits, including RF splitters and combiners, and optical transmitters and receivers
- Substantial equipment space and power requirements leading to facility/building expansions to accommodate more legacy optical nodes and systems
- Substantial headend cabling required for each of the nodes, impacting space and airflow
- Financial and space requirements needed to distribute equipment to additional locations

<sup>2</sup> <https://www.cablelabs.com/technologies/distributed-access-architecture>, July 27, 2019.

- Financial requirements to provide fiber path diversity from headend or hub to legacy optical nodes, if required

On the other hand, benefits driving the transformation to a vCMTS and DAA deployment include<sup>3</sup>:

- Reduced space, power, wiring and cooling
- Speed to complete node splits
- More frequent and shorter development cycles
- Sustainable capacity growth, elastic scalability and increased flexibility
- Improved total cost of ownership, including reduced operational (opex) and capital expenditure (capex)

In this deployment example, the vCMTS running on COTS x86-based 1-RU servers replaced legacy I-CMTS chassis in the same locations. The Remote PHY Nodes were deployed in place of legacy analog optical nodes and connected back to the vCMTS Cores over the Converged Interconnect Network or CIN. This approach allows operators to deploy the Remote PHY Nodes in existing locations or while performing node segmentation, leveraging the existing HFC infrastructure (i.e. leveraging existing amplifiers downstream of the node).

## Deployment Details

The deployment details describe the ending point after replacing the legacy I-CMTS chassis with the vCMTS Core servers, looking at a specific hub location and taking into account all data, video, out-of-band and proactive network maintenance services required for a production deployment.

Let's describe each device in this deployment type and where the device is located.

1. vCMTS: the CMTS Core functionality is implemented on a set of COTS x86-based servers. Each vCMTS Core replaced legacy I-CMTSes at each location and are connected to the Converged Interconnect Network (CIN).
2. Video Engine: functionally similar to an edge QAM, supporting broadcast video and video-on-demand. However, a Video Engine outputs video services over an IP-based L2TPv3 (i.e. DEPI) tunnel, over the CIN, to each RPD. In this deployment, there is no additional statmux, rate shaping or video processing performed by the Video Engine.
3. NDF/R Engine: narrowband digital fiber and return "digitizes a small portion of the spectrum, and sends the digital samples as payload within packets that traverse between the [vCMTS] and the RPD. This approach works with any type of OOB signal as long as the signal can be contained within the defined pass bands."<sup>4</sup>
4. QPSK Modulator/Demodulator: legacy SCTE 55-2 equipment that is used unchanged from legacy deployments. In this deployment, the existing SCTE 55-2 equipment is used in combination with the NDF/R Engine, which digitizes the RF to IP packets traversing over the CIN.
5. Core Routing Engines (CRE): standard COTS networking equipment are the switch fabric connecting the vCMTS Core servers to the rest of the CIN. The CREs are co-located in the same location as the vCMTS Core servers and are part of the Layer 2 CIN. They also perform the function of aggregation, sometimes referred to as the Distributed Access Architecture Switch (DAAS).

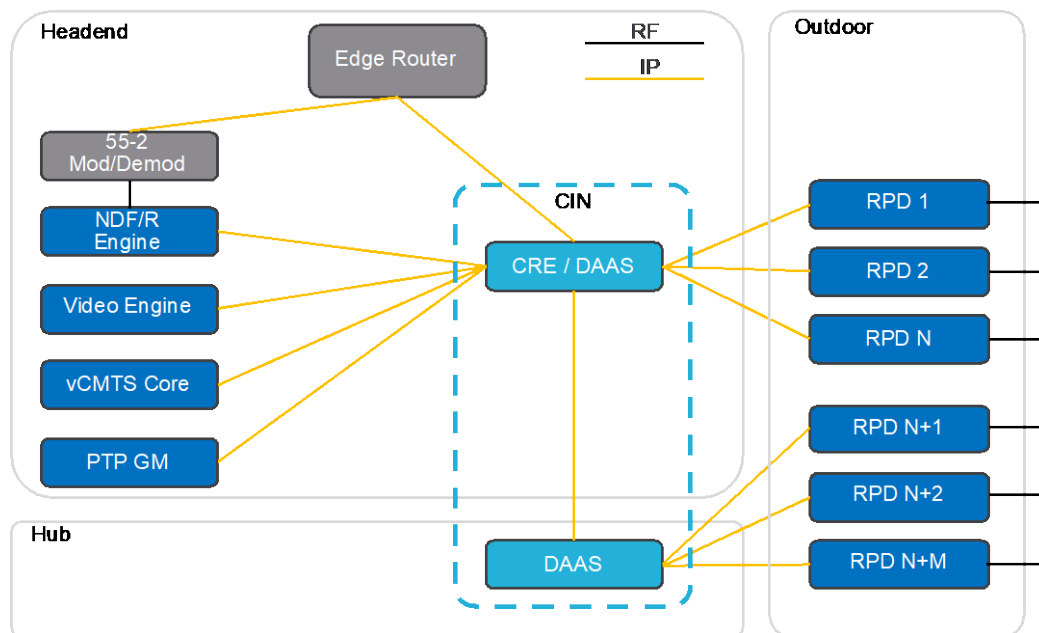
<sup>3</sup> Real-World Deployment of a Virtual Cable Hub, pg. 5

<sup>4</sup> Remote Out-of-Band Specification, CM-SP-R-OOB-I09-180509, pg. 14

6. Edge routers: large-scale edge routers connecting the access network with the core backbone network, which are typically previously established.
7. Remote PHY Nodes and RPDs: each RPN outdoor enclosure can house one or two RPDs. Each RPD connects to the vCMTS Core servers over the CIN, and outputs RF over the existing HFC infrastructure. A test RPN (without production subscribers) is deployed at each hub.
8. IEEE-1588 PTP Grandmasters: “Remote DTI provides timing synchronization between CCAP Cores and RPDs based on the IEEE 1588v2 standard. The protocol supports the basic synchronization between the CCAP Core and Remote PHY Device for DOCSIS/video/OOB services.”<sup>5</sup>

Figure 2 shows the devices and connectivity between the vCMTS Core, Video Engine and NDF/R Engine over the CIN, which is the network between the CCAP Core and the RPD. In this deployment type, the CIN traverses between the headend, each hub and the RPDs over a Layer 2 network.

One of the benefits in this deployment type is that the CIN infrastructure is transported over digital fiber which has the benefit of longer distances, more wavelengths and increased performance compared with legacy HFC transport. While the RPDs deployed in this deployment leverage the existing HFC infrastructure south of the Remote PHY Node (i.e. existing external amplifiers), the same CIN can be used to deploy fiber deep Remote PHY Nodes (i.e. high-power amplifiers are built into the node and no additional external amplifiers are needed downstream).



**Figure 2 - vCMTS Deployed in a Distributed Access Architecture**

<sup>5</sup> Remote PHY Specifications, CM-SP-R-PHY-I12-190307, section 5.5.2, pg. 30

# Real-World Considerations

While there are many benefits and opportunities with this new Remote PHY-enabled vCMTS in a DAA deployment, there are real-world considerations that should be considered. In particular, the following topics are described: services, compute location, networking, timing, network maintenance and operations.

## 1. Services

Cable operators deliver multiple services to subscribers, notably broadband, video and voice services. Delivery of the different services to a subscriber's cable modem, set-top box and other IP-based devices is generated and processed by different headend equipment, transported over the HFC network as RF signals and use different technologies and protocols.

In a traditional deployment, each service is processed and generated at the headend with RF output over a non-conflicting frequency span (with other services). For example:

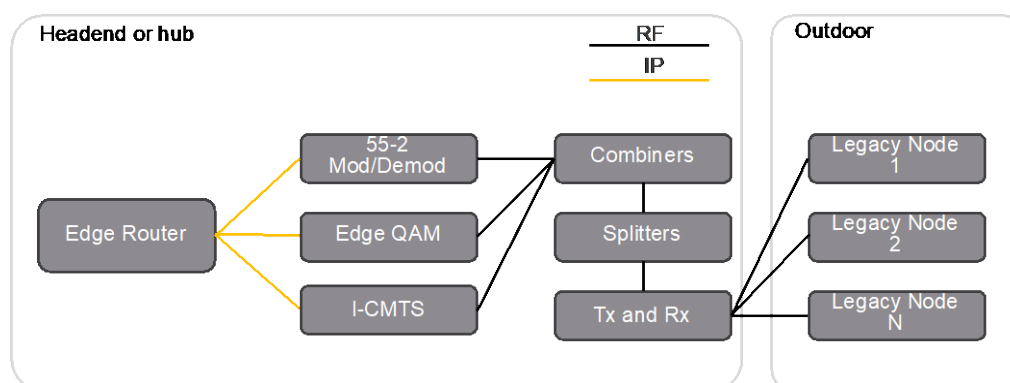
- Broadband and voice: an I-CMTS transmits and receives RF for DOCSIS-based data for broadband and voice services.
- Broadcast video and video-on-demand: an edge QAM transmits RF for broadcast video and video-on-demand.
- Set-top box control: Out-of-band systems (e.g. SCTE 55-1 and SCTE 55-2) transmit and receive RF for controlling set-top boxes.
- Network maintenance and monitoring: Test equipment transmit and receive RF for network maintenance and monitoring purposes.

The RF for each of the services is combined at the headend, or at each hub, to have all services available when received at each subscriber's devices. After the RF is combined, it is then split to different service groups and analog optical nodes. Additionally, transmitters in the headend send the RF signals to the analog optical nodes as well as amplified over the traditional HFC infrastructure. In the reverse direction, receivers in the headend acquire the upstream RF signals.

This is a simplified description of how, for decades, subscribers receive all services over a single coaxial cable in the home, as well as what equipment is used by the cable operator to deliver these services. This description is important to understand, specifically for the sake of contrasting service delivery and transport over a traditional HFC infrastructure and DAA.

A key difference between a traditional HFC deployment and DAA:

- A traditional HFC deployment transmits and receives all services from the headend over RF, as shown in Figure 3
- A DAA deployment transmits and receives all services from the headend via IP over Ethernet



**Figure 3 - A Traditional HFC Deployment Transmits and Receives RF to Analog Optical Nodes**

With DAA, and unlike a traditional deployment, all the services and signals are transported over an Ethernet-based IP network, the CIN. Transporting over the CIN means that what was once transmitted over RF from the headend or hub, is now transported over IP, eliminating the splitters, combiners, transmitters and receivers at the headend or hub. The facility savings (i.e., equipment, space, power and cooling savings) is substantial, as well as the savings due to reduction in ongoing cabling, maintenance and operational expenses associated with continuous bandwidth growth and annual expansion.

Let's briefly revisit each service described in a traditional deployment and the changes required to deliver each with DAA, over the CIN:

- Broadband and voice: A CMTS Core (in this case, a vCMTS) transmits and receives IP for DOCSIS and Remote PHY-based data for broadband and voice services.
- Broadcast video and video-on-demand: A Video Engine transmits Remote PHY-based IP for broadcast video and video-on-demand, and is independent of the vCMTS Core and Out-of-band engines.
- Set-top box control: Out-of-band systems (e.g. SCTE 55-1 and SCTE 55-2) transmit and receive Remote PHY-based IP for controlling set-top boxes. Alternatively, out-of-band systems can transmit and receive RF, similar to traditional deployments, with an NDF/R Engine performing the conversion from RF to Remote PHY-based IP and vice versa.
- Network maintenance and monitoring: Test equipment transmit and receive Remote-PHY IP for network maintenance and monitoring purposes. Additionally, some RF signals such as pilots and alignment carriers can be generated at the Remote PHY Device.

Remote PHY-based IP refers to the different Remote PHY protocols specified by CableLabs, specifically, Downstream External PHY Interface (DEPI), Upstream External PHY Interface (UEPI) and the Generic Control Plane (GCP).

## 2. Compute Location Options

One of the key opportunities when transitioning to a DAA with a vCMTS Core is where the CMTS Core is located, enabled by two technologies: Remote PHY and virtualization. The Remote PHY specifications enable separating the CMTS Core from the RF and provide the operator an opportunity to decide where the CMTS Core should be located. I-CMTS equipment perform DOCSIS processing, transmit and receive RF in the same chassis. There is no choice in location selection as I-CMTSs are deployed at each hub location connected to the access network, regardless of the size and scale required to support the



nearby footprint of subscribers. The Remote PHY specifications enabled separating the CMTS Core from the RF and provide the operator an opportunity to decide where the CMTS Core should be located.

The second technology enabler for flexibility in CMTS Core location is virtualization. A vCMTS Core solution provides the operator an opportunity to determine where the vCMTS Cores should be located, with options such as installing the vCMTS Cores at each hub location with the RF (similar to I-CCAP) or consolidating the vCMTS Cores at a few hubs or even a single centralized hub location. Having at least two vCMTS Core server locations may also provide geographical redundancy.

In this real-world deployment example, the vCMTS Core servers were deployed in the same locations as the legacy I-CMTS chassis. This approach allowed the operator to reduce space, power and cooling requirements at existing facilities while expanding their capability to deploy faster broadband services. Additionally, it was the simplest transition strategy from I-CMTS to vCMTS and DAA, as it required the least amount of change. The same replacement approach, installing Remote PHY Nodes at the same points in the network and HFC infrastructure as legacy analog optical nodes was used when transitioning from analog optical nodes to Remote PHY Nodes.

Expandability with more Remote PHY Nodes can be accomplished by adding each node to an existing vCMTS Core server with spare capacity, via the DAAS or by adding extra vCMTS Core servers for more capacity.

“More flexibility in deployment location and scalability requires increased attention on the scalability of each element in the end-to-end network. While expandability is easier and provides quicker time-to-market to add capacity, operators must pay attention to the scale limits of the compute and networking resources separately, as each type of resource type may require additional devices when those limits are reached. On the other hand, when scale limits are reached for a particular resource, those can be expanded in a focused method and the operator won’t have to scale everything at once. For example, when I-CCAP chassis reaches any of its scale limitations, another I-CCAP needs to be installed. In comparison, when a DAAS reaches a scale limitation, such as port count, another DAAS switch can be added without adding additional vCMTS Core servers. Nevertheless, an operator must pay attention to each element’s specifications and plan a network for existing and additional subscribers.”<sup>6</sup>

### 3. Converged Interconnect Network

It’s worth repeating, that a key difference between a traditional HFC deployment and DAA:

- A traditional HFC deployment transmits and receives all services from the headend over RF
- A DAA deployment transmits and receives all services from the headend using IP over Ethernet

“The network between the CCAP Core and the RPD is known as the Converged Interconnect Network (CIN). The CIN encompasses either or both the hub access network and the optical access network. The CIN can contain both Layer 2 switches and Layer 3 routers.”<sup>7</sup>

While the benefits of an Ethernet-based IP network (digital fiber) have been noted, there are a few CIN related decisions that need to be made:

1. Location and number of CIN switches between the vCMTS Core and the RPD
2. Layer 2 or Layer 3 configuration

---

<sup>6</sup> Practical Lessons of a Centralized Virtualized CMTS, pg. 8.

<sup>7</sup> Remote PHY Specifications, CM-SP-R-PHY-I12-190307, section 5.2.5, pg. 26

3. Traffic prioritization
4. Redundancy: Network redundancy is another deployment decision with many options available, including, link redundancy, chassis redundancy, and line card redundancy within a chassis. This topic is beyond the scope of this paper.

### **3.1. CIN Switch Location Options**

The simplest way to describe the CIN is as a logical network switch that connects the vCMTS Core, Video Engine and OOB Engine with the RPDs. In fact, the CIN can be deployed as a single switch co-located with a few ports connected to the vCMTS Core, Video Engine and OOB Engine, and many ports connected to RPDs. In the case of a vCMTS deployment, there are three types of traffic sources/destinations that are supported via the CIN:

- The vCMTS Core, Video Engine and OOB Engine
- The RPDs
- The Edge Router – the same one connected to legacy I-CMTSes

The CIN serves as the fabric that connects the northbound content sources, the devices performing the data and video processing (i.e. vCMTS Core, Video Engine) and the southbound RPDs. Consequently, the CIN can also be deployed as a collection of switches, creating this fabric.

The decision of whether to directly connect the Edge Router, vCMTS Core and RPDs to the same CIN switches is determined on a few factors:

1. An existing network infrastructure: it's possible to leverage an existing network to transport the CIN traffic.
2. Facility location of Edge Router, vCMTS Core, Video Engine and OOB Engine: there is typically a switch co-located with the vCMTS Core, Video Engine and OOB Engine.
3. Distances between devices: switches can be used to extend the distances between the vCMTS Core and the RPDs, if necessary, as well as provide fiber route redundancy.
4. Number of devices needing to connect to each switch: the main factor of switch ports is the number of RPDs, factoring in the initial set of installed RPDs and potential future RPD growth.

In this deployment, there are two locations for the CIN switches: centrally located at the headend and at remote hubs. The CIN switch located at the headend is connected to the Edge Router, Video Engine and OOB Engine. The CIN switch located at the remote hub is connected to the headend CIN switch, the vCMTS Core and the RPDs. In this case, the headend CIN switch partially performs the function of the CRE (the networking fabric between the Edge Router and the multiple vCMTS Core servers), while the remote hub switch partially performs the function of the CRE and is the DAAS (physical connectivity to the RPDs).

In this deployment, layer 2 networking was selected, as it was the quickest and simplest to deploy with CIN switches deployed at only two locations: at the headend and hub.

### **3.2. Traffic Prioritization and Capacity Management**

Remote PHY and DOCSIS demand that latency sensitive traffic receives the correct priority and handling in the CIN. Specifically, IEEE 1588 PTP packets used for timing synchronization between the vCMTS Core and the RPDs need to meet latency and jitter requirements based on how PTP is configured and deployed between the Timing Server Grandmaster and each PTP slave device (e.g., vCMTS Core, Video

Engine, RPD). DOCSIS MAC management messages such as the MAP also have real-time timing information and require proper traffic priority.

In an isolated CIN, where there is no other unrelated network traffic handled by the CIN switches, the design considerations are simpler, as congestion can be managed with good network design and capacity management practice, especially when planning expanding CIN network resources, before the congestion actually occurs in deployment. However, if the CIN can't be guaranteed to be congestion free, capacity management isn't sufficient. Congestion management is required.

If congestion does occur, congestion management is performed by traffic prioritization. Traffic prioritization effectively allows higher priority control and user packets to survive the congestion. There are multiple ways DOCSIS and IP-based traffic prioritization can be used to maintain the Quality of Service (QoS) when manageable congestion occurs in the CIN. However, if congestion reaches the point that it is impactful to critical control packets, such as DOCSIS MAP MAC Management Messages (MMMs), no amount of traffic prioritization will resolve the impact on stable network operations.

The Remote PHY specification also guides: “to prevent the MAP from being slowed down by other traffic in the CIN, the DOCSIS traffic (or a subset containing the MAP messages) may be sent in an independent L2TPv3 flow that can have a unique DSCP. The value of the marked DSCP value should be consistent with a configured "per hop behavior (PHB)" that will provide MAP messages with the highest priority and lowest latency across the CIN to the RPD.”<sup>8</sup>

## 4. Timing<sup>9</sup>

With all Remote PHY deployments, timing specifications such as R-DTI need to be adhered to, regardless of the location of the RPD. “The MHAv2 version of DTI (i.e., R-DTI) defines how to distribute phase and frequency information from the CCAP Core device to remote PHY devices within the HFC network.”

“For Ethernet based networks, IEEE 1588 allows both phase and frequency information to be transferred between nodes across an existing packet network with switches or routers, thus making it ideal for R-DTI.

In order to reduce any phase offset introduced by latencies through the network, IEEE 1588 defines a protocol for calculating the latency across sections of the network, and then compensating for those latencies. The latency calculations assume that the link is symmetric, and therefore the protocol works well for traditional full duplex Ethernet networks. IEEE 1588 also defines a protocol for determining the latency through any intervening switches or routers within the network, but the device is to be IEEE 1588 capable [referred to as PTP aware]. If the devices are not IEEE 1588 capable, the phase offsets and convergence times within the network will be greater [referred to as PTP unaware].”<sup>10</sup>

Devices in the network which are not IEEE 1588 capable are considered as non-participating. For the IEEE 1588 capable devices, they can operate as either a transparent clock or boundary clock. A transparent clock device modifies the PTP message, accounting for the processing time of the PTP message within the device, while a boundary clock receives the PTP message as a slave device and re-generates the PTP message as a master.

This real-world deployment experience is across a PTP unaware network, as not all existing network devices were IEEE 1588 capable. While the convergence times within the network are greater, as

---

<sup>8</sup> Remote PHY Specifications, CM-SP-R-PHY-I12-190307, section 5.6, pg. 30

<sup>9</sup> Practical Lessons of a Centralized Virtualized CMTS, section 3.1, pg. 9.

<sup>10</sup> Remote DOCSIS Timing Interface, CM-SP-R-DTI-I07-180509, pg. 6, 16

expected, they have not been operationally significant to justify immediate replacement of all network devices to be IEEE 1588 capable. However, it is critical to evaluate and consider the jitter and latency conditions of the CIN regardless of selecting an PTP aware or unaware mode. Both PTP deployment modes demand meeting jitter and latency requirements, which can be impacted by the CIN network device capabilities, and the number of network hops and congestion conditions.

As timing is critical for Remote PHY operation, there are a couple of other options to consider when deploying IEEE 1588 PTP grandmaster(s) in the network, which transmit the synchronization information to the other clocks in the same network. The first option is the reliability of the PTP grandmaster, as there exist a range of products which are small (SFP form factor) and less reliable as compared to full carrier-grade products which have redundant input/output clock (IOC) cards. The second option to consider is whether to use the best master clock algorithm (BMCA), which determines the highest quality or “best” clock within the network, in case the grandmaster clock quality is compromised or fails.

In this deployment, two carrier-grade grandmasters (each independently deployed at different headends) with high availability features were deployed. High availability features include redundant inputs, outputs, clock and power supplies.

## **5. Access Network Maintenance and Monitoring**

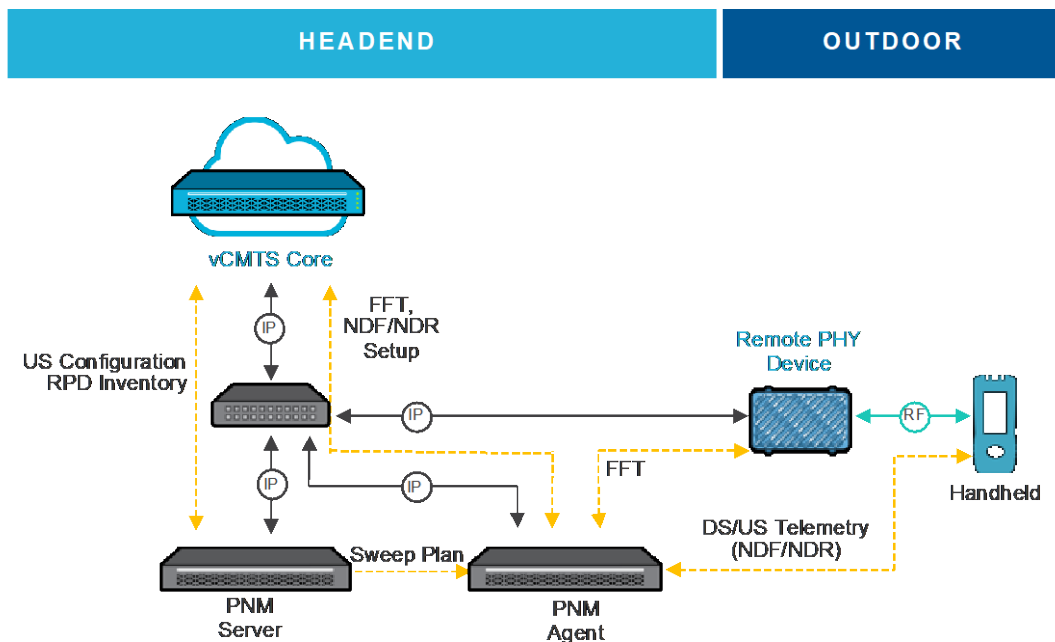
Access network maintenance and monitoring is vital to effectively install, maintain and monitor the end-to-end system. Performing this has historically been performed by injecting or inspecting RF signals with test or monitoring equipment in the headend or in the field. With DAA, RF inspection is no longer available at the headend, as there is simply no RF to inspect with all services running over an IP-based CIN.

Typical proactive network maintenance (PNM) includes upstream spectrum analysis, return sweep, ingress detection and common path distortion. Typical test signals that need to be injected include leakage detection, AGC pilots and alignment carriers. With DAA, these test signals can be generated at the RPD. However, other test signals can also be converted from RF to IP and vice versa with an NDF/R Engine so that they can be transported over the CIN.

In a DAA deployment, PNM examples of the upstream spectrum analysis and return sweep include the following devices:

- vCMTS Core: configures the NDF/R functionality for FFT data to be transported from the RPD to the PNM server.
- RPD: performs FFT processing and transmits FFT data over the CIN to the PNM server.
- CIN: the IP-based network connecting the different devices
- PNM Server: processes the FFT data to display upstream spectrum and sweep results
- Handheld device: telemetry to/from via NDF/R

Similar types of processing can be used to perform other types of PNM, such as CPD and ingress detection. Note that, as an outside plant device supported by maintenance technicians with deep backgrounds focused on RF technology, there is an operational transition period. Some fundamental expectations of prior generation devices should be anticipated, such as RF test points and modular design, in order to ease this transition.



**Figure 4 - Upstream Spectrum Analysis and Sweep with DAA**

## 6. Operations

New technologies such as Remote PHY and virtualization bring many benefits, but they must be supported by real-world operational practices that enable deployment. In the experience of deploying a vCMITS in a DAA, there is an immediate opportunity and need to start by consolidating monitoring of the vCMITS Cores, the CIN and RPDs in a single consolidated operational view. With the CIN as the fabric connecting all the devices the access network, the data can be gathered and visualized with other modern technologies.

Streaming telemetry data coupled with Grafana, as an open visualization tool for analytics and monitoring, is used to give operational organizations a continuous stream of health and performance for the system. The vCMITS streams its data, as well as the RPD data it gathers via the Remote PHY GCP. Other vital system information is gathered and displayed via customizable dashboards. With devices distributed in the network, having a single tool that visualizes a continuous stream of data across topic-based dashboards allows DAA to be operationalized in scale. This is a departure from the tried-and-true command-line-interface CLI-based “show” approach, which is limited to the text being displayed on a screen, a human being inspecting the output at a single instance in time. CLI still has its place for troubleshooting very specific issues but is also isolated to connecting to individual devices and piecing together the story to get a complete understanding of the end-to-end system health and performance.

With a software-based virtualized CMTS Core, there are many points of inspection in the software that can be made visible based on field experience and extended over time, based on new findings. Continuous monitoring improvements will also be augmented over time with configuration, deployment and automation capabilities. In addition, future operations, services, or CMTS upgrades – channel re-mapping, adding DOCSIS carriers, PMA optimization, Full Duplex DOCSIS (FDX) – all can take advantage of the significantly enhanced platform agility to increase the velocity of introduction of such changes, features, and services.

# Conclusion

## 7. Lessons Learned

Looking back at the field experiences from deploying vCMTS in a DAA, there are some key lessons learned. vCMTS has been deployed in different environments over the past few years, and consequently most of the lessons learned recently are related to scaling DAA deployments and specific operational aspects related to moving from a lab to the field. In particular:

1. New technologies require substantial lab testing.
2. DAA technologies may cross typical operational group responsibility boundaries, including installation and support. Operational responsibilities require planning and discussion amongst different organization before deployment.
3. Out-of-band designs, such as SCTE 55-2, need to be planned. Specifically, there are options for centralizing or distributing the out-of-band systems across different headend and hub sites.
4. Proactive network monitoring integration should be planned to gain operational familiarity with new DAA methods of performing sweep, upstream spectrum analysis and reverse ingress monitoring.
5. Disaggregation of the end-to-end solution components, as well as the separation of software from customized hardware results in separate hardware and software roadmaps and multi-dimensional scalability. With disaggregated swim lanes, the end-to-end solution may be multi-partner and multi-dimensional, as well as some swim lanes managed by the operator.
6. Tools, tools, tools: visibility into all aspects of software is a priority at the outset and is essential to effectively debug and develop utilizing Agile-based iteration management.

## 8. Summary

Cable operators have an existing footprint that they continue to grow and improve upon but it can't be overhauled overnight. Technologies such as Remote PHY and virtualization continue to extend the tool set that cable operators can use to stay competitive with capacity growth demands and challenging market environments. This paper has described a transition from I-CMTS in a traditional HFC deployment to a vCMTS in a Remote PHY-based DAA deployment.

DAA addresses historical challenges by eliminating some headend equipment altogether (e.g. splitters, combiners, transmitters, receivers) and by moving RF processing from the headend to the field, further reducing power, space, cooling and cabling requirements at the headend.

Additionally, when vCMTS is coupled with DAA, cable operators immediately benefit from additional savings of space, power, cooling and cabling, while unlocking the path to sustainably growing capacity, adapting quickly to customer demands, and a solution that is flexible and elastic enough to dynamically augment and shift resources to the most in-demand applications.

In summary, vCMTS deployed in a DAA is operationally and financially compelling, addressing today's access network demands as well as looking into the future, with no end in sight to continuous broadband and service demands.

## Abbreviations

CapEx	Capital Expenditure
CCAP	Converged Cable Access Platform
CIN	Converged Interconnect Network
CLI	Command Line Interface
CMTS	Cable Modem Termination System
COTS	Commercial Off-The-Shelf
CPE	Customer Premise Equipment
CPU	Central Processing Unit
DAA	Distributed Access Architecture
DEPI	Downstream External-PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
Gbps	Gigabits Per Second
GCP	Generic Control Plane
HFC	Hybrid Fiber-Coaxial
HW	Hardware
I/O	Input/output
MAC	Media Access Control
NFV	Network Function Virtualization
NIC	Network Interface Controller
NOC	Network Operations Center
OOB	Out-of-band
OpEx	Operating Expenditure
OS	Operating System
PHY	Physical
PNM	Proactive Network Maintenance
RF	Radio Frequency
RPD	Remote PHY Device
RPN	Remote PHY Node
RU	Rack Unit
SCTE	Society of Cable Telecommunications Engineers
SDN	Software Defined Networking
SW	Software
TCO	Total Cost of Ownership
TTM	Time to Market
UEPI	Upstream External-PHY Interface
vCMTS	Virtual CMTS
vCPE	Virtual CPE
VOD	Video on Demand

## Bibliography & References

[1] DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-118-, April 22, 2019, Cable Television Laboratories, Inc.

[2] Matatyaou, Asaf. Practical Lessons of a Centralized Virtualized CMTS. San Jose: Harmonic, 2018. Web.

[3] Matatyaou, Asaf. Real-World Deployment of a Virtual Cable Hub. Web. San Jose: Harmonic, 2017. Web.

[4] Matatyaou, Asaf. Transforming the HFC Access Network with a Software-Based CCAP. San Jose: Harmonic, 2015. Web.

[5] Modular Headend Architecture v2 Technical Report, CM-TR-MHAv2-V01-150615, June 15, 2015, Cable Television Laboratories, Inc.

[6] Remote DOCSIS Timing Interface, CM-SP-R-DTI-I07-180509, May 9, 2018, Cable Television Laboratories, Inc.

[7] Remote Out-of-Band Specification, CM-SP-R-OOB-I09-180509, May 9, 2018, Cable Television Laboratories, Inc.

[8] Remote PHY Specification, CM-SP-R-PHY-I12-190307, March 7, 2019, Cable Television Laboratories, Inc.



# **The Generic Access Platform**

## **What's in it for me?**

A Technical Paper prepared for SCTE•ISBE by

Roger G Stafford  
Principal Architect III

Charter Communications, Inc.  
CTEC I, 14810 Grasslands Drive, Englewood, CO 80112.  
Tel: +1 704 681 2799  
[roger.stafford@charter.com](mailto:roger.stafford@charter.com)

Joint Co-chair SCTE Generic Access Platform Working Group

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
1. What is the Generic Access Platform?.....	4
Content .....	6
2. Why a change in design philosophy becomes the new advantage .....	6
2.1. Distributed Access Architecture (DAA) - Enabling components deeper into the network. ....	6
2.2. Why will node housing become more common?.....	6
2.2.1. HFC Bandwidth increases.....	7
2.2.2. Deeper into the Network with Optical Nodes.....	7
2.2.3. Emerging Cellular Markets .....	7
2.2.4. Smart-City Applications.....	8
2.3. An opportunity to Unify DAA Components.....	10
2.4. Improvements in end-of-line Signal Quality .....	11
3. GAP is a more modular design approach .....	11
3.1. Modular Components with upgradability.....	13
3.2. A Traditional Node Lifecycle .....	13
3.3. The GAP Lifecycle.....	15
4. Examining the Economic Differences .....	17
4.1. The Traditional Node .....	17
4.1.1. From a node vendor perspective; .....	17
4.1.2. From a service provider perspective;.....	17
4.2. The GAP Approach .....	18
4.2.1. From a node vendor perspective; .....	18
4.2.2. From a service provider perspective;.....	18
4.3. Comparing GAP versus Traditional Node Costs.....	19
4.3.1. Assumptions used in the model.....	20
4.3.2. Comparing the cost factors for the design and production (R&D).....	21
4.3.3. Comparing the deployment and network sustaining cost factors for the service provider .....	22
Conclusion .....	25
Abbreviations.....	26
Bibliography & References .....	27
References.....	27

## List of Figures

Title	Page Number
Figure 1 - An Increasing number of application devices over time, which could be housed in a GAP enclosure. ....	5
Figure 2 - Bandwidth Compound Annual Growth Rate increases the number of Nodes .....	7
Figure 3 - Expansion of Cellular Networks .....	8
Figure 4 - Categories for Connected devices in emerging markets .....	9

Figure 5. Internet of Things (IoT) connected devices installed base worldwide from 2015 to 2025 .....	9
Figure 6 - An example of current strand-mount DAA applications housings .....	10
Figure 7 - A modular approach reduces the number of styles. ....	12
Figure 8 - A modular approach to node design.....	13
Figure 9 - Lifecycle for a typical node.....	14
Figure 10 - Lifecycle for a GAP Node or any GAP enclosure.....	16
Figure 11 - Cost factors for nodes.....	20
Figure 12 - Traditional node: vendor design, development and production costs by factor .....	21
Figure 13 - GAP node: vendor design, development and production costs by factor .....	22
Figure 14 - Traditional node: service provider deployment and sustaining costs .....	23
Figure 15 - GAP node: service provider deployment and sustaining costs .....	23
Figure 16 - Future traditional node deployed volume and cost over time.....	24
Figure 17 - Potential upgrade paths for GAP nodes .....	25

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Potential uses for a GAP housing enclosure.....	11

# Introduction

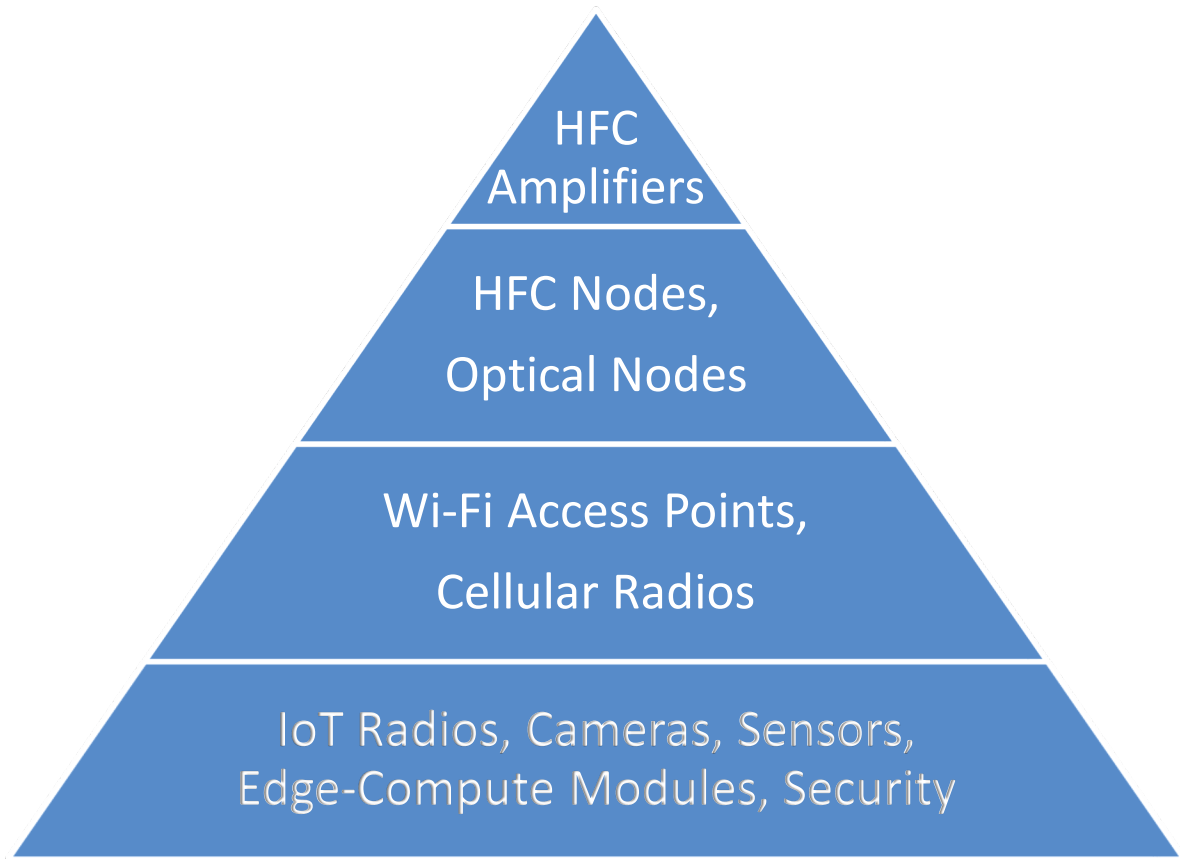
## 1. What is the Generic Access Platform?

Simply stated, the Generic Access Platform is an outdoor housing enclosure that can be built in a number of ways to address multiple applications.

When a Distributed Access Architecture (DAA) was being discussed as a new approach to the way Cable and Telecommunications operators could redistribute the new components of either purely optical or Hybrid Fiber-Coaxial (HFC) networks it became apparent that DAA would lead to a much greater number of network access devices being placed in the outside plant portion of the access network. Since most outdoor components are strand mounted or pedestal mounted within the United States and cabinet mounted in many other parts of the world, this in turn leads to a greater diversity of equipment vendors producing and deploying a variety of different types of equipment. Today, MSOs deploy strand-mounted amplifiers, numbering tens of millions, but these are limited to just a few vendor designs. DAA introduces the concept of remote physical layer nodes to convert deeper penetrated digital fiber into HFC radio-frequency spectrum, closer to customer's homes than previous network design and deployment approaches.

Additionally, the GAP housing can be used for other applications beyond just an RF node. It is also intended to standardize the housing design for other outdoor equipment; 4G & 5G Small-Cells radios, Wi-Fi Access Points, remote OLTs and ONUs to support EPON and GPON networks, Edge-Compute Nodes and other smart-city applications such as IoT radios, traffic-light and pedestrian monitoring, and smart-sensing, as shown in Figure 1. In fact, the GAP housing can be used to accommodate multiple functions in the same housing (considering some thermal and power constraints) such as being an RPD with an IoT radio included for example. This greatly reduces the need for multiple node housing on the same coaxial strand.

There has been a progression from simple HFC amplifiers to a myriad of network-edge applications, combined with an increase in the available bandwidth across a network. Many techniques for increasing bandwidth, such as node-splits, also requires additional components that need to be housed outdoor.



**Figure 1 - An Increasing number of application devices over time, which could be housed in a GAP enclosure.**

However, operational complexity arises from having so many applications. In a conventional approach the solution for each application will be provided by multiple vendors, each with a custom design, and each inside a custom housing enclosure. Technical momentum will drive a high rate of replacement for these new technically advanced “nodes”. The innovation cycle will become very challenging for cable and telecommunications service providers because the rate of deployment, and extraction to and from the field will rapidly increase over time. Each custom housing cannot be re-used at the end of its useful life, where the end of the useful-life could be a consequence of technology upgrades or component failures in the field.

The Generic Access Platform is designed to address the life-cycle challenges by producing a single housing that can be re-used for the same initial purpose, be upgraded to a new technology for the same functional purpose, or be completely repurposed as new device for a new technical function.

# Content

## 2. Why a change in design philosophy becomes the new advantage

### 2.1. Distributed Access Architecture (DAA) - Enabling components deeper into the network.

DAA has a number of advantages:

- Network efficiency
- Increased network capacity and simpler outside plant maintenance
- Supports Node evolution with Remote PHY, Remote MAC-PHY and Remote 10G EPON or GPON OLT
- Better end-of-line signal quality, higher modulation rates, higher bit-rates
- Better spectral efficiency, more wavelengths per fiber
- Operational and capital expenditure benefits
- Reduced headend power, space and cooling requirements
- Hub consolidation
- Add QAMs without changing the RF combining network
- IP convergence
- Extend IP network to the node
- Alignment with FTTx build-out
- Ability to leverage standards-based interconnectivity and economies of scale
- [1] [2]

The GAP housing does not aim to replace any DAA technologies. In fact, it encourages the development and deployment of DAA technologies in a standardized way. The GAP housing enables a standardized approach to the hardware enclosure. It can;


- Reduce the number of custom designed and manufactured housings.
- Address the market needs ahead of a large growth in outdoor equipment predicted by DAA and Smart-City applications.
- Reduce operational expenditure for Cable and Fiber Main Service Operators, and Telecommunications companies.
- Increase longevity for deployed housing due to re-purposing rather than housing replacement.
- Increase availability and ability to integrate advanced technologies within a modular approach.
- Facilitate inter-operability between different vendor technologies.
- Introduce a common industry-wide approach to outdoor deployed devices.
- Increase access to market-share for new technology providers.
- Reduce the need for multiple node housing on the same coaxial strand.

### 2.2. Why will node housing become more common?

The following section describes how the number of housing increases over time due to a number of factors.

### 2.2.1. HFC Bandwidth increases

The compound annual growth rate (CAGR) of user-demand for bandwidth is approximately 43-45%, which is driving service operators to divide nodes, as shown in Figure 2.



**Figure 2 - Bandwidth Compound Annual Growth Rate increases the number of Nodes**

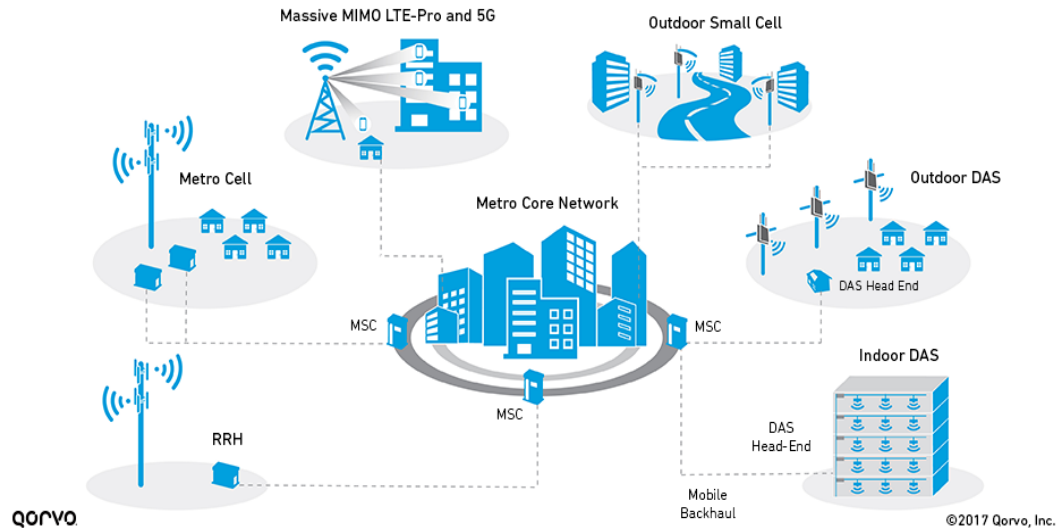
### 2.2.2. Deeper into the Network with Optical Nodes

The deployment of optical nodes, especially in the trend towards N+0 network architecture, will also increase the number of deep-fiber optical links and hence fiber nodes.

### 2.2.3. Emerging Cellular Markets

As the cellular market expands into the Citizens Band Radio Service (CBRS) band 48 at 3.5GHz, a greater number of new cellular nodes are needed to bring those services closer to the customer. Developments are on-going within the mobility market to design and deploy Radio Access Network (RAN) Small, Pico and Femto-Cell radio nodes into this emerging infrastructure, as shown in Figure 3.

## Wireless Infrastructure: A Heterogeneous Network




**Figure 3 - Expansion of Cellular Networks**

[3]

### 2.2.4. Smart-City Applications


Additional housings will be needed for Smart-City applications to enable sensing and monitoring devices such as IoT radios and surveillance cameras. These monitoring and sensing applications are still emerging but there is no doubt that they will continue to increase in number and will drive the need for more network node connection points. Higher traffic capacities and lower latencies are normally achieved by placing nodes closer to the user groups or customers. Data published by Priceonomics, Figure 4 and Statista, Figure 5 shows the emerging market shares for different IoT categories, by market segment, and predicted growth for IoT devices which could be used as a proxy for the relative growth in new technology node devices.





[4]

**Figure 4 - Categories for Connected devices in emerging markets**



[5]

**Figure 5. Internet of Things (IoT) connected devices installed base worldwide from 2015 to 2025**

The GAP housing aims to limit the number of housing styles by making available a modular, upgradable and re-usable housing that can accommodate one or more node functions at any time. The approach will provide a uniform outdoor enclosure platform for a variety of different types of devices.

### 2.3. An opportunity to Unify DAA Components

The predicted growth in strand-mounted technologies will put a strain on overhead plant infrastructure. Most service operators in the United States are confined to using strand (see Figure 6) or pole mount equipment due to problems with getting permits to use city-owned street-level mounting positions.

The GAP housing will help to alleviate the amount of strand-mounted infrastructure by combining two or more box solutions, or collections of associated strand-mount devices into a single enclosure. An example would be an outdoor Wi-Fi Access Point backhauled by a strand-mount DOCSIS cable modem. Both items could be co-located inside the same housing if designed in a modular form factor. Other application technology use-cases are shown in Table 1.



[6]

**Figure 6 - An example of current strand-mount DAA applications housings**

A modular form-factor allows for an even higher potential for integration. Multiple new functions could be achieved in the same node. For example, a traffic monitoring camera or CBRS small-cell could be incorporated along with a Wi-Fi AP with DOCSIS modem backhaul within the same housing. This benefits strand-loading from a reduced weight perspective and lower power consumption, compared to using two separate node enclosures.

**Table 1 - Potential uses for a GAP housing enclosure**

DOCSIS	R-PHY	R-MAC-PHY	vRouter	LoRA WAN	Wi-Fi
DAA	Traffic Monitoring	Edge-Compute	Security Cameras	Environmental Monitoring	R-ONU
PON	DWDM	Coherent Fiber	R-OLT	CBRS	FWA
DPI	Surveillance	Smart-Network Diagnostics	Earthquake Detection	Edge-Caching	Flood Detection


## **2.4. Improvements in end-of-line Signal Quality**

There are a number of advantages to the DAA architectural approach. One of the advantages is to use digital optics to deliver digital signals that are converted by the remote PHY device (RPD) to analog radio frequency signals such as OFDM and QAM used by coaxial cable systems, as well as those used by 4G, 5G, Wi-Fi and IoT radio applications. In the cable case, the elimination of noise and distortions produced by conventional analog intensity modulated optical links from the hub to the optical node will result in a higher MER at the end-of-line and consequently higher QAM modulation orders can be used such as 4K-QAM, 8K-QAM, and potentially 16K-QAM.


## **3. GAP is a more modular design approach**

The GAP enclosure present a very different lifecycle compared to the traditional node design. Technologies become housed in custom designed modules using a common modular form factor, as shown in Figure 7.

*Multiple Housings with Incompatible internal components*




*The GAP Housing*



**Figure 7 - A modular approach reduces the number of styles.**

Modules are mechanically supported by the housing which also provides thermal dissipation for up to 220W of input power. Power is distributed to each module through a backplane that modules plug into. The proposed design also offers a high-speed (PCIe) and a low-speed (I2C) bus for module-to-module data transfers and module management functions.



**Figure 8 - A modular approach to node design**

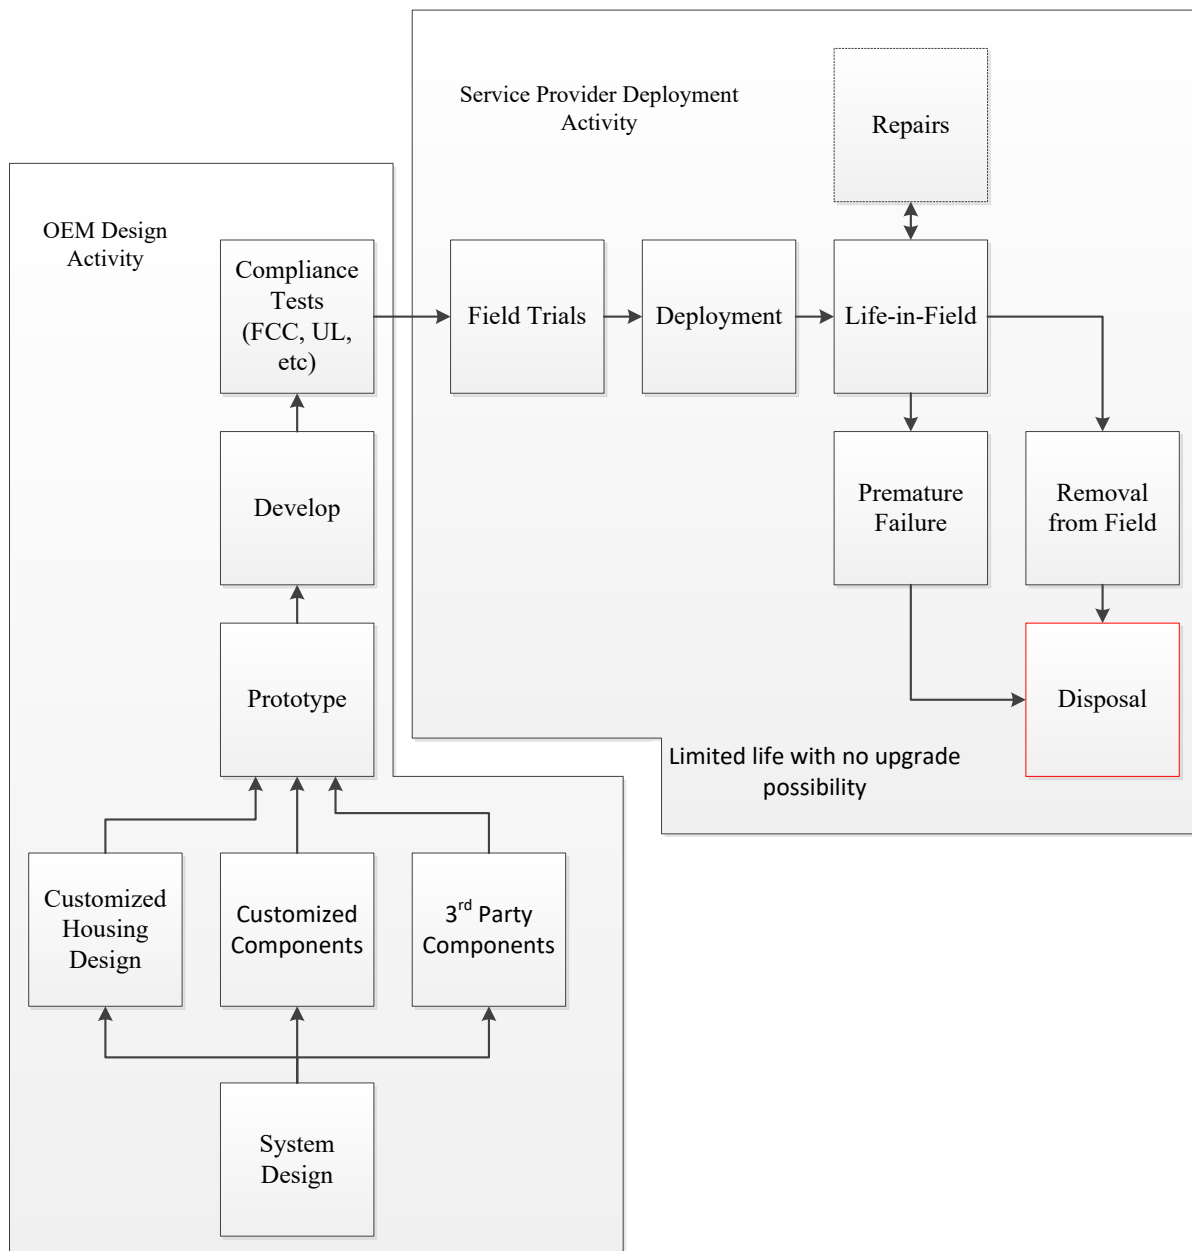
The external ports are basic threaded entry holes that can support either RF entry ports up to 3GHz capability, optical fiber entry, or shielded-Ethernet cable. A fourth option is for any RF port to be used to connect an array of antenna, mounted on the external surfaces of the housing via mounting spigots. This is to support outdoor Wi-Fi, CBRS and IoT applications requiring MIMO arrays. The internal RF interconnections remain customized as they are very application dependent so as to allow flexibility for module to module connections that go beyond the basic backplane requirements.

### 3.1. Modular Components with upgradability

- Housing: A clamshell design that can be re-purposed. Mechanical design features such as replaceable RF gaskets, silicone weather seals and removable entry/exit ports.
- A modular High-Speed Data backplane: A PCIe 5.0 based interconnect bus. As technology advances this backplane can be single part upgraded to PCIe 6.0 and so on.
- Module Slots: These remain the same width are interchangeable between the lid and the base, for flexible system design options.
- Power Supplies: Standardized voltage rails. Located in a position that optimizes cooling and thermal dissipation for heatsinking. Can be customized for some applications.
- Power-plane: A modular, replaceable, PCB that interfaces power to the modules though fixed and optional voltage rails.

### 3.2. A Traditional Node Lifecycle

A traditional node design is based on a single-use design philosophy that results in a single deployment and eventual disposal after a relatively short service life, as shown in Figure 9.



**Figure 9 - Lifecycle for a typical node.**

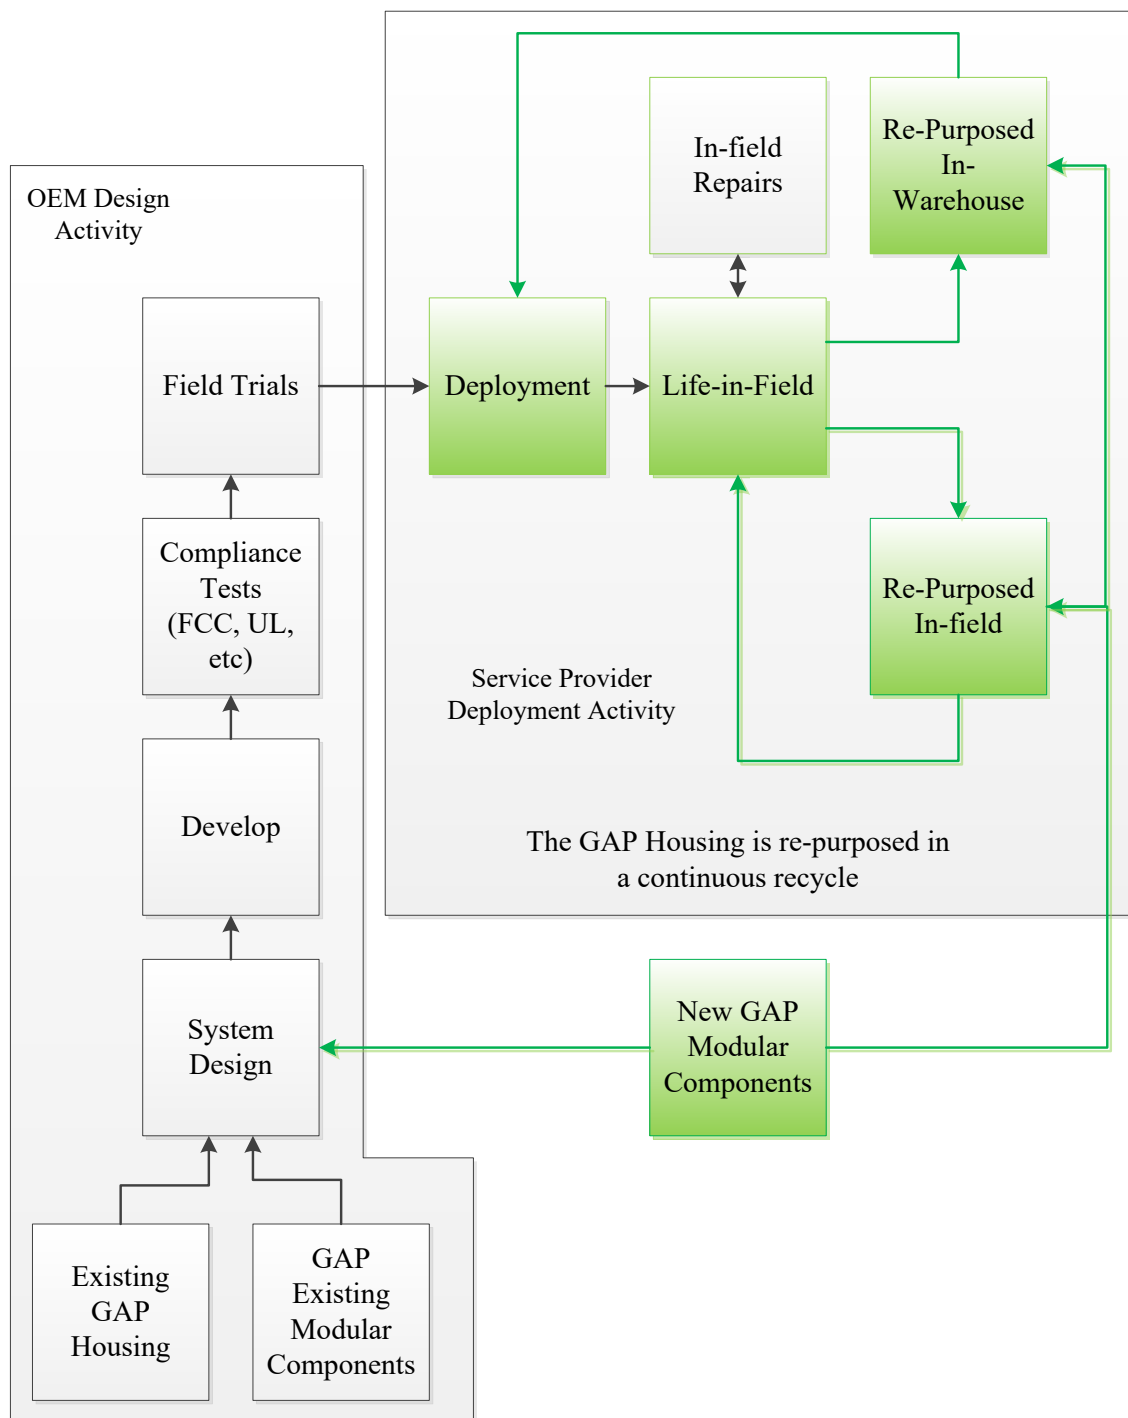
At present, DAA devices such as nodes begin as a custom design, using specially developed components such as an optical transceiver, embedded CMTS, power-supply, and various RF modules and connections. Once developed and tested for applications and regulatory compliance, the final design moves onto a field trial phase. After successful trials and any subsequent improvements, the device is deployed to the field, where it may go through some infrequent repairs, but ultimately gets removed from the field and has no further use, leading to disposal of the hardware, at a capital loss to the service operator. Only very recently has the idea of an upgradable node housing become more popular, with some node vendors already starting to propose re-usable housings.

An analogy for the GAP housing is the personal computer. The first personal computers of the 1970's and 1980's were customized designs that integrated all of the necessary components such as the CPU board with on-board memory, a visual display unit and a keyboard. It became apparent that the rate of development was being slowed by the sequential design work needed integrate all of the new, technically-advancing components. The solution derived in 1985 was to produce the Advanced Technology (AT) form-factor motherboard with common interfaces that could be housed in a standard chassis. The AT motherboard could be equipped with plug-in sub-components such as LAN/WAN cards, memory modules, disk interface cards and other peripheral devices. [7].

Development then centered on the incremental improvement to each of the peripheral modules. Over the following years, the motherboard capabilities were increased but there remained a backwards compatibility in terms of supported interface standards. The development process moved from a single vendor being the only developer, to a multi-developer environment with a higher rate of new product availability. The GAP project is essentially leveraging a similar approach except that cadence is being exchanged for the ability to develop technical functionality while retaining as much existing hardware as possible, thereby lowering the operational expenditure for service providers. By 1995, the AT form-factor was revised into the Advanced Technology eXtended (ATX) motherboard which continued the upgrade path for the next generation PC industry. [8]

### **3.3. The GAP Lifecycle**

A GAP-based node design is based on a multiple-use design philosophy where the housing and modules are re-used and therefore will have a much longer service life, as shown in Figure 10.



**Figure 10 - Lifecycle for a GAP Node or any GAP enclosure.**

One of the aims for the GAP housing is to make use of a modular approach, to design and deploy nodes and other DAA equipment. This approach has the following benefits;

- Reduced system design time because the whole enclosure is no longer a custom design.



- The enclosure can be re-used by repaired more easily by replacing failed or damaged modules.
- The enclosure can be re-used by refitting it with new modular components.
- The technical function of a GAP housing can be incrementally upgraded or completely repurposed as a new technical function.
- It is an easier technology upgrade if only certain modules or technology functions need to be upgraded.
- Repurposing, repairs and re-fitments can be done without having to remove the housing from the strand, thereby making it easier to maintain at a lower overall operational cost.
- Modules can be re-warehoused and re-used in other application housings.
- It is a better model for sustaining new technology deployments into the field, especially where both the volume and specialty of those technologies is rapidly increasing, which offers some future-proofing between network access upgrades.

Longevity in the field can only be achieved by a GAP enclosure if it offers a high degree of flexibility for the adoption of any particular set of new technologies. It must be designed to be upgradable if it is to adequately fulfil this role into the future. Today, it must be able to house DOCSIS3.0/3.1 with a 1.2GHz spectrum for example, but also accommodate DOCSIS4.0, Extended Spectrum DOCSIS and Full-Duplex technologies, up to 3GHz capability.

## 4. Examining the Economic Differences

### 4.1. The Traditional Node

#### 4.1.1. *From a node vendor perspective;*

Each generation of node does bring new technologies and a more advanced set of features, and this is an essential part of the reason for deploying new nodes. Traditional nodes require a large amount of engineering expertise to design and construct. Currently, nodes are developed in a flow similar to those shown in Figure 9. A group of system designers are needed to define all of the individual components needed for any node or application housing. These components are usually housed in individual modules to isolate certain functions from a power or RF perspective, or simply to divide up the design work among multiple design engineers or design and production companies. Mechanical engineers produce the overall structure, while individual teams customize and optimize the engineering to meet performance and cost objectives. Each module is effectively designed from scratch, and each set of modules that make up any node are discarded with each generation of housing. Similarly the housing style is also discarded and recreated between generations. There is very little re-use of the previous design, especially the mechanical components.

These design cycles increase the amount of design effort needed to produce each generation.

#### 4.1.2. *From a service provider perspective;*

Each new node technology is anticipated to need a large operational expenditure, in addition to a capital expenditure, in order to do a network access plant upgrade.

The initial phase focuses on specification work conducted between silicon vendors, node-providers and service providers. The next phase consists of a lengthy period of prototype work (1-2 years) and a period of design verification, software feature upgrades and field trials (1-3 years), that results in a deployable product. The overall development and deployment cycles are very lengthy and cumbersome.

The preferred method for network access upgrades is to do them incrementally. Usually upgrades are performed service group by service group rather than attempting a large-scale or nationwide upgrade. This is because the service-group approach requires a lesser initial capital expenditure, uses less service technicians at any given point in time, and spreads the capital expenditure over a longer period of time. Another preferred approach is to take a longer term view of each service group upgrade by installing equipment that will enable that site not to have to be revisited within a 5-8 year period after an upgrade takes place. This means the future phase must have already been contemplated, planned and resourced ahead of any upgrade. With the traditional node architecture the amount of prior planning is limited to what the current generation of technology can offer. The traditional node is deployed and the network remains constant for a long fixed duration, typically greater than 8 years, until the next generation becomes available. At that point the traditional node is disposed of and entirely replaced and disposed of.

A GAP housing can be re-used many times by replacing either damaged or failed modules, subject to normal AFR rates, or be upgraded using new function modules. In either case, the housing will last considerably longer. There will be some housings that will need to be replaced due to physical damage such as from natural disasters, lightning or accidental connector damage, for example, but the quantity will be very low compared to the total deployed population. For this reason, a low in-field disposal rate of 0.1% was used in the cost-model.

## **4.2. The GAP Approach**

### ***4.2.1. From a node vendor perspective;***

The GAP housing constitutes a different design philosophy compared to a traditional node design. The intention is to re-use as much as possible from a previous GAP node or device housing. A system design is still needed for the new components as these will most likely be the new technology and will involve some design effort to produce them in a GAP modular form-factor. The rest of the system will likely re-use existing power supply modules, backplane and the enclosure.

This new approach means the traditional node vendors migrate from being a complete system design and production entity, and instead become a producer of modules. They retain the module design and new technology development aspects however. They preserve their intellectual property in their traditional core areas. The modular approach allows those vendors to use modules from other vendors, such as the PSU or RF interface modules.

Additionally, it opens the node market to new OEMs, to build sub-sections of a complete node such as an iCMTS as a new module. Further, the GAP approach allows a market entry option to new companies that want to take on either the system design, build, test, compliance, complete node construction or some combination of these roles.

### ***4.2.2. From a service provider perspective;***

The service provider has the potential to take on the system design role which may be advantageous to the larger service providers, but there is also the option for a traditional node vendor to take on the system design role on behalf of a service operator. The service operator also becomes the system integrator and is ultimately responsible for the technical and compliance testing, although there is also the option for any 3<sup>rd</sup> party to perform these roles as an additional service.

Why would a service operator take on the role of being the system designer? The answer to how much or which areas of involvement depends on the complexity of the node design. For a new function or design, the traditional node vendor or OEM produces the new module specification using the GAP enclosure and

module specification as a basis. The OEM produces an ‘encapsulated’ module that a service provider can then integrate with other existing modules currently in inventory. For example, a service operator might already have stock of GAP housings, suitable GAP module PSUs and RF interface modules, and these can be put together by a service operator system designer who adds a new iCMTS module recently introduced by an OEM. In this example, the OEM would have designed the iCMTS taking into account already available system components, i.e. existing modules vital to the final design. The service operator might also take on the compliance and testing aspect of the final node design.

For a mature design, the components are already well understood and relatively easy to integrate to produce a variation on an existing GAP node design. Some system level design is needed but a new node with extended functions can be produced without going back to the very beginning of the design process.

### **4.3. Comparing GAP versus Traditional Node Costs**

The analysis was done by breaking down the factors that go into both the OEM’s or vendor’s design and production costs, and the service operators deployment and network infrastructure sustaining costs. Each are separately broken down in Figure 11.

### Vendor Design, Development & Production Costs

---

Customized Housing Design  
Customized Components Design  
3rd Party Component Design  
Prototyping  
Compliance Testing

#### Total R&D Cost

---

Material Costs  
Manufacturing Cost

#### Total Unit Manufactured Cost

---

#### Unit Sales Price

---

### Service Provider Capital and Operational Costs

---

#### Set-Up Costs (\$)

---

System Design - Service Operator  
Compliance Testing - Service Operator  
Field Trials set-up  
Training

#### Total Set-up Costs

---

#### Deployment Costs (\$)

---

Unit Deployment CapEx – Average node HW including fittings  
Unit Deployment OpEx Cost each node including hardline re-connection

#### Total Deployment Costs

---

#### Network Sustaining Costs

---

Annual Repairs at % AFR  
Annual Removal from Field  
Disposal at \$ per unit  
Capital Write-off due to failures  
Node Upgrade Cost

#### Total Network Sustaining Costs

---

#### Total Service Provider Cost

---

**Figure 11 - Cost factors for nodes**

#### ***4.3.1. Assumptions used in the model***

Two economic models were created; one for a traditional node and a second for a GAP node. The following assumptions were used in each:

- A traditional node would need to be designed from scratch.
- A GAP node would re-use an existing housing design, and all other components except one major module would change with each node design iteration.
- The deployed-node failure rates would be the same for each, at a 2% AFR.
- A failed node of any type would need to have all hardline connections replaced during the exchange.
- A traditional node in-field exchange would need each hardline connections to be re-terminated, whereas a GAP node would not.
- Node warehousing and site permitting costs were excluded because they are broadly the same for both housing types.
- The node upgrade cost assumes the majority case of one module (see Figure 17 for other assessments).

#### 4.3.2. Comparing the cost factors for the design and production (R&D)

An analysis of the node vendor or OEM's cost structure can be seen in Figures 12 and 13 below. Note that the vertical scale are the same in both figures, giving a view of relative costs.




**Figure 12 - Traditional node: vendor design, development and production costs by factor**

The customized component design represents the largest cost, followed by the development and prototyping cost. A significant cost factor is the housing which in the case of a traditional node, is customized for each application, and is not re-usable across different applications. Compliance testing is costly because it involves development and testing all of individual components that make up the node.

In the case of a GAP-based node the same factors apply but in different degrees of cost. There isn't customized housing design because a standard GAP housing design is being utilized. Third-party component design is also not a factor. Compliance testing is now done by the service provider, or its system design agent, so no longer figures as a cost item. There is still the need to develop the main

application module where the mechanical aspects are also standardized using the GAP form-factor. A cost still exists for prototyping, which is less given the reduced scope of electrical and mechanical changes.



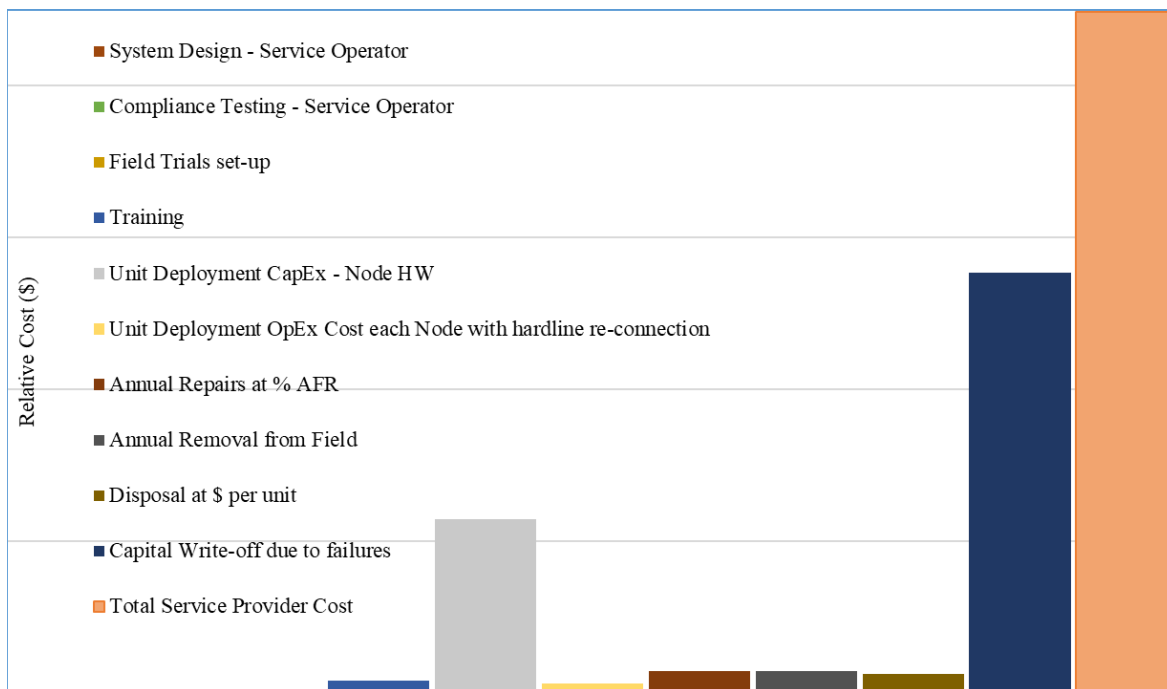
**Figure 13 - GAP node: vendor design, development and production costs by factor**

Overall, there is a large reduction design and development costs because of the large amount of component re-use when modules are available in a standardized form-factor. The use of a modularized approach does incur additional component costs however. A completed GAP node has a higher unit cost compared to a traditional housing. However, the impact of this increased costs be weighed against the lower overall cost for a service provider.

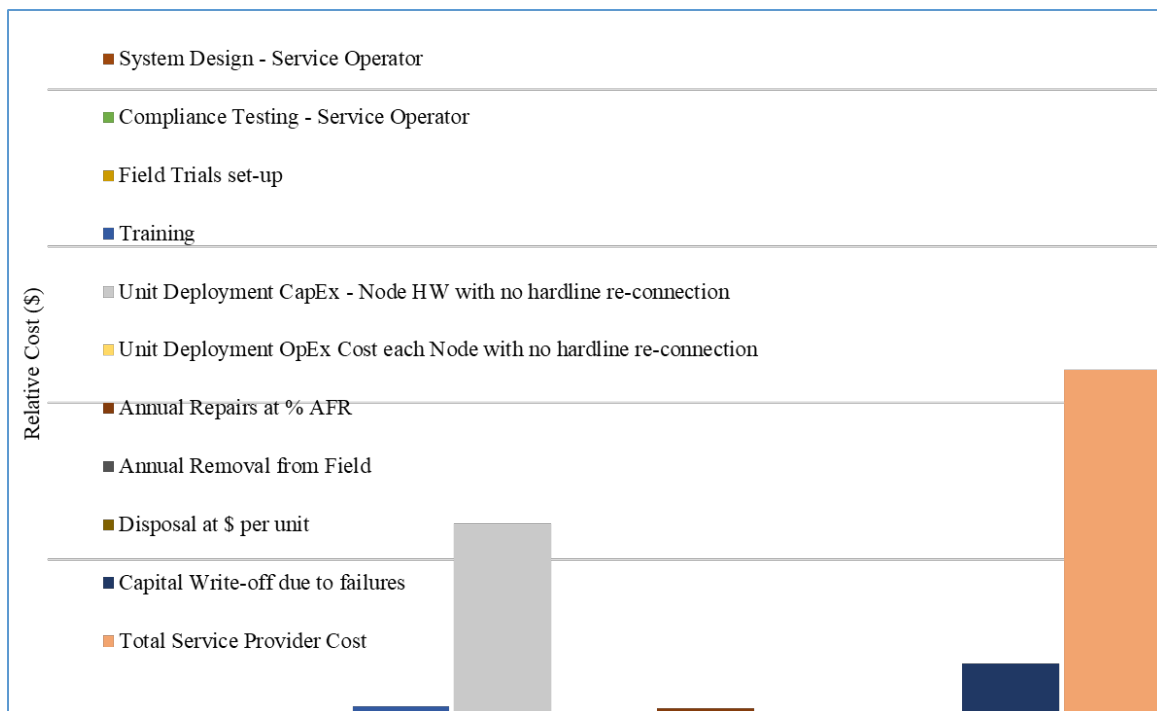
#### ***4.3.3. Comparing the deployment and network sustaining cost factors for the service provider***

In this analysis a cost model created that compares the deployment and aspects for sustaining a network evolution for both a traditional and a GAP housing approach. Again, the same cost factors were assessed to compare their relative values.

In the case of the traditional node deployment, some of the costs such as system design and compliance testing are zero because these were part of the cost associated with the R&D phase above. Significant costs are experienced for training, deployment and repairs. Repair of a traditional node usually involve removing the node from a strand and replacing it with a new node. This means all of the hardline connections need to be replaced and re-terminated with a new connector – which adds significant time and cost. By far the greatest cost is the capital cost write-off associated with replacing a node either because it has reached the end of its functionally useful life as a technology, or due to annual field failures.



**Figure 14 - Traditional node: service provider deployment and sustaining costs**




**Figure 15 - GAP node: service provider deployment and sustaining costs**

Both Figures 14 and 15 are drawn on the same relative cost scale for cost comparison purposes. The costs for deploying and sustaining GAP-based nodes are significantly lower as can be seen from the far-right

hand bar; total service provider costs which is a cumulative total of the preceding bars to the left. The central bar indicating unit deployment capital expenditure is higher, by about 5% for GAP-based node compared to a traditional node, however the costs associated with deployment, repair and in particular, capital write-off costs are significantly lower. The cost-model assumes the same rate in deployed volume increase for both traditional and GAP nodes and the same rate of repair based on a 2% AFR. For GAP-based node deployment, the cost-model included an additional factor not include in the traditional node model. That is, 2% of nodes are being upgraded each year with new features by replacing old modules for new modules.


Figure 16 shows a hypothetical scenario where there is a cut-over to GAP-based housing node deployments in year 4, coincident with an increase in demand for application nodes beyond the traditional fiber/HFC type of node.



**Figure 16 - Future traditional node deployed volume and cost over time**

The cost analysis assumes that one module change only is needed to upgrade any GAP node to a new function, so it represents one of the lowest cost scenarios. In reality, there will be a blend of different options ranging from a single component or module upgrade through to a complete change of all of a GAP housing's contents. In some situations, it will be advantageous to remove and replace an entire lid or base that has been pre-built with the new modules in situ, thereby reducing the technicians in-field upgrade time. Figure 17 shows a typical blend and predicted percentage of these module exchanges. A further extension for this paper would be to examine the costs associated with re-building a lid or base with new modular components to create a new node function, taking into account the construction location costs, labor, warehousing and distribution prior to a technician doing an in-field GAP node exchange.





**Figure 17 - Potential upgrade paths for GAP nodes**

## Conclusion

The Generic Access Platform (GAP) housing will change the way nodes are designed, developed and deployed into cable, optical and telecommunications networks. Network operators are on the cusp of a rapid increase in applications that are best suited for closer-to-the-edge node enclosures. GAP provides a means to deploy multiple types of nodes which are being created to address these new applications such as Wi-Fi, cellular, mobility, edge-compute and edge-storage.

The GAP housing offers the ability to use modularized and standardized components that will ultimately lead to operational cost reductions for service operators, while offering the ability to be upgradable to suit new technologies as they emerge. GAP nodes will have a longer in-field life and offer much reduced capital write-off when nodes need to be upgraded. Furthermore, the modular approach reduces the development time for new node features and provides a much faster time-to-market for node component vendors.

The cost-model shown in this paper reveals the design and development costs are also reduced for system designers and the vendor community. It enables an existing vendor to give greater concentration to their core competencies. The GAP modular approach also serves to allow the entrance of new node component vendors, as either node builders or modularized technology providers, such as for CBRS radios or edge-compute services.

# Abbreviations

4G	fourth generation cellular network
5G	fifth generation cellular network
AFR	annual field failure rate
AP	access point
AT	advanced technology (form-factor)
ATX	Advanced technology eXtended (form-factor)
bps	bits per second
CAGR	compound annual growth rate
CBRS	citizens band radio
CMTS	cable modem termination system
DAA	distributed access architecture
DOCSIS	data over cable service interface specification
DWDM	dense wavelength division multiplexing
EPON	version e of a passive optical network
ESD	extended spectrum docsis
FTTx	fiber to the x, where x is curb, premise, or home
FWA	fixed wireless access
GAP	generic access platform
GPON	version g of a passive optical network
HFC	hybrid fiber-coax
HW	hardware
Hz	hertz
iCMTS	Integrated CMTS
IoT	internet of things
ISBE	International Society of Broadband Experts
LoRa	long range (radio wan)
LTE	long term evolution cellular network
MAC	media access control
MER	modulation error ratio
OLT	optical line termination
OEM	other electronic manufacturer
ONU	optical network unit
PC	Personal Computer
PHY	physical layer
RAN	radio access network
RMD	remote mac device
R-OLT	remote optical line termination
R&D	research and development
RPD	remote phy device
SCTE	Society of Cable Telecommunications Engineers
WAN	wide area network
Wi-Fi	a family of radio technologies commonly used for wireless local area networking

# Bibliography & References

## References

- [1] ARRIS International Limited, "Evolving Networks with Distributed Access Architecture," August 2019. [Online]. Available: <https://www.arris.com/solutions/distributed-access-architecture/>.
- [2] AGC Research, "Modernizing the Cable Service Delivery," August, 2019.
- [3] Tuan Nguyen, Quorvo, "Small Cell Networks and the Evolution of 5G (Part 1)," 17 May 2017. [Online]. Available: <https://www.qorvo.com/design-hub/blog/small-cell-networks-and-the-evolution-of-5g..>
- [4] Priceonomics Data Studio, "The IoT Data Explosion: How Big Is the IoT Data Market?," 9 January 2019. [Online].
- [5] Statista, "Internet of Things (IoT) connected devices installed base worldwide from 2015 to 2025 (in billions)," 10 12 2017. [Online]. Available: <https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/>.
- [6] E. Dylag, "The Universal Node," in *SCTE Expo*, 2019.
- [7] "IBM AT form factor," 1 August 2019. [Online]. Available: [https://en.wikipedia.org/wiki/AT\\_\(form\\_factor\)](https://en.wikipedia.org/wiki/AT_(form_factor)).
- [8] "ATX form factor," 1 August 2019. [Online]. Available: <https://en.wikipedia.org/wiki/ATX>.

# **New Sensing Techniques For Advanced IoT Applications**

A Technical Paper prepared for SCTE•ISBE by

**Arun Ravisankar**  
Senior Engineer, Comcast Labs  
Comcast Corporation  
1800 Arch St, Philadelphia, Pa 19103  
Phone: 2152867558  
Arun\_Ravisankar@comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Sensors in IoT .....	4
New Sensing Techniques .....	7
1. RF Sensing for IoT Applications .....	7
2. WiFi as RF sensor?.....	10
3. Audio-based Sensing .....	12
4. Predictive Analysis using Advanced Machine Learning.....	14
Conclusion .....	15
Abbreviations.....	16

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Components of a Typical IoT Application .....	4
Figure 2 - Sensors used in a home security application.....	5
Figure 3 - Sensors in a Home Security Application.....	5
Figure 4 - Example Sequence Diagram of an IoT application .....	6
Figure 5 - Doppler Signatures for different motions .....	8
Figure 6 - Range Data for Walking and Falling using RADAR.....	9
Figure 7 - Home Security/Automation using RF Sensing .....	9
Figure 8 - Example MIMO system and WiFi RADAR Hardware and Software.....	10
Figure 9 - Components of a Wi-Fi RADAR system .....	11
Figure 10 - Doppler Signatures of Wi-Fi and traditional RADAR system .....	11
Figure 11 - Confusion Matrix for various sounds .....	12
Figure 12 - Example Spectrogram of an Audio Signal .....	13
Figure 13 - Edge Compute System .....	14

# Introduction

History is witness to the evolution of civilizations and how humans continue to discover and innovate things that would propel everyone to newer levels of technological advances, as we aspire to attain a higher intellectual state. Industrial revolutions are key indicators of how humankind continues to seek techniques that would improve lifestyles and bring advancement to civilization. IoT (Internet of Things) applications play a significant role in providing peace of mind to customers/users and help improve lifestyles. Residential IoT applications bring peace of mind to customers by offering security applications or improving lifestyle by offering a suite of home automation applications that enable users with a worry-free experience that also optimize usage of resources. Commercial and Industrial IoT applications help organizations increase productivity and help optimize resources.

Regardless of IoT application type, there are some major components that power them and help provide all the services that these applications are designed and intended for. Figure 1 shows the components that typically make up an IoT application. All these components play an integral role in providing the desired result.


Rules would define how applications would act on data sent from sensors. Rules form the basic construct of IoT applications. For example, rules set on a home automation application to determine the next course of action if a motion sensor detects a motion event.

Algorithms would analyze sensor data and provide insights that could be used in predictive analysis and advanced applications

Sensors are devices that interface with the physical world, and usually to convert physical parameters (for ex: temperature, humidity, magnetic field, light, sound, etc.) to electrical signals which a machine can understand and a software program can process.

Sensors on an automobile are vital in determining the health of important components and also are an integral part of the systems that provide reliable information for autonomous driving. Numerous sensors also aid the safe operation of airplanes.

This paper focuses on the role of sensors that are vital in providing the information necessary for software algorithms to process and apply the rules that govern the application's features. We examine how new sensors and sensing techniques could take these applications to a higher scale of efficiency.



**Figure 1 - Components of a Typical IoT Application**

## **Sensors in IoT**


IoT applications influence both residential and commercial landscapes and are an important aspect of the business services they offer -- be it providing peace of mind by protecting the customer's house, or by maintaining the production line at a factory. Sensors in these systems provide critical information that can determine further courses of action in the system.

Sensor choice, positioning and the data model depend on the product use cases and design. In-home security applications, for example Xfinity Home, use various types of sensors in a home to detect anomalies towards providing comprehensive security. The following are some of the sensors used in a typical home security application, which are depicted in Figure 2.

1. Motion Sensors
2. Door/Window Sensors
3. Security Cameras
4. Safety sensors (like smoke detectors)
5. Thermostats
6. Catastrophe sensors, for fires, floods, and weather-related dangers.


Most operators provide smart home applications packaged with sensors and detectors, like motion sensors, door window sensors, door locks, fire sensors, flood sensors and others as shown in Figure 3. The motion sensor is a passive infrared sensor (PIR). This sensor uses infrared to sense motion and sends data to a central control unit over a Zigbee network. The sensor detects motion events and passes this information over to the controller, which in turn processes it and translates it to a specific data model. The rules engine processes and determines if there is a need to send alerts. The door/window sensor uses magnetic fields between two ends of the sensor to signal when it is open or closed. Smoke detectors look

out for the presence of certain particles in the environment and generate alerts when those particles are present.



**Figure 2 - Sensors used in a home security application**

With advances in technology, most home appliances include a range of sensors and have the ability to connect via Wi-Fi or Bluetooth. Each of these sensors are designed to serve a specific purpose (Figure 3) and they connect back to a hub, typically on the home network. These sensors are usually paired to the hub via Zigbee<sup>1</sup>, Bluetooth<sup>2</sup> or Wi-Fi. Some sensors may have the ability to connect with an external network, so as to connect to a cloud-based application.



**Figure 3 - Sensors in a Home Security Application**

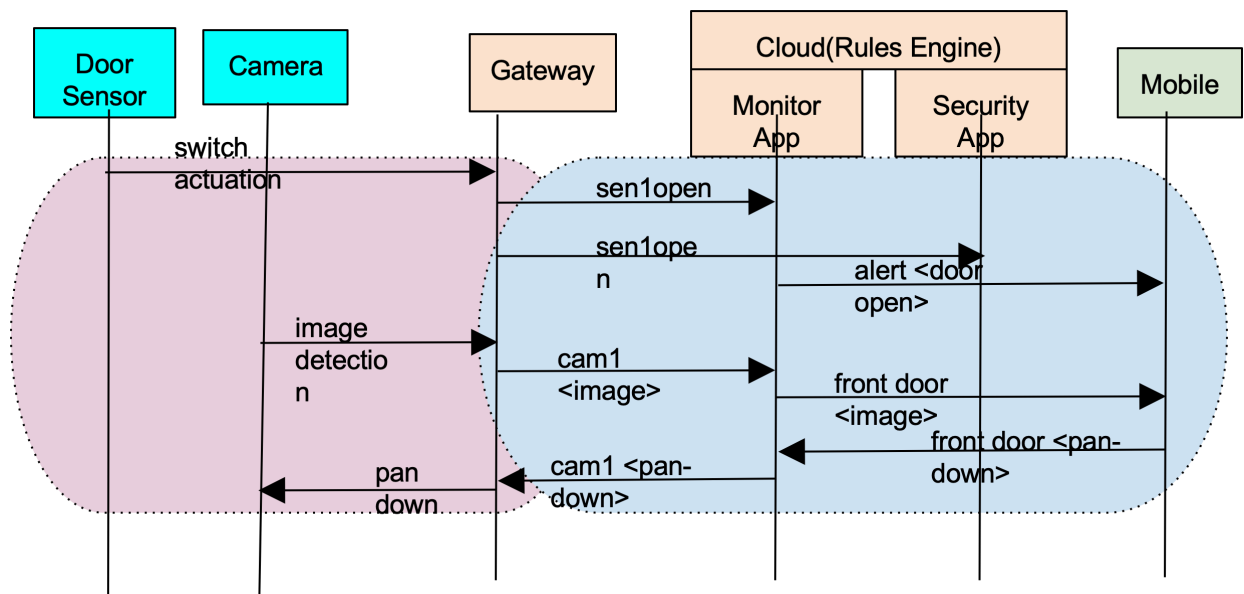
The example above is just one of many use cases that a typical IoT application would support. From a developer's perspective, this entails maintenance of all these components, on regular basis, so that the system functions without failure. Any update to one of the components has to be tested so that it doesn't impact the entire system. This includes updates to hardware and software components in the system.

<sup>1</sup> <http://www.zigbee.org/>

<sup>2</sup> <http://www.bluetooth.com/>



In any IoT network, multiple network protocols are involved, hence the need for a central gateway that can translate between protocols and forward packets for analysis. In the example above, the home gateway would be the IoT gateway, as it would have the necessary hardware and software to interact with devices which may use a different protocol. For example, there may be a rule set to turn on a light bulb upon detecting motion in the hallway. In this case, the motion detector could be a ZigBee-based device, and the light bulb could be based on BLE (Bluetooth Low Energy). In this case, the gateway would bridge the connection between the lightbulb and the motion sensor. Figure 4 shows a sample sequence diagram of an IoT application. In the following sections, we look at how newer sensing techniques and sensors could help bring in new features, improve efficiency, and ease system maintenance.



**Figure 4 - Example Sequence Diagram of an IoT application**

# New Sensing Techniques

The current family of sensors have been operationally effective for a long period of time. As a result, any changes to those sensing techniques involves a lot of research and require ample confidence in the technology in order for them to be adopted. This section lists some technologies that show a lot of promise, are not a “heavy lift”, in terms of adoption, and could work with the existing systems. The communication mechanism of these newer sensors with the rest of the system is compatible with the sensors and gateways already in the market (WiFi/Zigbee/Bluetooth.)

## 1. RF Sensing for IoT Applications

RF sensors have been around in the scientific community for decades, and are prevalent in avionics and military applications; they also hold great promise for residential IoT applications. We have been watching this space closely and have observed significant advancements, in terms of algorithms and sensors.

As mentioned, most IoT networks employ passive infrared/PIR-based motion sensors and cameras for surveillance applications. While these sensors work well and provide the desired results, there are some challenges. These are line-of-sight sensors, for starters, and thus cannot detect motion through a wall. Cameras are used mostly for outdoor surveillance, and are often perceived as being intrusive for indoor tracking (and, notably, activity monitoring in eldercare environments). These challenges could be easily overcome by instead using RADAR (Radio Detection and Ranging)-based sensing. RADAR has been historically dominant in aviation and military applications, but is increasingly applicable to residential and commercial applications. RADARs can detect presence and motion through walls, while providing good coverage, without cameras.

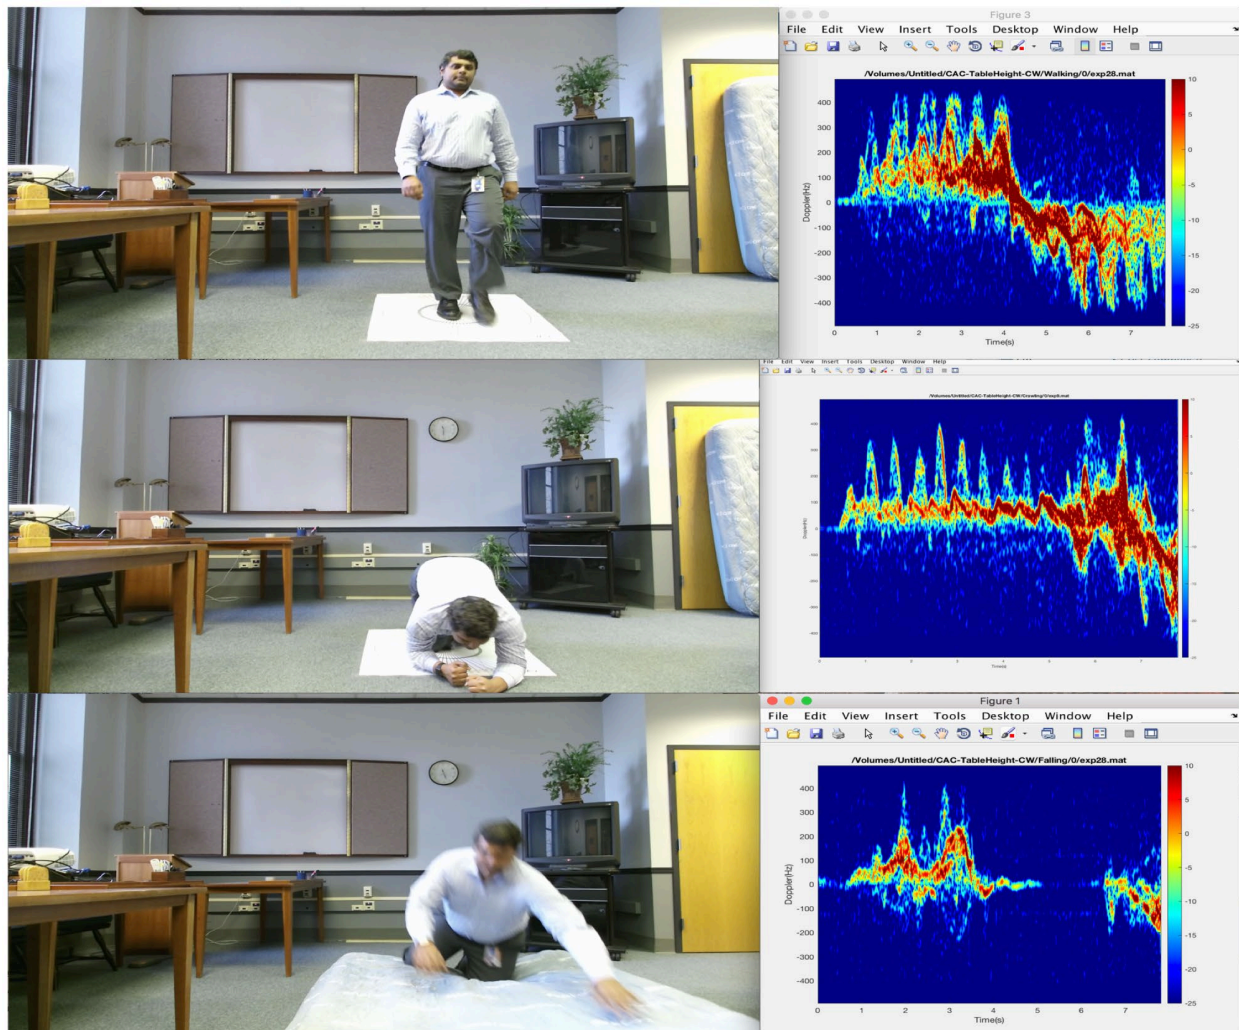
RF(Radio Frequency) devices that are compact, accurate, reliable, and inexpensive are currently commercially available. Over the past few years, attempts to apply such devices to biomedical measurements has increased. Although some studies applied these devices to medicine and health care, such research is still in its infancy. Nonetheless, radio-frequency sensing techniques -- originally developed for military applications, and later applied to search and rescue operations, such as locating earthquake survivors buried under rubble -- all carry plausible applicability for health care and home monitoring use cases, like aging in place and smart homes.

Doppler RADARs can be used to implement motion classification, which can be used for applications like activity monitoring, fall detection and Personal Emergency Response Systems/PERS. Traditionally, fall detection applications employ a wearable or push button device that needs to be activated after a fall occurs. The intent is to help the patient trigger emergency assistance. Using RADARs for fall detection is non-intrusive and does not need require manual intervention (pushing the button) to trigger an alarm. This matters because in many cases, people lose consciousness after falling, which obviates the applicability of such “push to enable” help calls. A RADAR, by contrast, monitors activities and can both auto-detect a fall, and raise an alarm. With machine learning capabilities, RADAR-based devices can also learn over time, to then more accurately detect falls.

Another form of activity monitoring enabled by RADAR is biometric, or the application of statistical analysis to biological data. Specifically, RADAR can provide effective, non-invasive and non-restrictive sensing techniques to acquire vital signs. For instance, RADAR could be used to monitor characteristics including body surface vibrations, mental state, and sleep apnea. For instance, RADARs can detect minute vibrations on the body surface, such as those induced by heartbeat and respiration. Simple equipment can

be used to self-monitor certain medical parameters or conditions, as well as to acquire related data required for senior living homes as well as medical facilities.


RADARs provide two specific attributes that provide accurate details about a person's locations and movement. FMCW (frequency-modulated continuous wave) RADARs provide range<sup>3</sup> and microdoppler<sup>4</sup> as the primary attributes. These parameters help an application to know the location of a target, and also effectively track movement of the target. Hence presence detection and motion detection are made possible, with a much higher resolution, to provide details on specific motion artifacts of the target (walking, sitting) including anomaly detection like fall detection. Doplar RADAR signatures are illustrated in Figure 5 (and the author wishes he had a nickel for every time he had to train the model to learn the "falling" artifact!)



**Figure 5 - Doppler Signatures for different motions**

<sup>3</sup> <http://www.radartutorial.eu/01.basics/Distance-determination.en.html>

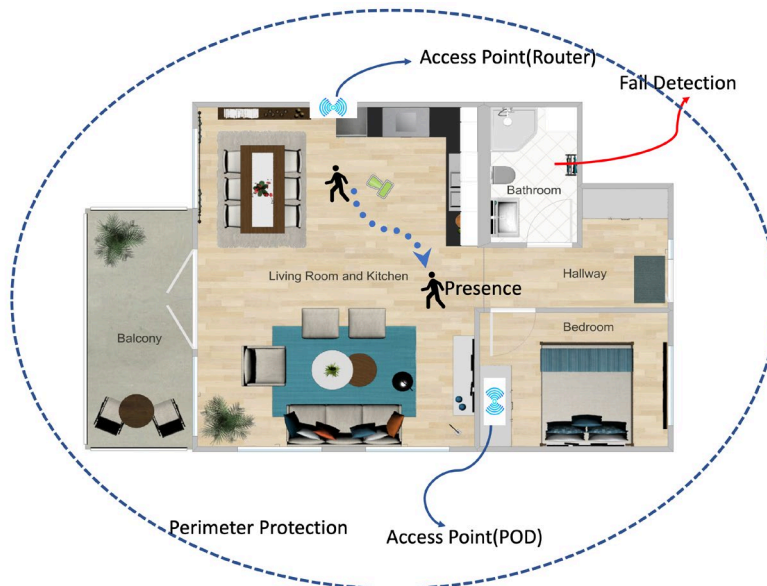
<sup>4</sup> <http://www.radartutorial.eu/02.basics/Continuous%20Wave%20Radar.en.html>



**Figure 6 - Range Data for Walking and Falling using RADAR**

6 shows the ability of a RADAR to detect motion artifacts, which could be used in many of IoT-based applications like healthcare and home security. 7 shows the ability of the RADAR to measure the distance from the AP (Access Point) and the target. This is helpful in detecting presence and location inside the home. The sensors would provide accurate position information, at the same time detecting postures. Perimeter protection is also possible with this technology. Since RF waves can penetrate walls, there is not a requirement for one sensor per room, which carries potential economic incentives that are beyond the scope of this paper but nonetheless worth acknowledging.

RADARs have already been used in cars to provide features like collision warning and prevention systems, as well as parking assistance and blind spot warnings. Extensive research has been done to ensure human safety while using RADAR systems, which have been proven to be safe. RADARs are thus a favorable sensing technology that IoT systems could use to enhance range and performance.




**Figure 7 - Home Security/Automation using RF Sensing**

## 2. WiFi as RF sensor?

Improvements in Wi-Fi technologies have made it possible to add more signal processing capabilities in the Wi-Fi chips, to function as RADAR. Specifically, a portion of the antenna could be apportioned to act as a RADAR, which would enable all the features supported by the RADAR previously discussed.

Newer Wi-Fi chips support elastic MIMO (multiple-input and multiple-output), which would give application developers the flexibility to dynamically choose the antenna pairs for multiple purposes. The system could use two antenna for RADAR applications, for instance, and the remainder of the antenna for Wi-Fi communications purposes. When there is no demand for communications traffic, more antennae could be in RADAR mode, for higher resolution imaging and better overall coverage. Figure 8 shows an example of an elastic MIMO configuration. Figure 9 shows the hardware and software components of such a sensing system. Using Wi-Fi is advantageous in that it makes use of existing hardware, and as such simplifies system maintenance.

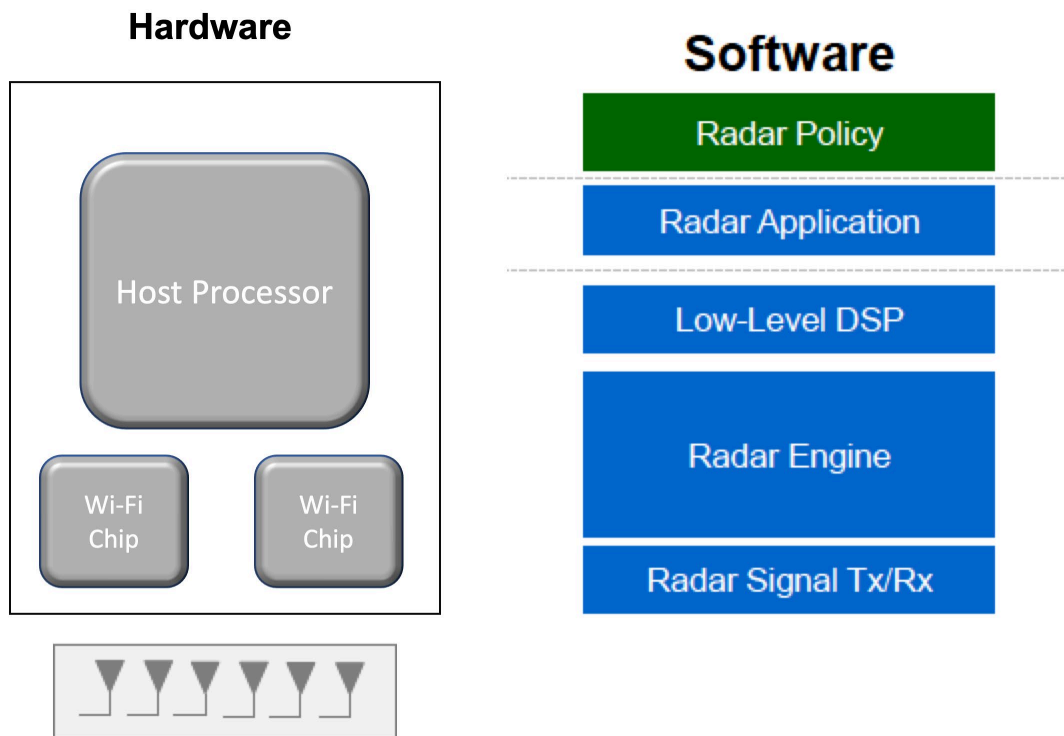
The data collected from an 802.11 AX<sup>5</sup> system shows that the Doppler signatures are similar to a traditional RADAR. Also, Wi-Fi pods help in expanding coverage and provide multiple data points to improve efficiency. Figure 10 shows Doppler signatures derived from a Wi-Fi system (top) and from a traditional 24Ghz RADAR system (bottom). These two look similar and hence algorithms that are developed and trained using traditional RADAR system could be easily ported across to a Wi-Fi system.



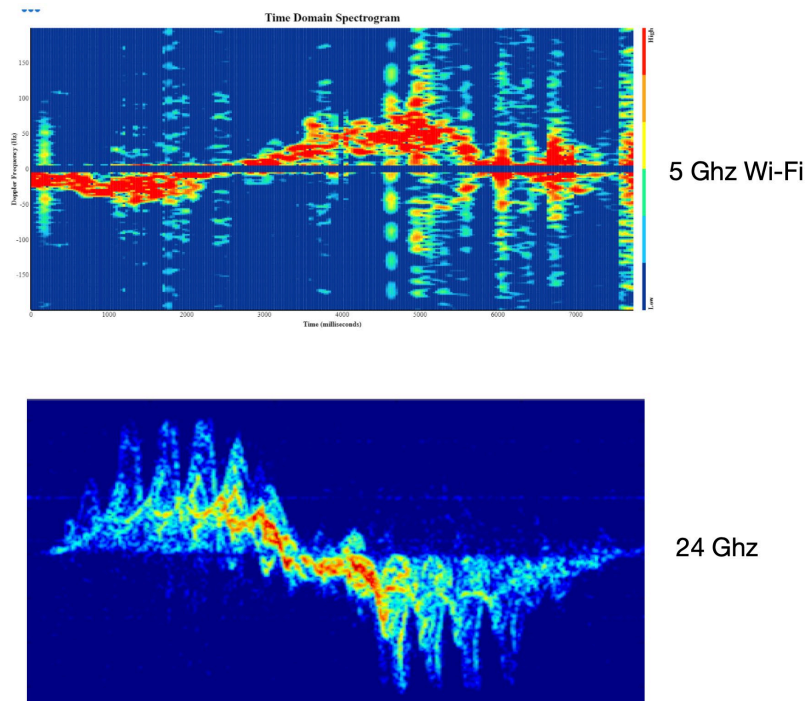
**Figure 8 - Example MIMO system and WiFi RADAR Hardware and Software**

<sup>5</sup> [https://en.wikipedia.org/wiki/IEEE\\_802.11ax](https://en.wikipedia.org/wiki/IEEE_802.11ax)





**Figure 9 - Components of a Wi-Fi RADAR system**




**Figure 10 - Doppler Signatures of Wi-Fi and traditioal RADAR system**

### 3. Audio-based Sensing

Acoustics analysis solutions can work when vision may be impaired, when obstacles are blocking the RADAR, and when sensors may not apply. Beyond the obvious solutions of using voice recognition, acoustic analysis applications can apply specific behaviors to the acoustic signature. Based on samples from an audio library, a list of events could be developed to track specific events. Using deep learning algorithms, sounds such as “door closing” or “water running” could be classified for correlation with daily activities.

Algorithms could be trained to detect anomalies like glass breaking, garage door opening, fire alarms and even calls for help. Figure 11 shows an example of a “confusion matrix” with a variety of sounds.

The diagonal represents correctly detected sound samples. Apart from the diagonal, the higher the number, the higher the chance of sounds getting confused with one another, which represents a high correlation.



**Figure 11 - Confusion Matrix for various sounds**

Ambient noise can be widespread in the home environment, and as such must be accounted for. Any acoustic monitoring technology must use a machine learning algorithm, to distinguish the targeted sounds correctly. Audio attenuation can aid in echo location, but can create challenges for acoustic analysis. Another challenge is that some sounds are quite similar and are therefore difficult to distinguish, even with machine learning.

Examples of highly correlated sounds:


- Children playing outside vs street music
- TV vs. conversation within a house
- Gunshot vs dog bark
- Engine idling vs. air conditioning

There are currently no commercially available products for IoT applications which use acoustic analytics. While the use of deep learning algorithms in acoustic analysis is a fairly common practice, their use requires the development of individual modules for detecting sounds applicable to IoT applications. The technology itself is quite nascent.

Microphones capture dialog in addition to the environmental noises. Microphone use can also raise privacy concerns. Similar to the concerns around video analytics, as the technology evolves, there will be opportunities to abstract certain acoustic content, and secure it, so that privacy concerns are minimized.

Microphones for audio capture are often much less noticeable and intrusive than other monitoring technologies. Many devices are being added into the home which capture and process voice commands today. Amazon Echo™, Google Now™ and Apple Siri™ are examples of common voice capture solutions in use today.

There are several ways that acoustic analysis can be achieved, such as detecting sounds by running the raw audio input through a Fourier Transform. This calculates the frequencies per given unit of time, which can be displayed in a spectrogram. Once audio has been converted to a spectrogram, visual analytics engines, such as ConvNet, can be utilized for audio classification. Figure 12 shows an example of an acoustic event displayed as a spectrogram. Studies have shown that this technology could be applied to detect age, gender and also voice recognition, which would help to personalize the experience for the customer.



**Figure 12 - Example Spectrogram of an Audio Signal**



## 4. Predictive Analysis using Advanced Machine Learning

Significant progress is being made to adapt AI and ML techniques to IoT applications. These algorithms have played a significant role in offering new features and improving the efficiency of IoT applications. Algorithms exist, for example, to protect the customer's network and devices. That matters because the increase in IoT applications has led to an increase of attacks, such as bots that exist to compromise the home network – and also the entire network, if there are unknown network vulnerabilities.

Network security is an important aspect of providing peace of mind to customers. Many applications offer network security as a service, in order to protect a customer's data, network and devices. There has been significant research about how customer data could provide detailed insights and hence offer some of these features without using any sensors. It must be noted that such techniques carry huge implications for privacy, and as such need to be dealt with great care to ensure that no privacy-critical data is being used without proper consent from the customer.

Edge compute is another topic of interest, because newer chips can provide some of the advanced compute to the network edge and hence improve latency. An Edge compute system, shown in Figure 13, would also provide an enhanced environment of privacy – the data is being processed on-premise, and never leaves the home network. This would enable using sensors like cameras and applied computer vision technology. With this the model could be trained in the cloud and the model could be deployed at the edge.

Computer Vision-based applications could also be applied to anomaly detection and advanced features like facial recognition and object detection. (Facial recognition techniques are outside of the scope of this paper.)

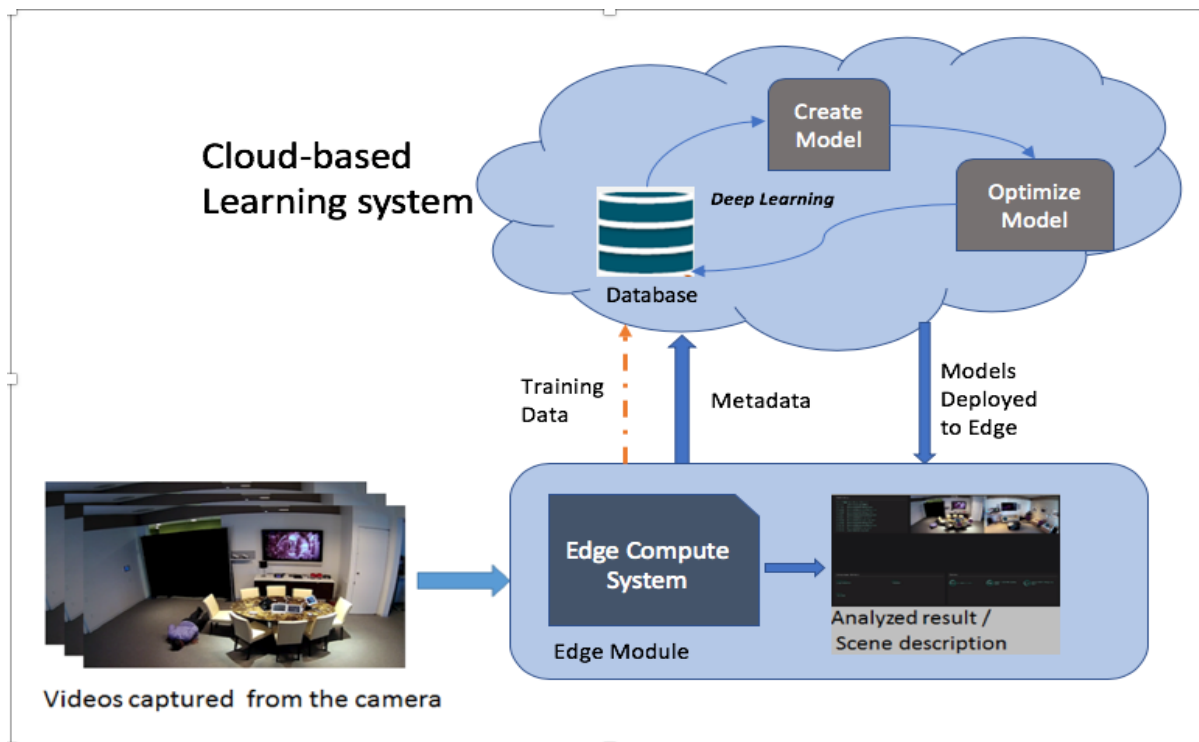


Figure 13 - Edge Compute System

# Conclusion

Multiple potential use cases and technological viability indicate that IoT technologies can definitively provide peace of mind to consumers. There are innumerable examples that can address the peace of mind factor in a consumer's life – whether it involves tracking a child, to ensure that he/she left school and reached home, or being able to check on elderly parents to ensure that they're doing ok, and, most importantly, to be alerted when they need help.

In this paper we examined the current state of sensors in an IoT applications, and then examined the potential of a variety of new sensing techniques. These new sensors are of interest because they provide advanced features while improving efficiency.

A wide variety of solutions, ranging from simple beacons using Wi-Fi, to elaborate, cutting edge solutions based on military technology, are edging to the forefront of the technologically-relevant in the IoT. The non-invasive nature and improved accuracy of RADAR, for instance, yields strong potential for IoT applications.

RADAR can track and monitor postures (activities) with reasonable accuracy. Some activities, like talking and reading, cannot be suitably monitored using RADAR, because RADAR's sensing capabilities do not cover acoustics or visual analytics.

The ability to use Wi-Fi signals to detect presence and to provide perimeter protection makes it easier to maintain, as it is a complete and widely deployed software solution that is easier to scale.

Acoustic analysis is very helpful in cases where visibility or RF penetration is low, and enables specific aspects in terms of security.

The paper also looked at the shift of critical data processing to the edge, which would reduce latency and also improve privacy, in that the data doesn't leave the home network. This would allow the application of Computer Vision- based techniques to a variety of applications.

Another important aspect that is under development is "Sensor Fusion". This involves employing different types of sensors and analyzing their data in unison to provide enhanced features – for example, combining RADAR, audio and/or video analysis. Such combinations would add another dimension to the sensor and provide greater overall detail.

Finally, we looked at predictive analysis, which, while interesting, carries considerable privacy implications which would have to be considered carefully and applied to very specific cases.

# Abbreviations

AI	Artificial Intelligence
AP	Access Point
BLE	Bluetooth Low Energy
FMCW	Frequency Modulated Continuous Wave
Hz	hertz
ISBE	International Society of Broadband Experts
MIMO	Multiple Input Multiple Output
ML	Machine Learning
PERS	Personal Emergency Response System
RADAR	Radio Detection and Ranging
SCTE	Society of Cable Telecommunications Engineers

# **Defining the Premise and the Edge**

## **A Managed Wi-Fi Service Application**

A Technical Paper prepared for SCTE•ISBE by

### **John Gammons**

Director, Wi-Fi Engineering and Operations  
Cox Communications, Inc.  
6305 Peachtree Dunwoody Rd. CTECH B09-105, Atlanta, GA 30328  
404-269-6992  
John.gammons@cox.com

### **Key Contributors**

**Michael Hurd**, Cox Communications, Inc.

**Kevin Klimek**, Cox Communications, Inc.

**Marco Lin**, Cox Communications, Inc.

**Jonilson Santos**, Cox Communications, Inc.

### **Acknowledgements**

**Nick Green**, Cox Communications, Inc.

**Brady Puckett**, Cox Communications, Inc.

**Altan Stalker**, Cox Communications, Inc.

**Drew Stravelli**, Cox Communications, Inc.

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
1. Managed Wi-Fi Services Overview.....	4
1.1. Use Cases.....	4
Content .....	5
2. The Problem .....	5
3. The Placement of Functions.....	5
3.1. Functional Overview .....	5
3.1.1. Access.....	6
3.1.2. Switching .....	6
3.1.3. Wireless LAN Controllers (WLC) .....	6
3.1.4. Wireless Access Gateway (WAG) .....	6
3.1.5. Routing .....	7
3.1.6. Firewall .....	8
3.1.7. Other Services .....	8
3.2. Considerations .....	9
3.2.1. General.....	9
3.2.2. Virtualization and Orchestration.....	9
3.2.3. Advanced Site to Site Network Requirements.....	10
3.3. Analysis.....	10
3.3.1. The Case for Centralization at the Edge .....	10
3.3.2. The Case for Distribution to the Premise Edge .....	12
4. Solutions by Vertical and Use Case.....	14
4.1. Edge - Public Hotspot and Guest Services.....	15
4.2. Premise Edge – Enterprise/School/Medical and Venues .....	15
4.3. Hybrid Use Cases.....	16
4.3.1. SMB and MDU.....	16
4.3.2. Multi-Service Customer .....	16
Conclusion .....	17
Abbreviations.....	18
Bibliography & References .....	19

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Next Generation Service Provider Network Model .....	5
Figure 2 - Functional Overview .....	6
Figure 3 - WAG Functions .....	7
Figure 4 - Service Overlays by OSI Layer 2/3 .....	9
Figure 5 - WAG Deployed in the Edge .....	10
Figure 6 - WAG Deployed in the Premise Edge.....	11
Figure 7 - Distributed Premise Edge Architecture.....	13
Figure 8 - Customer Segmentation .....	14
Figure 9 - Advanced Customer Integrations .....	14
Figure 10 - The Hybrid Approach .....	16

# Introduction

The next generation of Managed Wi-Fi Service offerings have distinct requirements and challenges. These offer unique challenges to both the business and the architects defining the underlying technology solutions. As architects and engineers are beginning to create solutions for these services, there are numerous requirements which need to be considered. The Service Provider will need to find the right balance between deploying these services in the Edge and the Premise Edge. The Service Provider may need to leverage multiple edges depending on the customer requirements. Careful analysis of the offering, the functions and the benefits will be required as Service Provider's aim to solve these fundamental questions surrounding these Managed Services. This analysis will review the offerings, the functions, analyzing the benefits of each architecture, and will provide directional guidance regarding the ideal deployment scenarios by use case, making the case for both Edge, the Premise Edge and a hybrid architecture for Service Providers.

## 1. Managed Wi-Fi Services Overview

The Managed Wi-Fi Service is a relatively new offering from Service Providers, and an important one. These services are driving new use cases and new demarcations for the Service Provider and the customer. Access speeds are now commonly approaching 1Gbps and beyond, and quickly eyeing 10Gbps on the horizon. The available cloud services which are being provided have made it easy for users to consume complex services with relative ease and reduced cost. Cellular and mobile networks have improved their reliability and general availability. Furthermore, device and network standards have evolved to improve compatibilities. These trends have converged in bringing about the BYOD (Bring Your Own Device) era, where the traditional Internet Service connection has transformed into the ability to easily consume wireless or wired access, in several geographical locations, and in a reliable manner. Consumers are also demanding more features and integration options in an intuitive and easy to use App-like interface where support is available but rarely required. This is the new Service Provider landscape and the new Managed Wi-Fi Service. These trends have continued to evolve across multiple industry verticals including:

- Public Hotspot
- Hospitality/Amenity/Guest Services
- SMB (Small-Medium Business)
- MDUs (Multiple Dwelling Units)
- Enterprise
- School/Campus
- Medical
- Venues


### 1.1. Use Cases

As we dive deeper into this offering, let's first analyze the use cases Service Providers are being challenged to deliver for these verticals. Typical use cases for these services include the following:

- Guest Wireless Local Area Network (WLAN), Local Area Network (LAN), Internet access service and associated services for end-user devices.
- Private Wireless Local Area Network (WLAN), Local Area Network (LAN), Internet access and associated services for end-user devices.

- Security and content filtering solutions
- Application Visibility, Monitoring and Analytics
- Advanced layer 3 routing and layer 2 transport protocols
- Network Authentication services
- Integration with Customer or 3<sup>rd</sup> party systems and services

In order to deliver these use cases successfully to the consumer, in a way that satisfies their BYOD (Bring your own device) mindset, the new Managed Wi-Fi Service offering is creating a need for a next-generation network deployment model for the Service Provider. As the demands on the network and the cloud increase, services in the SP Cloud and even the Internet are distributing to take advantage of edge computing. Gartner defines edge computing as “solutions that facilitate data processing at or near the source of data generation”. For Service Provider to client communications, this could mean distributing to the edge of their cloud, or the premise itself. This traditional SP premise delivery architecture is being augmented to include one or more access points and switches deployed downstream of x86 based commodity hardware as access and services meet in these new Managed Wi-Fi Service offerings, as illustrated in Figure 1 below:



**Figure 1 – Next Generation Service Provider Network Model**

## Content

### 2. The Problem

While this next generation Service Provider network has advantages, the new capabilities offer new challenges and questions. Most notably, the following questions surface:

- Where does a Service Provider deploy these Managed Services functions?
- When does the Service Provider consider Edge deployments?
- When does the Service Provider consider Premise Edge deployments?

These questions are bound to generate discussions, some of which may even get heated amongst Network, Datacenter and System Architects.

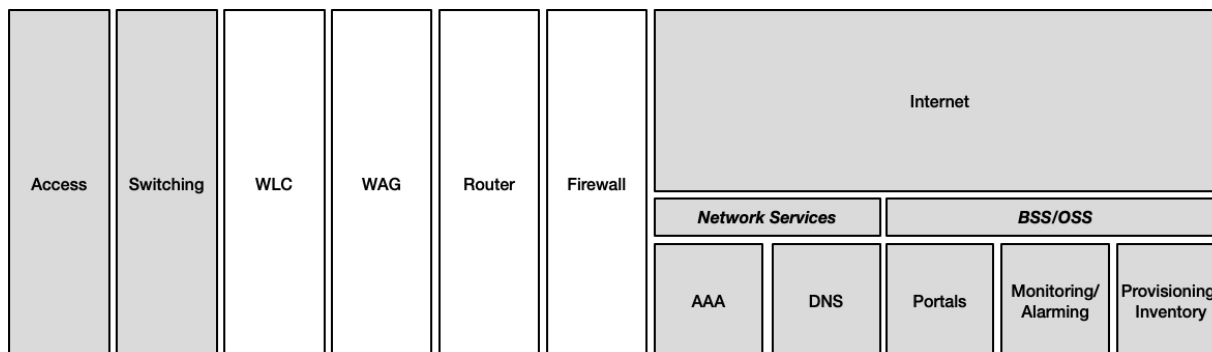
### 3. The Placement of Functions

#### 3.1. Functional Overview

Before we decide where functions are placed within the network, let's first analyze each function within the typical solution to ensure we understand the functions to be performed. The below illustration in Figure 2 shows the necessary functions to be covered as part of this analysis. The services which will serve as our primary focus are shown in white, while the others are shown for reference in gray to ensure



a proper understanding of the entire landscape of the solution. Each functional area is shown in its logical place within the network, and then subsequently defined below.



**Figure 2 - Functional Overview**

### **3.1.1. Access**

The access network is delivered through 802.11 RF based Wi-Fi or wired copper ethernet connections. These are physical components which are not able to be virtualized and are required to be on-site. These are not directly analyzed within this paper.

### **3.1.2. Switching**


The ability to have multiple access mediums connected to common or independent layer 2 broadcast domains. Some aspects of switching could be virtualized through a variety of protocols, and that is assumed to be the case. Since the construct could not be virtualized completely, it is shown to be out of scope.

### **3.1.3. Wireless LAN Controllers (WLC)**

Many enterprise and carrier-grade access points manufacturers incorporate a Wireless LAN Controller function into their offering. These systems provide intelligence in the form of Radio Resource Management (RRM) and Self-Optimizing Network (SON) functionalities for the access points. When WPA2, 802.1x, or 802.11r Fast-Secure Roaming are enabled, the WLC will also serve to ensure keys are cached and exchanged quickly without sacrificing the network integrity. These solutions are often available in virtualized container forms.

### **3.1.4. Wireless Access Gateway (WAG)**

The Wireless Access Gateway (WAG) is a critical component within the typical Service Provider Wi-Fi offering. It may also be referred to as a TWAG (Trusted Wireless Access Gateway) per the 3GPP and Evolved Packet Core (EPC) standards.



**Figure 3 - WAG Functions**

#### **3.1.4.1. DHCP**

This function assigns IP addresses to end user devices on the network for both wireless and wired mediums across the Guest and Private use cases explained above.

#### **3.1.4.2. Whitelist**

This is the ability to allow only certain protocols or destinations through the WAG while a session is in an unauthenticated state. An example of this allowing the ability for a user to render a splash page portal for registration or authentication on the network.

#### **3.1.4.3. Network Address Translation (NAT) or Carrier Grade NAT (CGN/CGNAT)**

Network Address Translation as described in RFC 1631 Traditional Network Address Translator is a function providing the ability for clients to exist within RFC 1918 Private Address Space using 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16 or RFC 6598 Shared Address Space using 100.64.0.0/10 but still reach the Internet. NAT is also “often accompanied by application specific gateways (ALGs)” to monitor application payloads and correct as necessary (Srisuresh).

#### **3.1.4.4. Session Policy Enforcement and Accounting**

This functionality refers to the ability to create and tear-down sessions as necessary to enforce network policies on individual devices for the purposes of authentication, authorization, and accounting. Sessions may include both wired and/or wireless network devices.

#### **3.1.4.5. Gateway and Router**

The WAG serves as a Layer 3 IP Gateway for the end-user devices. It also routes packets in between multiple networks or forwards packets upstream through POPs or Peering locations to the Internet.

### **3.1.5. Routing**

This function implies the ability to have multiple customer networks and the ability to optionally route between multiple networks and the hosts within them using traditional Layer 3 routing techniques. This is listed separately from the WAG in case there are more advanced requirements such as MPLS, VPLS, VXLAN, SD-WAN or other requirements. This is shown upstream of the WAG however it could alternatively exist downstream of the WAG, especially if it is a Layer 2 transport protocol.

### **3.1.6. Firewall**

The firewall function has the ability to analyze and limit both external and internal communications across the physical interfaces. Packet analysis may include source, destination of both the hosts and the ports/protocols. Tracking state or session awareness is optional and not pertinent to this paper and analysis.

### **3.1.7. Other Services**

The following other services may be included in a Managed Wi-Fi Services offering but have limited scope within this specific analysis.

#### **3.1.7.1. Authorization, Authentication and Accounting (AAA)**

This service is responsible for authenticating and authorizing devices in accordance with the defined network policies or proxying as appropriate to upstream systems including partners when policy and business rules dictate. The AAA is also responsible for accounting policies and processing near-real time on the network at defined intervals and delivering those records to a data warehouse infrastructure as necessary. For the purposes of this paper, we will abstract any potential upstream authentication, authorization or accounting systems into this function for simplicity. Communications will typically occur over RADIUS or DIAMETER between the Session and Policy Enforcement point and the AAA. This is assumed to reside in the Service Provider's Edge or their Cloud, depending upon their specific requirements.

#### **3.1.7.2. DNS Resolution (Proxy)**

This functionality provides the end user device with the ability to resolve Domain Names to IP addresses as described in RFC 1034 and 1035. This is assumed to be provided by the Service Providers traditional network services infrastructure. It is only referenced as consideration for proxy services lower within the network, likely in the WAG as an extension of the Gateway function.

#### **3.1.7.3. Portals**

Portal functions include Graphical User Interfaces which are available over HTTP/HTTPS protocols and allow unknown end users to authenticate to the network. They also include the ability to allow customers or Service Provider operational support personnel to manage their end user experience, manage their network configurations and to also view the status and analytics of the platform.

There are extreme cases such as venues where a significant number of authentication transactions may occur within a short amount of time. While these scenarios may warrant distribution of a portal application to the Edge or Premise Edge, that will not be a significant consideration in this analysis.

#### **3.1.7.4. Monitoring and Alarming**

Any Managed Service offering delivered by a Service Provider should include the ability to monitor and alarm upon any event that is deemed actionable or any event that either the SP or the Customer have deemed to require awareness. The exact placement of this function could vary, but it is assumed that the SP has existing infrastructure or requirements which would dictate the placement.

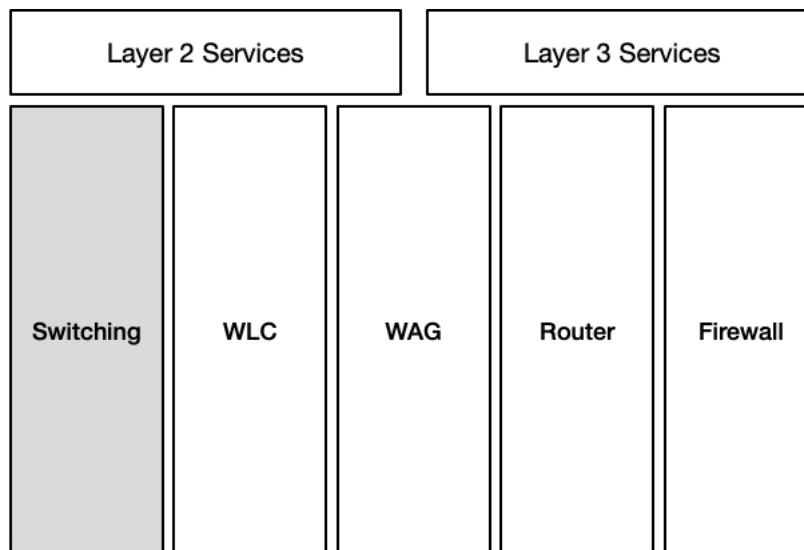
### 3.1.7.5. Provisioning and Inventory Management

Service Providers will require back office integration with both Provisioning and Inventory management systems. These systems are out of scope for this analysis.

## 3.2. Considerations

### 3.2.1. General

As we first consider these functions, and their placement, it is important to note the network layers for each function. Consider the following logical diagram of functions, which makes a number of assumptions and simplifications on the WLC, WAG, Router and Firewall placement for the purposes of generalization:



**Figure 4 - Service Overlays by OSI Layer 2/3**

Depending upon configuration the ordering of these functions can certainly shift. The assumption for this analysis is that these are directionally correct, and more importantly linked directly or through Service Function Chaining (SFC) in either the Premise Edge or Edge to ensure these functions are compatibly linked to one another. The Service Provider would have deployment flexibility to stitch these together in a number of different ways to meet the individual customer requirements. As an abstract, some part or all of these functions are assumed to be required for deployment to fulfill the Managed Wi-Fi Service requirements of the customer. The remainder of the analysis will often times generically refer to an abstract functional representation of the WLC, WAG, Router and Firewall as the Managed Service.

### 3.2.2. Virtualization and Orchestration

Both the Edge and Premise Edge architectures can take advantage of Virtualization, Orchestration and Service Function Chaining. The ability to deploy container-based services and functions on-the-fly for customers is attractive in all cases. The Premise Edge deployment scenario would require more “clusters”, which would drive complexity of those solutions and overall costs to deploy and operate orchestrated virtualization environments such as PaaS (Platform as a Service), IaaS (Infrastructure as a Service) and NFV (Network Function Virtualization) solutions. However, the more distributed Premise Edge model would also likely improve the value proposition of orchestration for the Service Providers.

This is due to the fact that as the number of environments they need to manage increases and those environments become more dispersed, their business would see more overall value in orchestrated provisioning of those services.

### **3.2.3. Advanced Site to Site Network Requirements**

Any Enterprise customers which require advanced site-to-site layer 2 networks such as L2VPN, Metro Ethernet, or L3VPN networks could be ideal candidates for either deployment methodology. While there are benefits to centralize and leverage an SD-WAN or MPLS/VPLS technology to aggregate these networks at the Edge, there is also the opportunity for the Premise Edge to create these overlay networks. Due to this, many of the above factors would have to be considered alongside their exact site-to-site requirements as well in order to best meet the customers individual needs.

## **3.3. Analysis**

### **3.3.1. The Case for Centralization at the Edge**


There are numerous arguments for centralization of the Managed Service functions. The arguments which typically drive centralization are based within the following principles:

- More easily deliver larger mobility domains
- Improved operations and reliability
- Reduced total cost of ownership

We will discuss and analyze these aspects further in the sections that follow.

#### **3.3.1.1. Mobility**


The traditional role of a WAG within a wireless architecture is to function as the border of a mobility domain to anchor sessions higher in the network and create session mobility across small or large portions of a network. To accomplish this WAG terminates Layer 2 and serves as the IP gateway for end user networks. In order to have WAG exist in the Edge, either the layer 2 network domains will need to be extended from the Premise to the Edge, or an equivalent layer 2 over layer 3 overlay or tunneling protocol will be required. The end user will have a seamless experience as they roam from Premise A to Premise B, assuming ubiquitous RF coverage as the session is anchored in the SP Edge. This is illustrated in Figure 5 below:



**Figure 5 - WAG Deployed in the Edge**

As the WAG is distributed to the SP Premise Edge, the mobility domains are also distributed. In order to maintain a seamless experience, sessions will need to be anchored between Premise A and Premise B WAGs either through the use of direct tunnels or another layer of aggregation and anchoring such as is available within an EPC, or Local Mobility Anchor (LMA) in the case of PMIPv6. This adds additional

cost in the form of deploying more WAGs in addition to more layers of complexity to provide a larger mobility domain.



**Figure 6 - WAG Deployed in the Premise Edge**

In both cases, the architecture is very similar, only potentially introducing different levels of overhead to transport Layer 2 packets to the WAG and altering the location of the deployed equipment. As the WAG location is determined, and more importantly where end-user layer 2 is terminated, one can determine how overlay managed services are applied for the customer. It is important to note that in both cases there are standardized means to enable session mobility across multiple WAGs, however these technologies invoke additional cost as well as additional complexity and overhead which will need to be balanced.

It is important to understand the implications of mobility. A customer or service which has requirements for a large mobility domain will typically require some centralized Edge or Cloud elements in order to accomplish mobility across numerous locations. This should be one of the primary considerations when determining if a function or service should be deployed on the Edge or Premise Edge. Another consideration in mobility, is security of the RF layer. WPA2/3, 802.1x (and all EAP types therein), as well as 802.11r all have implications to mobility as well. These requirements will need to be well understood and the interconnections between the functions will need to be designed to meet the exact customer requirements. For all of these reasons, use cases where a mobility domain is required over a larger area prefer deployments within the Edge.

### **3.3.1.2. Improved Operations and Reliability**

As solutions centralize, the Service Provider can create larger clusters that serve more customers. Even if they are supporting the same number of virtual instances, having them coexist on less physical hardware will utilize less operational resources. Additionally, having them in fewer geographical locations would enhance their supportability and overall reliability of the service for the customers. Centralization should therefore directionally improve the overall operation of the solution. Quantification of the improvements would vary greatly based on the individual Service Provider circumstances and would need to be analyzed for each individual SP and for each service offered.

In general, for the use cases and customers that this paper is referring to, the Service Provider network is assumed to have improved datacenter facilities when compared to the customer premise. This will likely include more highly available and consistent connectivity, power, cooling, and physical security at the SP facilities. As the size of a customer's network declines so will their environmental protections. As an example, consider the following potential deployment examples of customer premise equipment:

- on a table or desk in the corner of a common area such as a breakroom where coffee is spilt
- setup under an employee's desk where the cables and equipment are stepped on multiple times a day

- plugged into an extension cord that gets unplugged in favor of a vacuum each evening

These are all too common occurrences for Service Provider support personnel. Inversely, consider the Service Provider facilities, where typically there are varying degrees of criticality assigned to datacenters and the greater the impact, the greater the environmental protections which are put in place to protect the facility and therefore improve the services they provide. With this logic, centralization will correlate to improved environment availability. This may also imply that smaller customers may see additional benefit in more centralization than larger customers, as they are less likely to have taken any precautions to improve the provisions at their facility.

### **3.3.1.3. *Reduced Total Cost of Ownership***

As resources including compute, memory and storage are centralized there are direct financial gains to the business. There are a number of reasons for this including:

- Economies of Scale
- Reduced Operating Expenses
- Improved abilities to utilize/repurpose available resources

Oversubscription of the network is an assumed part of the business case. Centralization allows what would otherwise be Premise Edge hardware to be centralized and therefore oversubscribed based on true peak utilization. Additionally, as the solution centralizes, there are opportunities for multiple customers to share the same virtual instances within the Edge as well. In addition to the hardware, this would lead to even greater utilization of the software and subsequent licensing. There are a number of factors to consider including the different market segments and what that may mean for peak utilization of the Edge. As an example, the same Managed Service delivered to a dinner restaurant and a coffee shop will likely have distinctly different peak times when customers and employees are utilizing the services. Allowing the same hardware/software to serve both segments will allow those resources to be maximized. If this is done across the scale of the entire Service Provider, this could mean significant reductions of resources and overall cost to the business.


While there are technologies which could take advantage of underutilized resources on the customer premise, these would result in additional utilization on the network which is less than ideal. Similarly, local traffic which does not need to traverse the backhaul or the core should not as this creates undue network utilization and drives network costs. Therefore, centralization is only beneficial under the appropriate use cases.

### **3.3.2. *The Case for Distribution to the Premise Edge***

There are many considerations and requirements which would drive the Managed Service to be distributed to the Premise Edge. In every case, individual customer requirements must be analyzed to ensure that all considerations are taken into account. Generally speaking, the requirements which drive the Managed Service functions to be deployed on the Premise Edge fall into the following categories:

- Intra-premise Communications
- Customer Security and Segmentation
- Advanced Customer Integrations
- Performance

We will discuss each of these in greater detail in the following sections. Please reference the diagram below in Figure 7.



**Figure 7 - Distributed Premise Edge Architecture**

### **3.3.2.1. Intra-Premise Communications**

The primary driver for distribution is the flow of customer traffic. Anytime there are requirements to route between multiple networks, the packets will have to traverse the Managed Service. The Premise Edge deployment is able to keep all premise communications local to the premise, whether they are layer 2 or across distinct layer 3 networks. These would include customers with local file shares, client to client communications or sharing, or even printing.

Key customer requirements that may drive this include:


- Substantial Client to Client communications
- Multiple premise networks with interconnectivity requirements
- Mission critical Client to Client communications even through an Internet disruption/maintenance activity

### **3.3.2.2. Customer Security and Segmentation**

Certain aspects of the solution may have specific requirements which require additional network segmentation. These situations may create considerations with regards to which functions are deployed on-premise in order to improve the segmentation between customers. As an example, consider a customer with locally hosted servers/services which are managed by the customer or an independent 3<sup>rd</sup> party but reside on the customer premise. These could range from an on-premise video camera termination appliance, to a server which terminates network authentication requests to 802.1x/LDAP/Active Directory databases. As these network, firewall or proxy functions move upstream into the Service Provider's Edge there will inevitably be multiple customer networks overlapping on the common Edge. While these can be virtualized to segment customers from one another, there is still a risk that these hooks into the customer's private network could be compromised by another customer.

As these functions are distributed to the Service Provider Premise Edge, the amount of separation between the customers increases. From a security perspective, physical separation is always preferred over virtual segmentation. The improved segmentation is certainly preferred by not only the Service Provider, but more importantly the customer. In the centralized model, the connections would have to be secured against both the public internet, as well as other customers. For extremely secure applications, when leveraging the Premise Edge, the Service Provider could recommend that the customer keeps certain networks completely isolated from the Internet and from other customers, to further limit the potential impacts. Consider the figure below which shows these potential use cases:

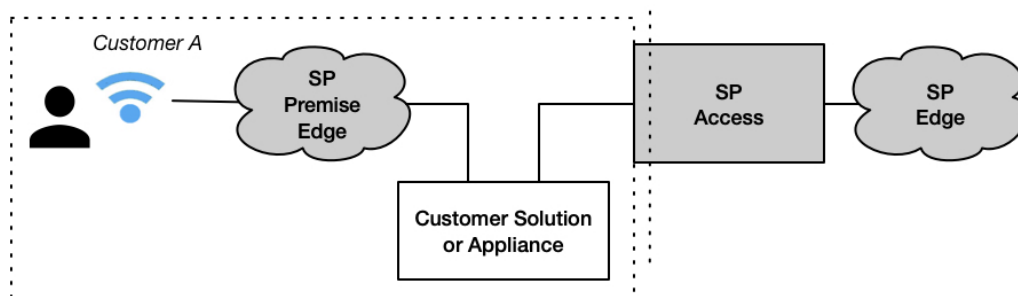




**Figure 8 - Customer Segmentation**

### 3.3.2.3. Advanced Customer Integrations

Customers may require that the Managed Service integrates with other premise-based solutions or services. These 3<sup>rd</sup> party solutions may be owned by the customer or may be provided by 3<sup>rd</sup> party services that the customer has chosen. These could include enterprise or venue customers which have special requirements for visibility to outbound traffic. These may exist out of line and simply require a tap, or a port devoted to the monitoring and replication of traffic such as a SPAN port. Other examples may include more intrusive integration such as a traffic shaping appliance or firewall as shown in Figure 9 below. It is important to acknowledge that the Service Provider may choose not to support these types of integrations. However, these complex integrations may also be a necessary challenge that is associated with deploying Managed Services.



**Figure 9 - Advanced Customer Integrations**

### 3.3.2.4. Performance

In any network architecture and across any medium, shortening the path always equates to less latency and a higher performing network. Inversely, adding to a path, especially with additional unnecessary hops or processing, degrades the network. Higher performance applications will almost always require additional dedicated resources devoted to making the best use of the network. In this case, the additional overhead and wrappers of tunneling are not preferred, and in these cases, neither would the shared infrastructure of the Edge. Therefore, higher performance applications will typically prefer distribution.

## 4. Solutions by Vertical and Use Case

We have a clear understanding of the problem, a solid grasp on the functions, and have explored the considerations which would drive the Managed Service placement. The analysis of the architecture benefits thus far has shown that centralization at the Edge is necessary when mobility requirements

demand it and is preferred whenever possible due to improved operations and reliability while also reducing cost and improved utilization of resources. Distribution to the Premise Edge is required as the solution increases in complexity either due to site integrations or demands for increased security or performance. Let us next consider the offerings of the Managed Service we are offering to our customers and apply those use cases to our analysis thus far.

#### **4.1. Edge - Public Hotspot and Guest Services**

Public hotspot services typically will require a Service Provider to provide basic Internet Access services to a high number of temporary mobile or nomadic users. These use cases will typically be highly susceptible to peaks and valleys of utilization based on their individual traffic patterns. These networks are typically openly accessible as well, which tends to drive significant walk-by user connections as there are a number of devices which are configured to attach to nearby open networks by default. These networks will rarely have local switching and routing requirements such as premise file-sharing or printing.

The demands for high capacity and limited accessibility requirements make these use cases highly attractive for centralized deployments within the Service Provider Edge or even higher within the network. Large aggregate CIDR blocks of IP addresses can be allocated in these central locations and reused across a high number of end-user subscriber devices. The highly mobile and nomadic nature also allows these IPs to be configured to expire and be relinquished often. All of these traits combine to limit the strain of hotspot users on a Service Provider's network. Also as previously discussed, multiple verticals can also be balanced against one another. For instance, there may be a large public hotspot which is available in a park with lots of daytime foot traffic, and another hotspot a few blocks away along a corridor of night clubs. The same hardware, software and network resources that feed the daytime hotspot can then be repurposed to serve the night club hotspot at night with no perceived impact to the end-users whom are consuming the service. Hotspot services are the ideal use case for an Edge deployment model which will ease operational support, improve the service availability, while also providing a lower cost service for both the Service Provider and their Customers.

#### **4.2. Premise Edge – Enterprise/School/Medical and Venues**

Private Enterprise services will typically have more requirements for accessing local resources for transfers including printing, servers or client to client communications. They will also tend to have more requirements for segmentation between specific departments or areas of the building, which will drive their design to be more complex than a single flat network while possibly still requiring interconnectivity. They are also more likely to have premise-based 3<sup>rd</sup> party solutions which they require to be integrated. This will especially be the case at Enterprise customers where they are in the process of transitioning to a Managed Service. These customers have a lot to risk and may not feel comfortable swiftly transitioning to a Managed Service model for the entire offering. They may also have legacy network components which are tightly integrated to their business or just require significant efforts to successfully migrate. Some of these legacy services may also have security segmentation requirements as previously discussed. As larger business customers realize the Managed Services value proposition, they are more likely to implement these strategies in a cap and grow model. This will require the Service Provider to be prepared to aid the customer during this transition, which will likely drive complex integration requirements on the premise. Similarly, large venues tend to have very custom requirements depending on the exact venue type such as a Convention Center, Stadium, or Amphitheatre, and sometimes even depending upon the event that they are hosting. In both cases, these types of business customers will likely require the Premise Edge architecture to be deployed to ensure the necessary flexibility to meet their requirements.

### 4.3. Hybrid Use Cases

The above scenarios reflect the extremes, which are beneficial in determining the outer bounds of the solution. Our architecture analysis confirmed that both the Edge and the Premise Edge architecture offers benefits in certain use cases. Customers whom have solution requirements somewhere between those of a guest/public hotspot service and the large enterprise or venue service are another challenge. These hybrid customers likely include SMB and MDU segments. We will also need to discuss customers who require multiple services including both a guest/public hotspot and a private enterprise service.

#### 4.3.1. SMB and MDU

The SMB and MDU use cases are quite challenging architecture propositions. In most cases, these have slightly more requirements than a guest/hotspot service, but less requirements than a private enterprise. In short, these verticals have requirements for elements of both solutions. Drivers for the Premise Edge architecture likely include some premise local communications, a desire for the benefits of customer security and segmentation, and also include the occasional advanced integration on premise. Similarly, they also would benefit from the Edge architecture with improved operations and reliability, while also reducing the cost and resources driving a lower cost product for the customer. Arguably, both segments will have users that likely fall on either end of the spectrum – pun intended. This will likely result in some customers falling into both architectures, and the SP having reason to support both architectures, and possibly a business process which helps direct the customer to the right solution early in the process. This could be in the form of sales questionnaires or surveys. Effectively driving towards a Premise Edge architecture for the advanced customer, and an Edge architecture for the basic customer, with the flexibility to upgrade the basic customer should they suddenly require a more advanced integration or premise-based communication.

#### 4.3.2. Multi-Service Customer

The simplest solution for the customer which requires both a guest/hotspot service while in tandem delivering a private enterprise service, would be to deliver this based on the most stringent requirements. Since the large business private network would drive a Premise Edge hardware deployment, the hardware would already be deployed on-site, and dedicated to the customer, so leveraging this same hardware for the guest/hotspot service would be a plausible solution. An additional consideration is to tunnel guest traffic to reduce the size of the hardware deployed on-site and better leverage resources which are centrally deployed, or if tunneling an increase is needed to the size of the mobility domain. A hybrid approach may be a viable consideration, especially if the SP decides to deploy guest services for basic tiers in the Edge. If those services exist already, there may be significant costs reduced at scale by leveraging the Edge resources for Guest services. This architecture is shown below:




Figure 10 - The Hybrid Approach

# Conclusion

The Managed Wi-Fi Service and all associated features are an important offering for Service Providers. These service offerings are ripe with challenges and opportunities for both the business and the solution architects. There are numerous functions that are required to be delivered in the solution to successfully manage an access service for the end-users. Elegance in the delivery and the solution architecture will be critical to the Service Provider's overall success. While either methodology could be leveraged to deploy to any of these verticals, there are benefits to both the Service Provider and the customer in both the Edge and the Premise Edge architectures. The Premise Edge architecture favors premise-based local communications, customer security and segmentation, advanced customer integrations, and high-performance applications, making it ideal for the increased demands of private enterprise or venue deployments and similar verticals. The Edge architecture improves operations and reliability, while also reducing cost and better leveraging available resources, improving both the availability of the service and reducing the cost to both the Service Provider and the customer. This makes the Edge scenario a preferred architecture for guest or hotspot services.

Between these extremes are the SMB and the MDU, which may be best served by creating a tiered offering which supports both deployment methodologies. The final consideration is the customer whom requires both a private network with complex local requirements for traffic and integrations, while also requiring a guest or hotspot service. These customers would likely require the Premise Edge which could be utilized to meet the needs of both services. Ideally the Service Provider could deploy elements at both the Premise Edge and the Edge architectures for these customers, creating a hybrid delivery architecture. This hybrid architecture carries the burden of providing infrastructure at both the Premise Edge and the Edge, but which places the Service Provider in the best position to deliver the Managed Services of today, while being prepared for the Managed Services requirements of the future.

# Abbreviations

AAA	Authentication, Authorization and Accounting
AP	access point
BSS	Business Support Systems
BYOD	Bring Your Own Device
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
EAP	Extensible Authentication Protocol
EPC	Evolved Packet Core
IaaS	Infrastructure as a Service
LAN	Local Area Network
LMA	Local Mobility Anchor
OSS	Operational Support System
NAT	Network Address Translation
NFV	Network Function Virtualization
PaaS	Platform as a Service
PMIPv6	Proxy Mobile IPv6
RF	Radio Frequency
RFC	Request for Comment
RRM	Radio Resource Management
SFC	Service Function Chaining
SMB	Small Medium Business
SON	Self-Optimizing Networks
SP	Service Provider
TWAG	Trusted Wireless Access Gateway
WAG	Wireless Access Gateway
WLAN	Wireless Local Area Network
WPA2/3	Wi-Fi Protected Access 2/3

# Bibliography & References

Firmin, Frédéric. “The Evolved Packet Core”. 3GPP. 2019.

<https://www.3gpp.org/technologies/keywords-acronyms/100-the-evolved-packet-core>. Accessed July 15, 2019.

Rekhter, Y., Moskowitz, B., Karrenberg, D., J. de Groot, G., Lear, E. “Address Allocation for Private Internets”. RFC 1918. February 1996. <https://tools.ietf.org/html/rfc1918>. Accessed July 10, 2019.

Srisuresh, P., Egevang, K., “Traditional IP Network Address Translator (Traditional NAT)”. RFC 1631. January 2001. <https://tools.ietf.org/html/rfc3022>. Accessed July 10, 2019.

van der Meulen, Rob. “What Edge Computing Means for Infrastructure and Operations Leaders”. Gartner, Inc. October 3, 2018. <https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders/>. Accessed July 1, 2019.

Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., Azinger, A. “IANA-Reserved IPv4 Prefix for Shared Address Space”. RFC 6598. April 2012. <https://tools.ietf.org/html/rfc6598>. Accessed July 10, 2019.

# **How to Leverage SD-WAN to Accelerate Time to Market and Revenue**

## **The Importance of Service Assurance**

A Technical Paper prepared for SCTE•ISBE by

**Cliff Lane**

Principal Systems Engineer  
VMware  
1503 LBJ Parkway, Ste. 700  
Farmers Branch, TX 75234  
cliffordl@vmware.com

**Kishan Ramaswamy**

Senior Product Manager  
VMware  
1503 LBJ Parkway, Ste. 700  
Farmers Branch, TX 75234  
kishanr@vmware.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Background .....	3
1. SD-WAN's Agility and Speed .....	3
2. Defining the SD-WAN Challenge .....	4
3. Five Service Assurance Use Cases .....	5
3.1. Topology Discovery .....	5
3.2. Availability .....	6
3.3. SLA Management .....	7
3.4. Business Impact .....	8
3.5. Automation .....	9
4. Two Practical Examples .....	10
4.1. Managed Services Outage .....	10
4.2. Cross Network RCA .....	10
5. How to Become an SDN Provider .....	10
Conclusion .....	11
Abbreviations .....	12
Bibliography & References .....	12

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - SD-WAN Service Provider Solution Architecture .....	4
Figure 2 - Topological Mapping of SD-WAN Deployment .....	6
Figure 3 - SD-WAN Branch Site with High Availability Link .....	7
Figure 4 - SD-WAN Branch Site with SLA-Driven Services .....	8
Figure 5 - SD-WAN Branch Site, Network and Services .....	9
Figure 6 - SD-WAN Automated Alerting .....	9



# Introduction

The agility, simplicity and cloud-friendliness of SD-WAN may seem self-evident, but to win customers over, MSOs also need to convince enterprises accustomed to high reliability that they can meet and exceed those expectations. Five key elements of service assurance in a software-defined environment include automatic discovery of network topology, identification of any loss of availability, SLA management, awareness of business impacts, and automated processes. With sophisticated multi-vendor service assurance tools and turnkey models for entering this market, MSOs have the means to quickly engage and generate real revenue with SD-WAN.

## Background

More than a decade ago, the cable industry in the U.S. began to accelerate its venture into business services. In 2007, MSOs were generating between \$2 billion to \$3 billion in this arena. By 2018, they had grown that to \$18 billion. In recent years, the rate of growth has declined to roughly 10 percent. While still respectable, the current growth rate has motivated industry leaders to consider new initiatives, such as targeting more of the enterprise segment and expanding into new verticals.<sup>1</sup>

One promising technology that hits the enterprise segment across multiple verticals and leverages the industry's service delivery infrastructure is SD-WAN. Well-suited to organizations that are embracing digital transformation, using more applications from the cloud and hoping to augment their existing WAN without adding cost and complexity, SD-WAN is poised for growth. IDC estimates that the global SD-WAN infrastructure market will grow at a 30.8 percent CAGR from 2018 to 2023, to reach \$5.25 billion.<sup>2</sup>

There are two primary ways to deploy SD-WAN. In North America, according to Gartner, more than 60 percent of deployments have traditionally followed a do-it-yourself (DIY) model; whereas much of the rest of the world prefers a managed services approach.<sup>3</sup> To understand the demand on both sides, let's review SD-WAN components, architecture and service characteristics.

### 1. SD-WAN's Agility and Speed

The software-defined nature of this technology enables service agility, rapid rollout and instant-on WAN that delivers significant benefits to business customers. SD-WAN simplifies today's increasingly distributed branch networking by automating WAN deployment and improving performance over private connections, internet and wireless links.

Enterprises are open to changing or augmenting their connectivity solutions because of the need for alternatives that better support access to cloud applications, increased bandwidth demands and optimized network performance. Cloud-based applications, for instance, drive enterprises to more heavy reliance on internet connectivity and change traditional traffic patterns from a branch-to-data center to a branch-to-cloud. The legacy network design of traffic flow from branch-to-data center causes inefficiencies that impact cloud applications, while driving up costs and complexities.

From a service provider's perspective, the multi-tenant architecture of an SD-WAN orchestrator can enable easy provisioning of new customers and service management across multiple customers. Each

---

<sup>1</sup> "US Cable Set to Rake in \$18B in Biz Services Revenue in 2018," Jeff Baumgartner, Light Reading, Nov.11, 2018


<sup>2</sup> SD-WAN Infrastructure Forecast, IDC, July 2019

<sup>3</sup> Magic Quadrant for WAN Edge Infrastructure, Gartner, October 2018

business can be given access to its own self-service web portal for monitoring and policy configuration. A CPE or vCPE edge device can be provisioned with zero-touch deployment, which enables a service provider to deploy a new branch's WAN within minutes, as opposed to weeks or longer with legacy solutions.

The branch edge is provisioned either as a VM or customer hardware shipped to the branch office, where a non-technical person can plug in a few cables. Activation, configuration and ongoing management are all handled in the cloud. The edge device efficiently integrates LAN, WAN, Wi-Fi and 4G-LTE connectivity, along with other dynamic virtualized functions or business applications, such as next-gen firewall, application performance monitoring or dynamic multipath optimization. (See Figure 1.)

Once enabled, an edge device can automatically detect circuit characteristics, such as bandwidth, latency and more, then build a secure overlay network with SD-WAN gateways across all available links and begin steering applications per the configured policy. While these features and components, including gateways, are not universally available from all providers, they are key to SD-WAN's value proposition.



**Figure 1 - SD-WAN Service Provider Solution Architecture**

## 2. Defining the SD-WAN Challenge

The number of companies offering SD-WAN technology and services either directly to enterprises and/or through service-provider channels has grown into the multiple dozens, which makes for a busy and fluid marketplace. Exactly what SD-WAN providers and their service provider partners offer has come under scrutiny, and the divergences led the Metro Ethernet Forum (MEF) to propose common terms.

Embodied in the MEF 70 Draft Standard, released in May 2019, are four primary components: 1) edge, whether physical or virtual; 2) controller, which maintains centralized management of edges and gateways; 3) orchestrator, which handles the lifecycle service orchestration (LSO) of SD-WAN and other services – and could incorporate the Controller; and 4) subscriber web portal, enabling subscriber service ordering and modification.<sup>4</sup> The MEF definitions also specify the SD-WAN user interface, the underlay

<sup>4</sup> Draft Standard MEF 70 (R1): SD-WAN Service Attributes and Services, May 24, 2019

connectivity service (UCS), tunnel virtual connections (TVC), visibility into the application layer, control over the application layer extending to dynamic path selection, analytics tools, and more.

In addition to the recent definitions, there are other reasons why the SD-WAN market seems to be a work in progress. One analyst at Light Reading’s Future of Cable Business Services event in November 2018 indicated that a lot of enterprises were “just kicking the tires at this point” and that some service providers were “giving it away.”<sup>5</sup> While it is true that many enterprises are still getting their heads around SD-WAN and service providers are figuring out how to sell – and in some cases trying to avoid cannibalization of existing network revenues – we can offer further commentary.

As a leading SD-WAN provider in this space with more than 60 service provider agreements,<sup>6</sup> we know that enterprises are doing more than just considering SD-WAN. They are purchasing infrastructure and services, from both technology companies and service providers. One challenge, however, is that to win customers, you also need to convince them that you can assure their service.

### **3. Five Service Assurance Use Cases**

Service providers should aim to go beyond baseline specifications of SD-WAN. Service assurance is one way to do so. Many enterprise IT leaders, for instance, value the high reliability of their existing connectivity, including MPLS. From their perspective, any replacement or addition to it needs to not just maintain but also exceed existing performance.

The right SD-WAN solution can drastically improve overall performance, resiliency and security of hybrid and internet wide area networking. But to many enterprises, the technology is still new. If you’re delivering SD-WAN, you should take pains to explain how, whatever the underlying connectivity, it is mature enough to handle any incidents. Important ways to assure SD-WAN service include being able to automatically map network topography, identify a loss of availability, actively manage SLAs, understand how incidents impact business, and automatically identify any issues that arise. Here is more detail on these five use cases, their data sources and likely scenarios:

#### **3.1. Topology Discovery**


It is hard to manage what is out of sight, and what service providers need first and foremost is some means of dynamically discovering SD-WAN devices, gateways, networks, tunnels, customers/tenants, services and how they interact. The data for building these maps come from the orchestrator, the cloud-computing virtualization platform, the IP/MPLS fabric, configurations and profiles. (See Figure 2.) Invoking this function would involve the capability of discovering:

- Devices – The SD-WAN edge, gateway and hub devices at the customer premise, along with redundancy for both physical and virtual edges.
- Connectivity – All overlay (tunnels) to underlay from the edge device over LTE/3G, MPLS, DIA and termination to gateways, orchestrator and other edges, hubs (edges), and cloud services.
- Services – LAN-side and Cloud SaaS services profiled at the customer/branch SD-WAN edge.
- Customers – Identity of customer to device edge and infrastructure mapping, for both physical and virtual via integration with the orchestrator.

---

<sup>5</sup> “Vertical Systems: SD-WAN Not Profitable Yet,” Light Reading, Nov 29, 2018

<sup>6</sup> Magic Quadrant, Gartner, *ibid*.




**Figure 2 - Topological Mapping of SD-WAN Deployment**

### 3.2. Availability

Just-in-time root cause analysis can help identify any loss of availability that could impact and/or degrade service. As with automatic topology discovery, the sources for discovering any impact on uptime include fault data from the orchestrator, cloud-computing virtualization platform, IP/MPLS fabric, as well as any wireless infrastructure. (See Figure 3.) The scenarios related to availability could involve:

- RCA – Symptom and signature-based RCA is used to isolate issues from a multi-domain service delivery stack.
- Edge Down – The reasons could include downing of box, power, primary gateway, underlay, links or VM.
- Link or Tunnel Down – Reason could include downing of underlay MPLS or ISP, edge-side adapter, or provider side adapter.
- Link Down, 3rd Party – Symptoms would require service provider visibility of WAN equipment and faults.
- HA Impacted – High-availability of redundantly configured edge devices and gateways could be impacted by downing of HA link, one of the edge devices, one of the edge VMs, one WAN link or one tunnel.

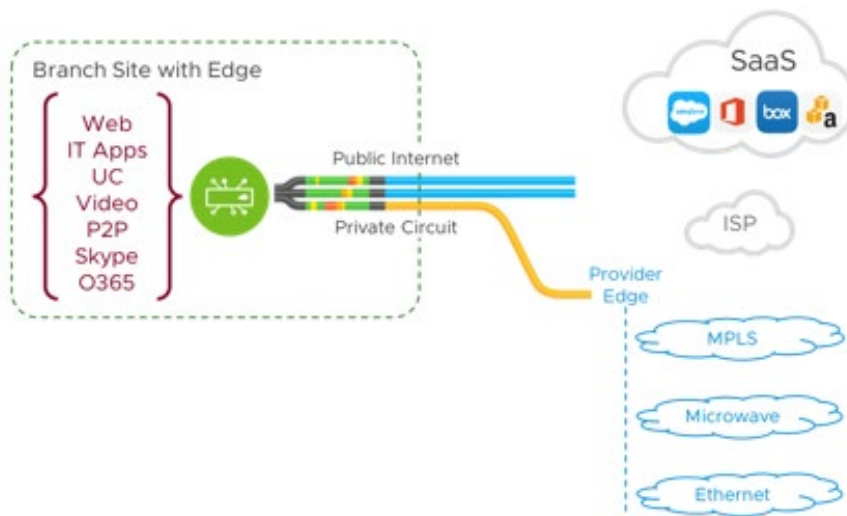


**Figure 3 - SD-WAN Branch Site with High Availability Link**

### 3.3. SLA Management

Each SD-WAN branch site has its own profile with unique services to assure. (See Figure 4.) The goal and sources with SLA management would be largely the same as with availability. The scenarios, however, are different and would include:

- SLA Violations – Reporting on SLA violations involving metrics such as customer type, throughput, downtime, number of faults by severity; and using CRM integrations and business parameters.
- Poor Availability Edges – Reporting and ranking of edge sites based on availability, criticality of downtime, maintenance downtime, and cost ratio as a function of overall weighted value.
- Prioritization – Reduction of penalty violation costs by prioritizing customer revenue value; assigning cost weights to services and infrastructure and prioritization of remediation based on higher-risk customers.

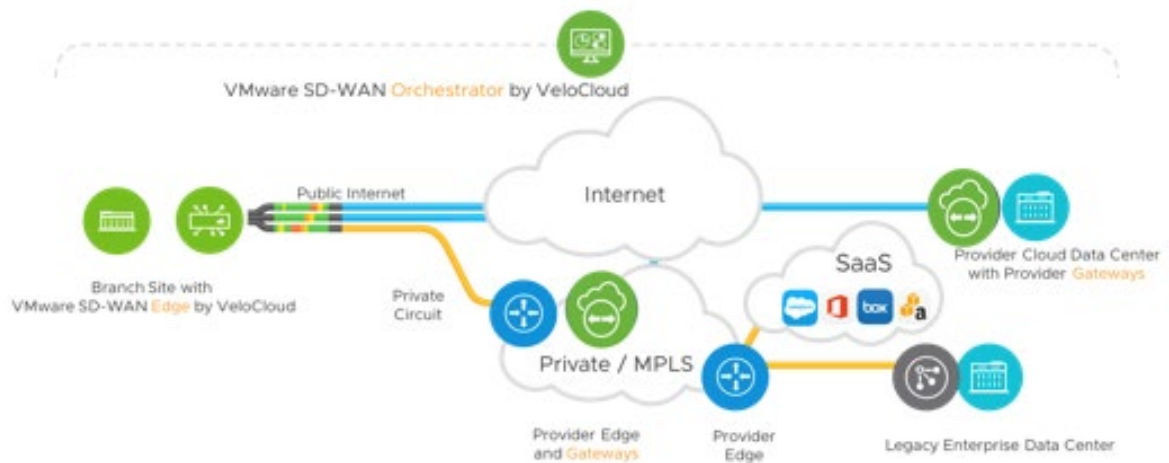


**Figure 4 - SD-WAN Branch Site with SLA-Driven Services**

### 3.4. Business Impact

The goal here is to understand the business integrity of the services offered to customers and any potential impact to revenue. Sources for making these assessments come from the orchestrator's fault data, trouble tickets, cost tables, and profiles of customer and edge device. (See Figure 5.) Ascertaining the impact would involve:

- **Prioritization** – In a multi-tenant environment with a mix of customer SLAs, the takeaway is to prioritize the remediation of impacted services.
- **Most Impacted Customers** – Identify the most impacted edges and customers to understand the operational integrity and reliability of infrastructure and services.
- **Lost Revenue** – Report on the business impact of customer and services based on standard penalty schedules, outages and severity.
- **Repair prioritization** – Prioritize network repairs and optimization based on impacted customers and their net value.
- **Edge Device Profiling** – Profile the edge devices to understand their reliability and serviceability based upon various parameters.

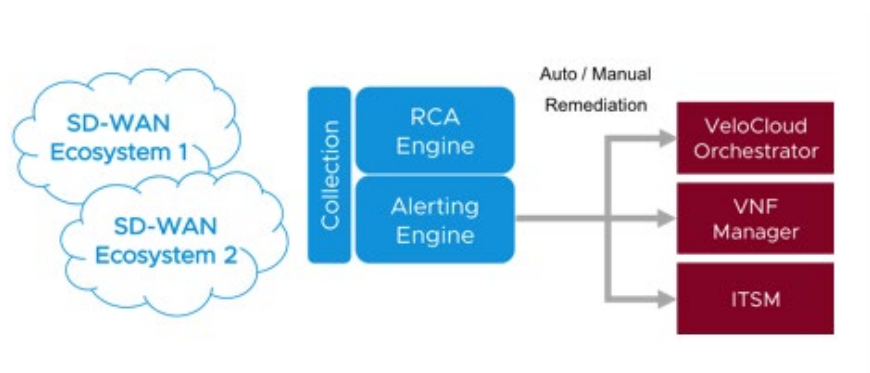


**Figure 5 - SD-WAN Branch Site, Network and Services**

### 3.5. Automation

The premise of this last use case is that the timely identification of issues, when integrated within an automated and self-optimized workflow, is crucial for maintaining SLAs and performance QoS. Sources include root-cause intelligence, policies and intent. (See Figure 6.) The scenarios involving this functionality include:

- FM alerts – RCA for fault management events are cross-correlated across application, virtual, physical, LAN and WAN transports. Alerts from the RCA signatures can be used to automate the remediation of availability issues.
- VNF Notification – Notify the VNF manager for a soft Edge deployment to carry out typical LCM remediations such as healing, scale out and scale in.
- Trouble Ticket – Integrate with an ITSM system to open trouble tickets as an automatic or manual trigger.



**Figure 6 - SD-WAN Automated Alerting**

## 4. Two Practical Examples

How do these functionalities play out in practice? Two examples, which illustrate the application of service assurance, involve triaging managed services and conducting RCA across an entire network.

### 4.1. Managed Services Outage

Let's say you're managing a plethora of vendor devices and multi-vendor SD-WANs with multiple management tools. Branch site of tenant 1 is down, and branch site of tenant 2 is down. Which branch and service do you prioritize? What is the cause of the outage? What applications are associated with it?

An optimal service assurance model would correlate the edge to authorized services and applications; and tenants to devices, virtual tunnels and services. Then assign an impact and cost score to each app/service/tenant. Finally, determine a root cause, such as a router misconfiguration, identify which customer is running which service and correlate the business impact.

The result would be an impact score that flags priority, based on tenant, applications, service and costs. Alerts are revised automatically in SA notifications, while suppressing extraneous alarms. The system would automatically trigger a trouble ticket via API integration with your ITSM. The upshot is to decide on tenant 2 and prioritize based on SLA and associated costs.

### 4.2. Cross Network RCA

In this challenge, five branches connected via SD-WAN cannot access the company database. There are multiple management tools for different services and devices. What is the cause of the outage? Switch, router, SD-WAN edge, hub or gateway, server, configuration, or application?

The optimal case is to discover and monitor complete network topology and relationships. Automated RCA pinpoints problem as a hub between five branch sites and downed data center, identifies the application/service associated with edge devices, and automatically triggers alerts. The resulting impact score in this case flags priority (based on application/service and/or VIP status) and results in a faster resolution for priority sites, improved MTTI/MTTR for faster resolution, and lower OpEx.

## 5. How to Become an SDN Provider

As the SD-WAN market matures, the need for comprehensive service assurance is becoming more clear. On the one hand is a very large number of SD-WAN solutions; on the other is interest from enterprise IT leaders in their capabilities, but strong attachment to certain aspects of legacy connectivity, such as reliability. Fortunately, service assurance tools supporting integrated monitoring of SD-WAN solutions exist and can lay the foundation for multi-vendor SD-WAN management.

MSOs who have already begun delivering SD-WAN should review their service assurance operations and capabilities. For those who have not, we recommend keeping these concerns in mind but launching SD-WAN with all deliberate haste. There are three basic models for doing so: over the top (OTT), hybrid and fully managed. The fastest is the OTT approach. In this scenario, the entire SD-WAN service is hosted by your technology provider, including physical and virtual appliances and support, employing either CapEx or OpEx models. The benefits to this approach include:

- Fast time to market to establish market presence and leadership
- Quick deployment with customers and prospective customer



- Immediate revenue stream
- Proactive enablement of technical, sales, and marketing teams proper to a full deployment

In hybrid scenario, you manage some of the infrastructure, such as gateways and edge devices, while your technology provider handles the more encompassing network orchestrator. The advantage is that the network is managed until all team members are fully trained and able to manage the system independently, thereby simplifying the transition to full integration.

In this full scenario, you assume management of all components, including support. The full-service integration allows you to truly differentiate your capabilities and layer in other services. As with some enterprise customers, service providers with certain size and capabilities might try the DIY equivalent and go the fully managed route. But for the rest, and for those looking for quick go-to-market strategies, the turnkey approach makes most sense.

## Conclusion

The strong interest in SD-WAN indicates an ongoing shift in thinking about enterprise connectivity. Gartner estimates, for instance, that by 2023, more than 90 percent of WAN edge infrastructure refresh initiatives will be based on vCPE platforms or SD-WAN software/appliances rather than traditional routers, which is a more than two-fold increase from today.<sup>7</sup>

It makes sense on multiple levels for MSOs to enter this market. For those who haven't, the turnkey OTT approach affords the quickest entry strategy. But SD-WAN, for all of the agility and efficiency that it can deliver to enterprises, will not necessarily sell itself. To avoid giving it away or turning it into a loss-leader for the additional sale of broadband connectivity, MSOs should aim high. Deploying SD-WAN along with rich service assurance functionality can help settle any doubts about reliability and protect the tremendous value that it can deliver.

---

<sup>7</sup> Magic Quadrant, *ibid.*

# Abbreviations

API	application programming interface
CAGR	compound annual growth rate
CapEx	capital expenditure
CPE	customer premises equipment
DIA	dedicated internet access
DIY	do it yourself
HA	high availability
ISP	internet service provider
ITSM	IT service management
LAN	local area network
LCM	lifecycle management
LSO	lifecycle service orchestration
LTE	long-term evolution
MEF	Metro Ethernet Forum
MPLS	multi-protocol label switching
MSO	multiple systems operator
MTTI	mean time to identify
MTTR	mean time to resolve
OpEx	operating expenditure
QoS	quality of service
SA	service assurance
SD-WAN	software-defined wide area networking
SLA	service level agreement
TVC	tunnel virtual connections
UCS	underlay connectivity service
vCPE	virtual customer premises equipment
VM	virtual machine
VNF	virtual network function
WAN	wide area network

## Bibliography & References

“US Cable Set to Reap in \$18B in Biz Services Revenue in 2018,” Jeff Baumgartner, Light Reading, Nov.11, 2018

SD-WAN Infrastructure Forecast, IDC, July 2019

Magic Quadrant for WAN Edge Infrastructure, Gartner, October 2018

Draft Standard MEF 70 (R1): SD-WAN Service Attributes and Services, May 24, 2019

“Vertical Systems: SD-WAN Not Profitable Yet,” Light Reading, Nov 29, 2018

# **A Roadmap for Virtualization in HFC Networks**

## **Use Cases and Considerations**

A Technical Paper prepared for SCTE•ISBE by

**Andrew Bender**

Global Solutions Consultants Leader,  
Telco and Media Service Providers  
VMware  
1503 LBJ Parkway, Ste. 700  
Farmers Branch, TX 75234  
+1 972 762 3399  
abender@vmware.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Background .....	3
Five Cases Involving Virtualization.....	3
1. Distributed Access Architecture (DAA) .....	3
2. Video on Demand and Network PVR.....	4
3. CDNs.....	5
4. 5G .....	6
5. MEC .....	8
Workload Placement Options.....	9
1. Regional and National Datacenters .....	9
2. Headend sites.....	9
3. Remote PHY .....	9
4. New Definition of the Edge.....	10
Framework Requirements and Common Platform .....	10
Conclusion .....	12
Abbreviations.....	13
References.....	15

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Remote PHY System .....	4
Figure 2 - Architecture for Advertising Supported VOD and Linear Video .....	5
Figure 3 - CDN Reference Model.....	6
Figure 4-a - Virtualized Radio Access Network Topology Options .....	7
Figure 5-b - 5G Functional Elements.....	7

# Introduction

The virtualization of software workloads to provide network functions is a concept that has arisen in multiple industries, including the MSO / MVPD operator community. Most initial implementations of Network Function Virtualization have followed a pattern of centralized sites that correspond to regional serving area or national data centers. But driven by the need to bring services, resources, and intelligence deeper into the HFC plant in closer proximity to subscribers the industry is promoting virtualization deeper in the access network. A leading example would be Distributed Access Architecture (DAA) or Distributed CCAP Architecture (DCA), and the associated Remote PHY Devices (RPDs). But there are other use cases and architectures driving demand for edge network intelligence including Cloud DVR, CDNs, as well as mobility driven multi-access edge computing (MEC) and 5G; all of which bear consideration for deployments as well. Going forward, a platform strategy and framework for virtualization, which anticipates multiple applications and software driven technologies spanning access and centralized datacenters will enable operators to enable new revenue streams and drive operational efficiencies across their service portfolio.

## Background

The Cable industry has considered virtualization in various forms for approximately a decade. In addition to leveraging web-based architectures to provide services for video subscribers – which generally depend upon virtualization - the industry has for several years been planning to apply virtualization to the broadband network through the Distributed Access Architecture (DAA) initiative.

The high-profile DAA initiative has now advanced into commercial implementations and deployments. At the same time, we can point to other areas where virtualization intersects existing operations, such as VOD, nPVRs and CDNs. Two additional developments originating outside the industry – 5G and Multi-access Edge Computing (MEC) – also merit attention for their potential applicability. After first reviewing these five cases (DAA, cDVR, CDN, 5G and MEC) we will share some strategic considerations about the growth of virtualization within the industry’s evolving HFC networks; and then conclude with thoughts about how a second-wave virtualization framework that extends beyond large data centers can continue to drive efficiency and reduce complexity.

## Five Cases Involving Virtualization

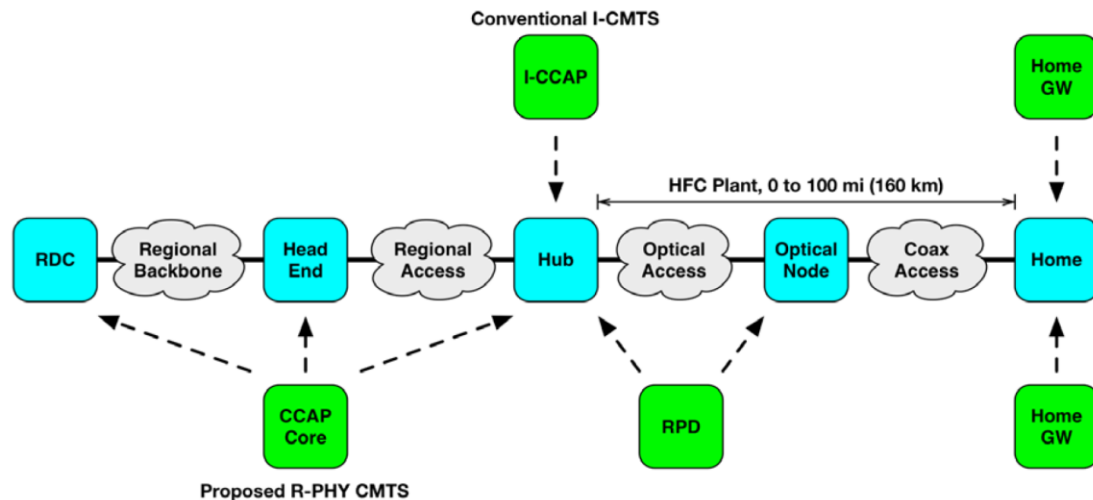
### 1. Distributed Access Architecture (DAA)

This framework has roots in the industry’s earlier modular headend architecture (MHA), which separated the Physical (PHY) downstream and DOCSIS Media Access Control (MAC) components.<sup>1</sup> (See Figure 1.) Using a new digital link extending from the CCAP core to an RPD, DAA enables the distribution and virtualization of network functions. In the new model, the CCAP core could reside at the headend or hub; and the RPD at a hub or node. Officially known by CableLabs as the Distributed CCAP Architecture (DCA), DAA aligns with other industry initiatives, including Full Duplex DOCSIS (FDX), Extended Spectrum DOCSIS (ESD), and the extension of fiber to points deeper into the network. CableLabs also associates it with higher spectral efficiency, Gigabit services, and increased access network performance

---

<sup>1</sup> Data-Over-Cable Service Interface Specifications, DCA – MHA v2, Remote PHY Specification, CM-SP-R-PHY-I12-1903307, March 7, 2019

and a much smaller footprint. By transforming the CCAP (or CMTS) from purpose-built hardware into software that could potentially run in a data center on COTS equipment in a private cloud, DAA becomes a classic case of a virtualized and software-defined infrastructure.



**Figure 1 - Remote PHY System**

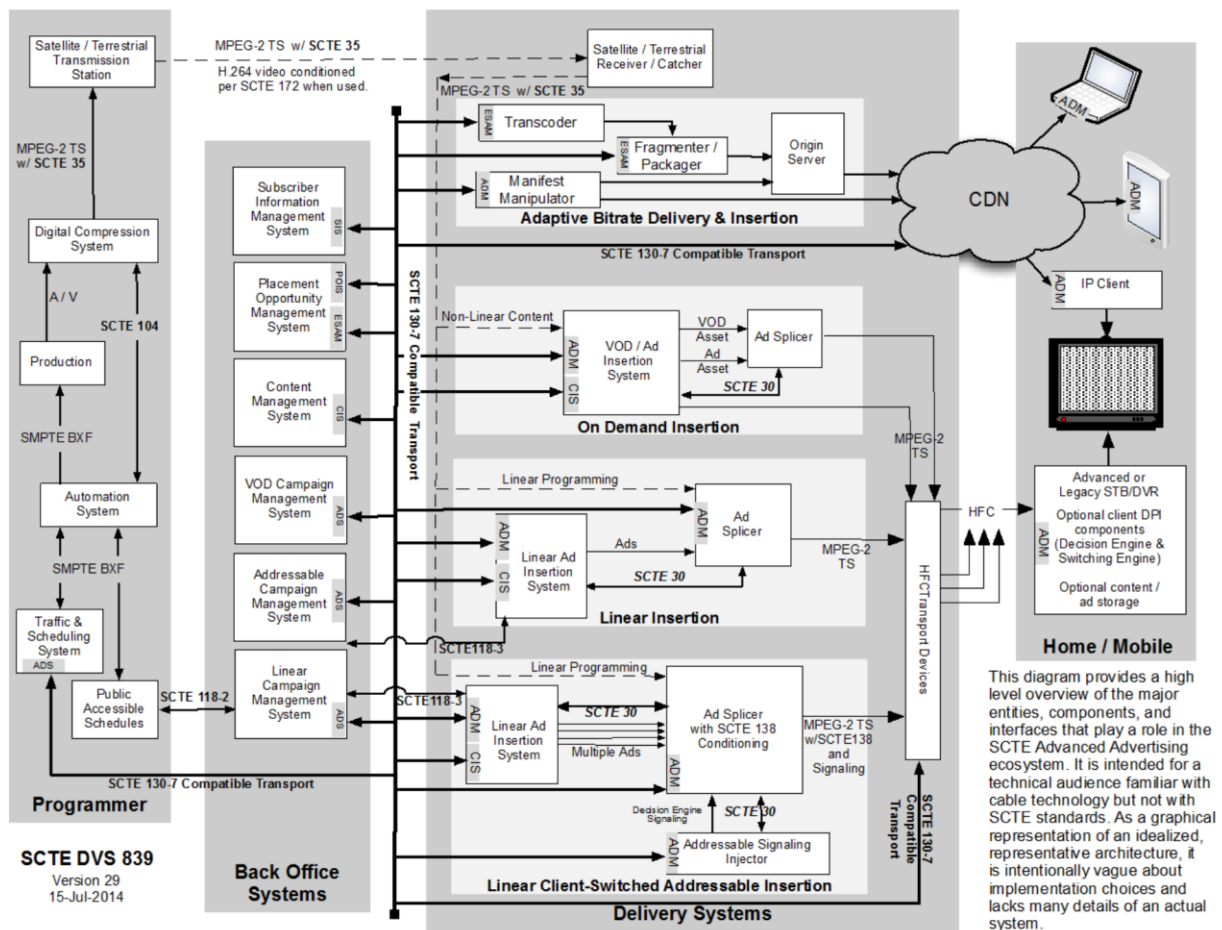
Source: CableLabs

## 2. Video on Demand and Network PVR

The architecture to deliver localized or personalized content through Digital Program Insertion (DPI) and Video on Demand (VOD) systems is generally implemented through a multitude of software-based elements, a number of which are naturally distributed to the network edge for functional reasons. See Figure 2.

Likewise, the migration of video content storage from purpose-built customer premises equipment (CPE) in the home to cloud infrastructure is another example of a service innovation enabled by virtualization. The network-based personal video recorder (or “cloud DVR”) approach facilitates simultaneous, efficient availability of private and catalog content in a multi-device and multi-network consumption model. Although control plane functions like schedulers, license servers, program guides and the like use web technology and interfaces that are readily centralized, the dataplane and network traffic requirements of origin servers, packagers, transcoders scale quickly according to subscriber demand.

Higher resolution, rate, and quality media formats now in use for streaming (UHD, HDR) and the associated codecs (H.265) call for an increased proportion of resources per active subscriber. These content types, coupled with growth of non-streaming and less predictable services also give rise to characteristically more variable compute, storage, and network demands. Thus the ability to dynamically allocate – and relocate – these resources for software workloads in response to network demand is highly desirable.

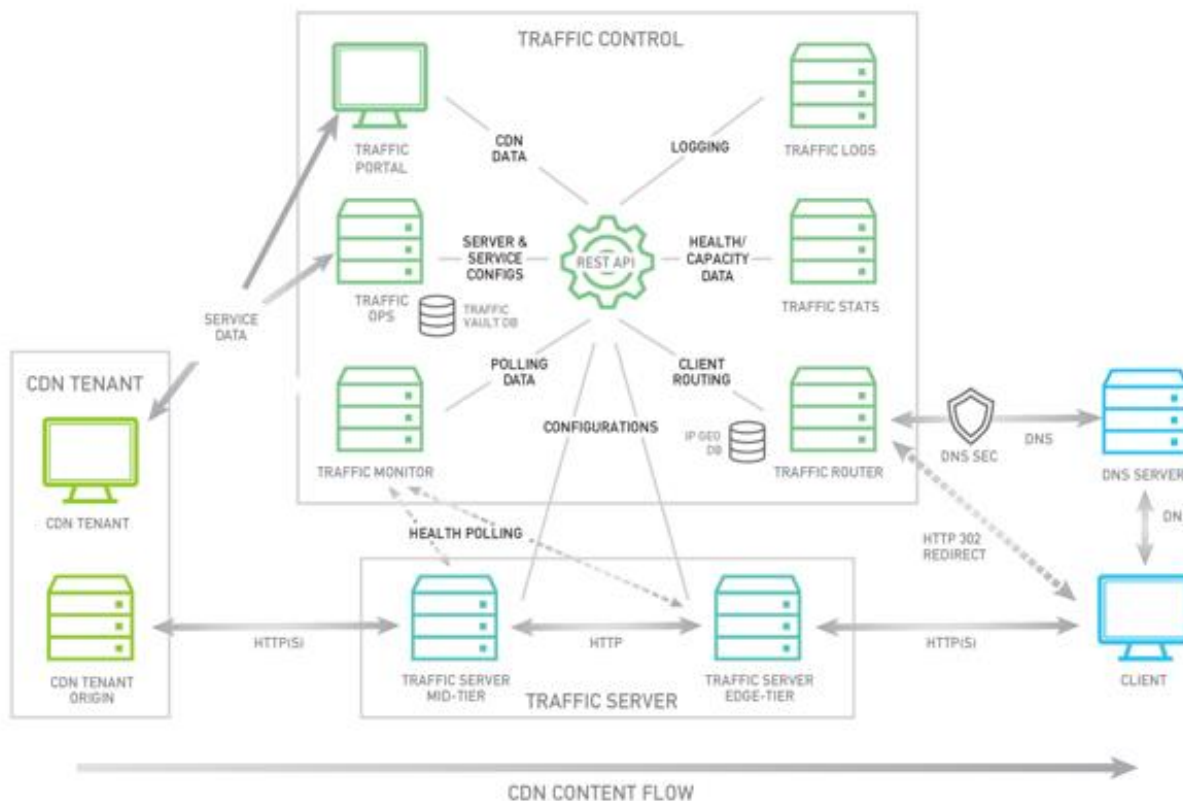


**Figure 2 - Architecture for Advertising Supported VOD and Linear Video**

Source: SCTE

### 3. CDNs

Content distribution networks (CDNs) are another application category that can leverage virtualization, and a runtime environment at the network edge. MSOs have long-established patterns of working with or deploying CDNs, which provide web, media, OTT streaming content delivery. The CDN reference model created by the Apache Software Foundation highlights a cluster of traffic control functions connected to each other and related servers on the data plane. (See Figure 3.) In a web centric application paradigm, RESTful APIs are the prevalent way for applications to interact and access resources; they also facilitate network portability and remote interconnection between these systems. However, shield cache tiers (or, “traffic servers” according to this model) are typically positioned in proximity to demand sources to benefit network efficiency and scaling. These dataplane candidates are natural candidates for distribution to the network edge.



**Figure 3 - CDN Reference Model**

Source: Apache Software Foundation

## 4. 5G

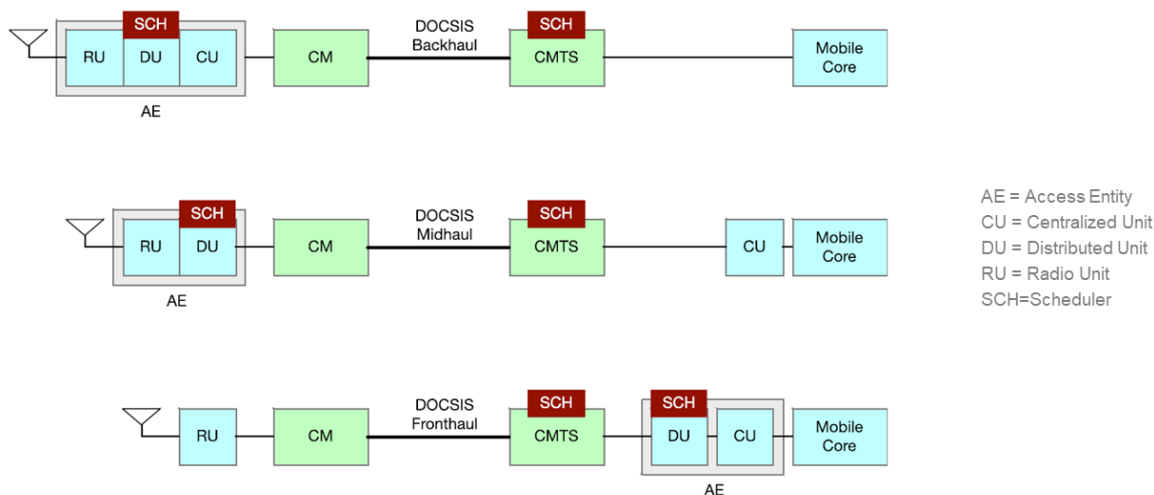
Using mobile networks and technology for service delivery within and outside the home has been a strategic focus for a number of MSOs<sup>2,3</sup>, supporting new applications and modes of consumption. The 5G initiative is the next frontier of focus for the mobile industry building upon the mobile broadband framework provided by 4G and LTE. This technology and standards regime will provide new capabilities for ultra reliable low latency communication, enhanced mobile broadband, and enhanced machine type communication for IoT, supporting both consumers and enterprise use cases. The question raised here is how, where and when to transition mobile virtual network operator (MVNO) or mobile network operator (MNO) operations to on-net solutions leveraging 5G. The new architecture associated with 5G differs from 3G and 4G, all the way down to the radio level and the new base station, or gNodeB (gNB). Given new frequency bands, air interface technology and propagation properties, as well as Radio Access Network (RAN) and network core architectures associated with 5G - architects and planners must determine where in the network footprint new user plane and radio units should be located. (See Figure 4-a.) The core and edge components of the 5G network are all expected to be IP-connected, software-defined and virtualized throughout... down to the radio baseband level. (See Figure 4-b.) So called

<sup>2</sup> <https://newsroom.charter.com/press-releases/charter-launches-spectrum-mobile-a-smarter-network-designed-for-the-future/>

<sup>3</sup> <https://corporate.comcast.com/news-information/news-feed/comcast-xfinity-mobile>

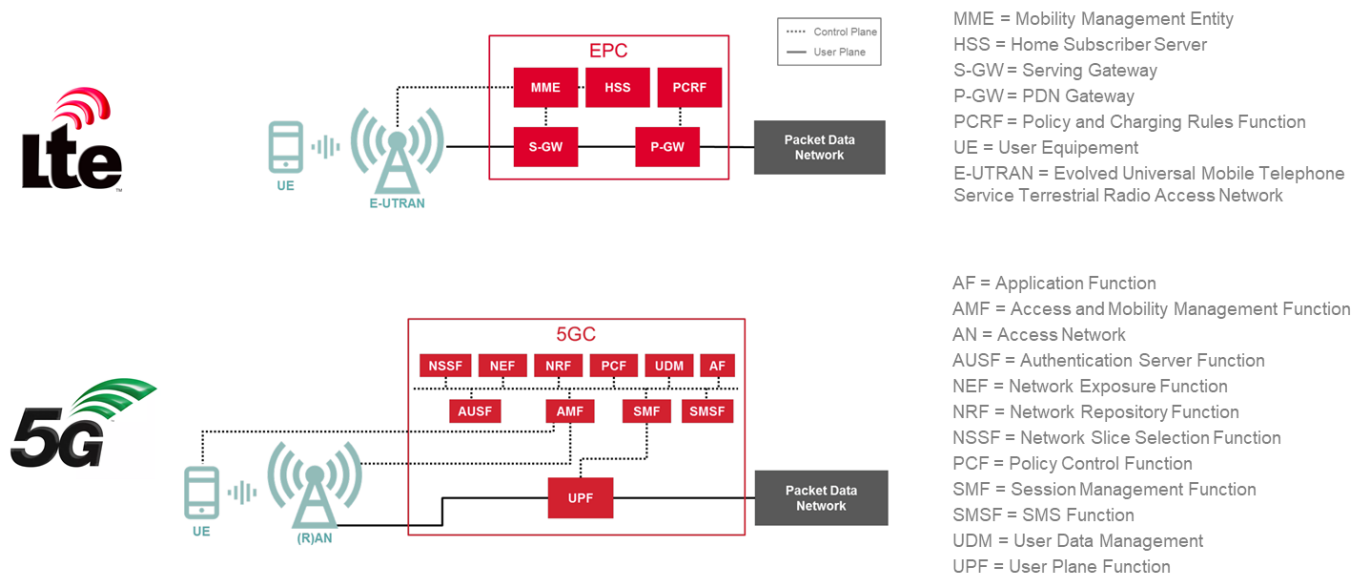


“densification” of the gNB radio units is anticipated, due to higher frequency spectrum, higher levels of demand and concurrent use – that in turn will drive radio placement choices in the home gateway, NID, street or neighborhood level for small cell gNBs versus the macrocell-heavy network footprint that typifies the 3G and 4G footprint.



**Figure 4-a - Virtualized Radio Access Network Topology Options**

Source: CableLabs



**Figure 5-b - 5G Functional Elements**

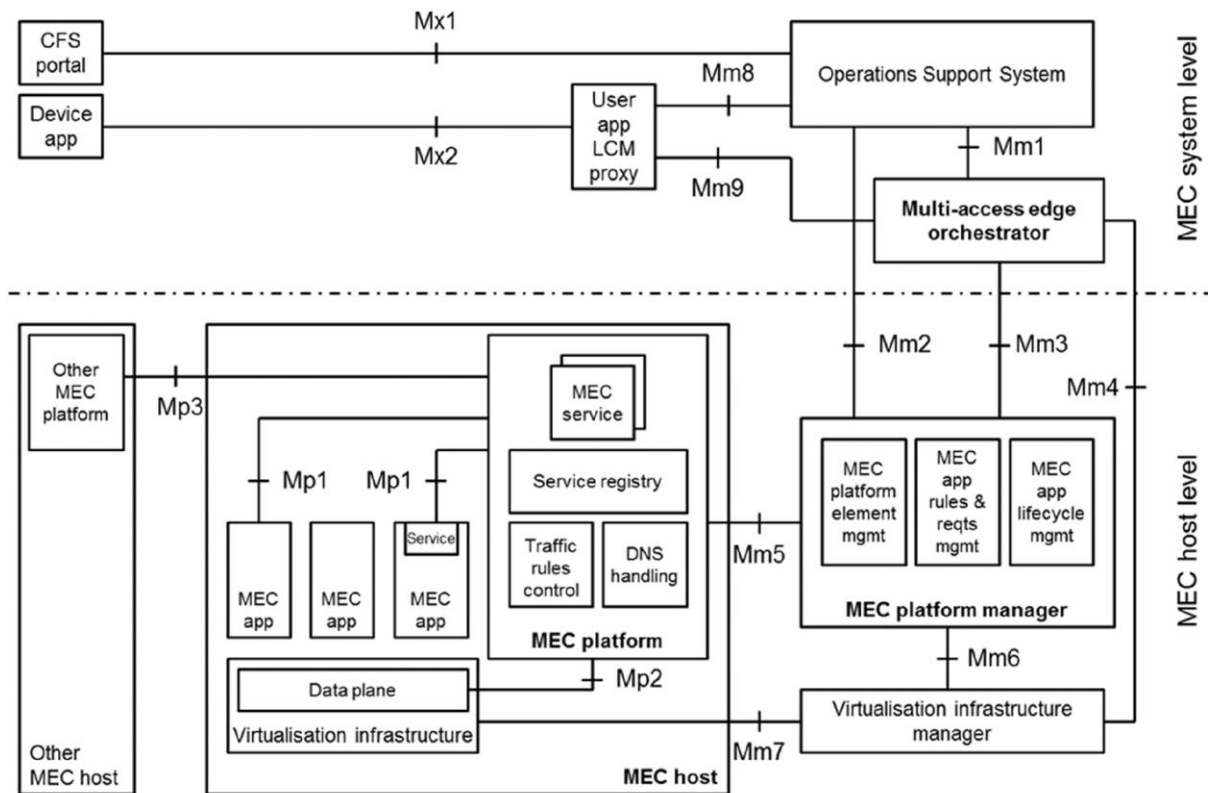
Source: GSMA

## 5. MEC

The Multi-access Edge Computing (MEC) model, established by ETSI, provides a framework for the distribution and coordination of edge computing intelligence (in the “MEC host”) to support software based network functions (“MEC applications”).<sup>4</sup> Although originally conceived in anticipation of mobile applications and use cases, MEC acknowledges the notion of a “multi-access” paradigm, using varying access and transport technologies.

A MEC platform provides Bandwidth management and prioritization, location, addresses the question of how to do edge computing and run workloads supporting mobile, but not necessarily over a mobile network.

While being developed by ETSI, this distributed computing model has implications for MSOs because it acknowledges the need for a platform strategy that pervades access technologies, and provides capabilities and services at the network edge. (See Figure 5.) This network edge could correspond with or complement the placement of Remote PHY Nodes (RPNs), Remote CCAP (MAC-PHY) elements, or Head End equipment, according to network and application requirements.



**Figure 5 - ETSI Multi-access Edge**

<sup>4</sup> Multi-access Edge Computing (MEC); Framework and Reference Architecture, ETSI GS MEC 003 V2.1.1 (2019-01)

# Workload Placement Options

In light of the use cases we have considered, engineering and planning questions naturally arise about what level within (or outside of) the HFC network is the correct one for a given functional element. In turn, if the majority of these elements are software-defined, and virtualized, a decision is required as to what framework shall be used to facilitate the orchestration, management and lifecycles of this collection of elements.

## 1. Regional and National Datacenters

The primary environment in the operator service network where virtualization has been implemented to date is the large, centralized regional (or super-regional) datacenter. Various architectures and frameworks have been conceived to support this paradigm, which include both open source community maintained and commercially supported implementations.

As the notion of software-defined network workloads is applied at scale in this context, specialized network capabilities, performance demands, and operational considerations have become issues for many deployments.

Resolutions have been brought forward for these issues, either from the community in reference implementations, or by commercial vendors. However, many of these implementations presume a deployment model that only favors a small number of discrete sites with a large, dense population of resources, and unconstrained “East - West” network bandwidth for the software-defined workloads.

Considering only this deployment type drives a model that may not be extensible in the opposite case – “sparse” network edge sites that have a low density of nodes or hosts across a large quantity of sites

## 2. Headend sites

Today’s headend sites are where edge-QAM devices, CCAPs and other legacy systems reside. Virtualized and software-defined functions are starting to be deployed at this level as well.

There are industry initiatives that seek to extend this trend to optimize and re-architect headends as datacenters for IP-based software workloads (as well as the edge sites of other operators, such as mobile network base station sites, telco central offices, and the like).

Regardless of the network type, these sites differ from the regional or national datacenter in various engineering and design parameters, including their geographical distribution, access to transport, available space, power and compute capacity. Although technical innovations continue to change the scope of these limitations, these differences ultimately determine the type and quantity of workloads that are suitable for headend site.

## 3. Remote PHY

The Remote PHY node is the new edge site for intelligent software devices defined in a DAA network. A site that might simply have been associated with an amplifier or regeneration in the past can now become a site for one or more bona fide compute elements running software.

Whether within the Remote PHY Device itself, or collocated in the enclosure or site equipment package, this can be a natural location for a small cell radio site or gNB, virtual RAN components or MEC hosts.

## 4. New Definition of the Edge

Going beyond RPN sites, there are other candidates for software defined network functions, including the last active, the NID, and even home gateways, set-top boxes, or other CPE devices – these are locations where equipment can be installed or upgraded with elements that provide compute capacity for additional software and network capabilities, and even this class of devices is now capable of supporting virtualization.

Because of the cost, environment, and various technical considerations, these devices often have less resources and capacity than network elements at the RPN, or any point upstream. However, there is still a requirement to provide management, security, software deployment, and lifecycle capabilities for these elements.

For operators who are implementing mobile, these “new edge” locations are obvious candidates for a small cell or combined device providing licensed and unlicensed radio access alongside the DOCSIS network.

With the advent of FDX, the asymmetric balance of the HFC network starts to become more upstream-oriented. This will enable new consumption patterns, as AR/VR hardware and applications gain traction, for example. This will lead to new equipment and software-defined capabilities in the “customer edge”.

# Framework Requirements and Common Platform

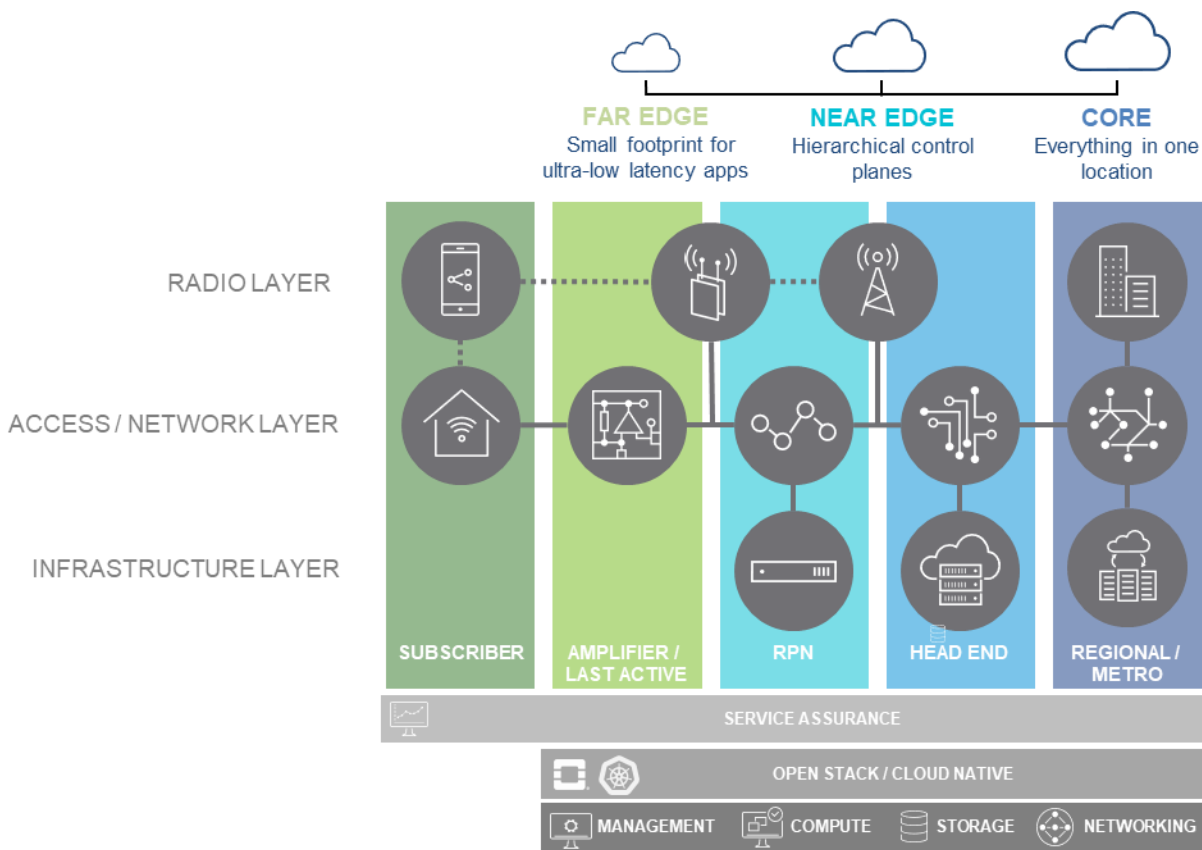
To enable a common platform for virtual functions – what we might call a “reference environment” – any such framework must provide a consistent set of capabilities:

- Disaggregated functions. Services are highly disaggregated so that control, data, and management planes can be deployed across the distributed topology. Edge clouds offer the performance advantages of low latency and data plane intensive workloads. While control and management plane components can be centralized with a regional and global scope.
- Functional isolation. Network slicing provides network and service isolation across different tenancy models in the reference environment. However, resource management considerations need to be made for shared network functions such as DNS, policy, authentication, and so on.
- Data intensive workload acceleration. The demand for throughput has increased exponentially with smart devices and immersive media services. Networking and compute expenditures continue to grow to meet traffic throughput demands. Support for acceleration technologies like DPDK, VPP, and hardware offload are required to make virtualization of data intensive applications feasible
- Cloud-native and hybrid form factor execution environments. Cloud-native approaches are dictating a new CI/CD paradigm and micro services application architectures. Container technology is a new lightweight execution environment option for delivery of these applications. While the fine-grained abstraction of applications might be a good fit for control plane functions in the reference environment, user plane functions may be required to execute as native VM

functions. This requires a cloud infrastructure environment to be heterogeneous enabling such hybrid execution environments for native VM and containerized applications.

- Federation options. The reference environment must provide a diverse set of federation options for end-points, private and public clouds, each with distinct ownership and management domains. Virtualized end-points provide better control and manageability, however they are not suitable for all types of use cases. Likewise, service functions need to be distributed and managed across private and public clouds.
- Service placement. The highly distributed topology allows for flexibility in the workload placement. Making decisions based on proximity, locality, latency, analytical intelligence, and other criteria are critical to enable an intent-based placement model.
- Workload life cycle management. Each cloud is elastic with workload mobility and how applications are deployed, executed, and scaled. An integrated operations management solution can enable an efficient life cycle management to ensure service delivery and QoS.
- Platform lifecycle management. The platform must be patched and upgraded by using optimized change management approaches for zero to minimal downtime.
- Carrier grade characteristics. Because Communications Service Providers (CSPs) deliver services that are often regulated, carrier grade aspects of these services, such as high availability and deterministic performance are also important.

The solution then (as shown in Figure 6) must be a multi- tiered hierarchical platform capable of addressing the requirements and workload types at each level within the service provider cloud – at the regional or national datacenter, the headend or “near edge”, and emerging “far edge” as well.



**Figure 6 - A Distributed Architecture for the CSP Core / Edge / Access Network**

Source: VMware

## Conclusion

Specific applications are what drove many initial virtualization deployments, sometimes for the use cases discussed here. That led, in turn, to these clouds being customized, tuned, or optimized in unique ways for specific workloads. When there are multiple instances of these clouds, and each is a bespoke environment with a diverging architecture the opportunity to realize a common platform across these applications is lost.

Whenever the architecture of a network changes, or a new cloud is implemented, a key consideration also becomes the visibility, operational tools, troubleshooting and service assurance framework that enables the environment to be managed. Each instance or cloud then requires a solution set for these capabilities which then makes the associated operational practices and support systems potentially different as well.

A fundamental reason for the drive toward virtualization and a common platform for network functions is 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.

# Abbreviations

5GC	5G core network
AE	access entity
AF	application function
AMF	access and mobility management function
API	application programming interface
AR	augmented reality
AUSF	authentication server function
CBRS	citizens broadband radio service
CCAP	converged cable access platform
CD	continuous development
CDN	content distribution network
CFS	customer facing service
CI	continuous integration
CMTS	cable modem termination system
CU	centralized unit
CPE	customer premises equipment
DAA	distributed access architecture
DCA	distributed CCAP architecture
DOCSIS	data-over-cable service interface specification
DNS	domain name system
DNSSEC	domain name system security extensions
DVR	digital video recorder
DU	distributed unit
EIR	equipment identity register
ESD	extended spectrum DOCSIS
ETSI	European Telecommunications Standards Institute
E-UTRAN	evolved UMTS terrestrial radio access network
FDX	full-duplex DOCSIS
FMA	flexible MAC architecture
HFC	hybrid fiber-coax
HSS	home subscriber server
IoT	internet of things
ISBE	International Society of Broadband Experts
LCM	lifecycle management
LTE	long-term evolution
MAC	media access control layer
MBR	maximum bit rate
MEC	multi-access edge computing
MHA	modular headend architecture
MME	mobility management entity
MNO	mobile network operator
MSO	multiple systems operator
MVNO	mobile virtual network operator
NEF	network exposure function
NID	network interface device
NFV	network functions virtualization

NG-RAN	next-generation radio access network
NRF	network repository function
NSSF	network slice selection function
OTT	over the top
OVP	online video platform
PCF	policy control function
PCRF	policy charging and rules function
PHY	physical layer
P-GW	packet data network gateway
QAM	quadrature amplitude modulation
RCA	root-cause analysis
RDK	reference design kit
REST	representational state transfer
RF	radio frequency
RPD	remote PHY device
RPN	remote PNY node
RS-DVR	remote storage DVR
RU	radio unit
SCH	scheduler
SCTE	Society of Cable Telecommunications Engineers
SMF	session management function
UDR	unified data repository
UDSF	unstructured data storage function
UDM	user data management
UE	user equipment
UMTS	universal mobile telephone service
UPF	user plane function
VIM	virtual infrastructure manager
VR	virtual reality



## References

- 802.3, I. (2018, May). *Beyond 10km Adopted Objectives*, Study Group. Retrieved from [http://ieee802.org/3/B10K/project\\_docs/objectives\\_180521.pdf](http://ieee802.org/3/B10K/project_docs/objectives_180521.pdf)
- CableLabs. (2018, June 29). *P2P Coherent Optics Physical Layer 1.0 Specification*. Retrieved from <https://apps.cablelabs.com/specification/P2PCO-SP-PHYv1.0>
- Forum, O. I. (2018). *Current Work done at OIF*. Retrieved from <http://www.oiforum.com/technical-work/current-oif-work/>
- G.694.1, I. (n.d.). *Spectral grids for WDM applications: DWDM frequency grid*. Retrieved from <https://www.itu.int/rec/T-REC-G.694.1-201202-I/en>
- Microsemi. (2017, March). *Microsemi Enables Terabit OTN Switching Cards for Flexible Optical Networks*. Retrieved from <https://www.prnewswire.com/news-releases/microsemi-enables-terabit-otn-switching-cards-for-flexible-optical-networks-300608509.html>
- OpenRoadm. (2018). *Open Roadm MSA*.

# The Promise of WiFi in the 6 GHz Band

A Technical Paper prepared for SCTE•ISBE by

**J.R. Flesch**

Director, Advanced Technology  
Commscope  
3871 Lakefield Drive, Suwanee, GA 30024  
678 473 8340  
Jr.flesch@commscope.com

**Charles Cheevers**

CTO/CPE  
Commscope  
3871 Lakefield Drive, Suwanee, GA 30024  
678 473 8507  
Charles.cheevers@commscope.com

**Kurt Lumbatis**, Commscope

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	3
The Exciting Promise of Indoor 6 GHz .....	3
1. Motivation for Exploit of 6 GHz as Unlicensed Spectrum Relief.....	3
1.1. Overcrowding in legacy unlicensed spectrum.....	3
1.2. MDU and near-neighbor ingress .....	4
1.3. The insatiable demand .....	4
1.4. Coverage gaps of the WAN gateway .....	6
2. Application of the 6 GHz Remedy.....	7
2.1. Prospective performance .....	7
2.2. Why low power? .....	15
2.3. Challenging the WAN capacity.....	17
3. Tinkering with expected performance in MDU environments.....	20
3.1. Napkin musings on the scope of the ask .....	20
3.2. Crunching some numbers.....	21
3.3. The Power of BSS Coloring in MDUs.....	23
3.4. External Interference by Inside WiFi at 6 GHz.....	24
Conclusion .....	25
Abbreviations.....	25
Bibliography & References .....	27

## List of Figures

Title	Page Number
Figure 1 – Spectral Crowding @ 2.4 and 5 GHz for MDU cases.....	4
Figure 2 – IP traffic demand expectations .....	5
Figure 3 – Consumer adoption of wireless IoT devices by year (blue component) .....	6
Figure 4 – WiFi test house, front view .....	8
Figure 5 – WiFi test house, rear view .....	9
Figure 6 – WiFi test house, top level floorplan with test cases .....	10
Figure 7 – WiFi test house, main/mid level floorplan with test cases.....	11
Figure 8 – WiFi test house, basement level floorplan with test cases.....	12
Figure 9 – TCP bitrate performance across six test cases in WiFi house .....	13
Figure 10 – 4 SS, 80 MHz UDP bitrate service radius at 5 frequencies in the 6 GHz band, 250 mW .....	14
Figure 11 – 4 SS, 80 MHz UDP bitrate service radius at 5 frequencies in the 6 GHz band, 1W.....	15
Figure 12 – Pending FCC NPRM showing consideration of non-AFC low power bands .....	16
Figure 13 – 250 mW WiFi6 UDP bitrate curve @ U-NII-5 & -8 versus 1W WiFi5 @ U-NII-3 .....	17
Figure 14 – Model of 6 GHz / 160 MHz BW in-home backbone trunk and service mesh .....	18
Figure 15 – Expected Performance of the exercise model.....	19
Figure 16 – Floor plan of “typical” 900 square foot apartment, AP location in red.....	21
Figure 17 – One floor of example MDU (6 units) showing CCI peak location .....	22

# Introduction

The 5.925-7.125 GHz band (colloquially “6 GHz band”) represents an immense opportunity for indoor WiFi to fully adopt the promise of WiFi6 in a green space environment and clear out the channel access baggage and heterogeneous technical epoch mix accumulated during the more or less organic growth of unlicensed, contention-based wireless services in the 2.4 and 5 GHz bands. In exploiting this clean break, it avoids disrupting the existing population of devices and their present state of interoperability (however suboptimal that may be). Spectrum leverages associated with multiple-user OFDMA, multiple-user MIMO and BSS coloring have the ability to promote low-latency spectrum scheduling, improved link margins and topographical channel re-use which will go a long way towards resolving the potentially thorny CCI environment represented by multiple-AP, dense client device deployments. Additionally, the wealth of new spectral piping available at 6 GHz, of itself, may provide all the solution required to wirelessly backbone data hauling in the home between WAN attachment and an opportunistically-placed AP/hub/extender, resulting in reliable (virtually OOB) trunking hauls between access points which enable whole-home LAN bitrates sufficient to meet anticipated WAN bulk connectivity budgets as these inflate via either DOCSIS or 5G mechanisms. Blanket WiFi coverage of multi-Gbps (as a services ensemble) ought to be achievable given the power, BW, link budget and spectral efficiencies available within the service radii posed by indoor residential environments.

## The Exciting Promise of Indoor 6 GHz

### 1. Motivation for Exploit of 6 GHz as Unlicensed Spectrum Relief

#### 1.1. Overcrowding in legacy unlicensed spectrum

The following figure details the spectral occupation associated with existing ISM and U-NII bands at 2.4 and 5 GHz which host 802.11-based wireless traffic across at least three technical specification epochs of that standard. MAC differentials across these epochs contribute to the access pathology by collapsing throughput to least-common-denominator type of wireless medium exploit in cases where heterogeneous client populations comprising some number of older legacy devices compete for airtime from the access point. Even in the cases where relatively high bitrate streaming traffic is shunted off of the 2.4 GHz to the 5 GHz band, crowding in both pieces of spectrum is becoming everyday more commonplace.



**Figure 1 – Spectral Crowding @ 2.4 and 5 GHz for MDU cases**


And the problem is exacerbated by IoT radios exploiting the 2.4 GHz ISM band to connect constrained end devices (potentially yielding status-critical telemetry in small packet traffic) to IoT mesh hubs (as discrete CPE or WiFi extender adjunct stackware) for ultimate backhaul over the 802.11 network. Home security and aging in place services represent two such IoT applications which can ill afford excessive attachment latencies (or worse – lost data). The upshot is that offset carriage spectrum needs to be mined in order to free up IoT access (for those services’ several NFC MACs operating at 2.4 GHz) by moving 802.11 traffic away from that highly contentious band.

## 1.2. MDU and near-neighbor ingress

Allowable legacy wireless power levels conspire to recruit ingressing, unwanted interferers in the case of near neighbors – or even more problematically, MDU structures. While clever amelioration techniques like EasyMesh can identify problem channel competition and provision better-case utilization of available spectrum, such techniques are rendered less effective by overlays of wireless networks representing disparate control authorities (whose closed circuit loop dynamics can conspire to orchestrate chaotic thrash in the mixed environment). Migration to more common adoption of these type of higher stack layer controls will help – but not as much as reserving pristine new spectrum and reserving it for exploit by devices compliant to only the most recent MAC initiatives implemented in WiFi6.

## 1.3. The insatiable demand

To compound the spectral crowding, in-home wireless bitrate appetite is only increasing. Note the implacable demand for ever more device connectivity expected in the immediate future, as witnessed by the accompanying figure:



**Figure 2 – IP traffic demand expectations**

And this is for native 802.11 traffic. As alluded to above, IoT home device use is on the rise as well and this will inject further traffic into the spectrum, particularly at 2.4 GHz (and with Zigbee, Bluetooth and Thread MAC behavior which may not be at all WiFi-aware):

## The Internet of Things (IoT) Units Installed Base By Category 2014 to 2020 (in billions of units)




Figure 3 – Consumer adoption of wireless IoT devices by year (blue component)

### 1.4. Coverage gaps of the WAN gateway

Multilevel single family dwellings represent a single AP wireless coverage challenge once floorplans exceed 2000 square feet or so; this is a function of service radius (set by losses to both in-air link endpoint distances and accumulated solid surface transitions of walls and floors), increasing number of client devices competing for shared air time and the common restriction that wireline WAN services are usually introduced to the home at an external wall. This aperture typically wastes half of the isotropic radiated antenna pattern of the gateway located just at the interior feed point of the WAN and requires a service radius extender somewhere toward the middle of the home's interior to support wireless backhaul (or the co-option of existing – or the pulling of new – wireline to provide front- and backhaul support to the extender). Absent the extender, whole sections of the home furthest from the gateway would be effectively blanked from network attachment (certainly from bitrate service which could support multiple large-format video streaming applications, say).

Wireline trunk upgrades to the home, however, are not the stuff of self-installation and have the capacity to generate significant consumer ill will – the common feeling being that the ISP has failed in its duty to provide wireless coverage as was expected. And home infrastructure upgrade/maintenance costs are seldom properly anticipated and never tolerated well. MoCA, Ethernet and powerline co-option

techniques do offer the promise of less-interference-susceptible trunking to the extender – but not necessarily to carrier-grade reliability levels (or without – excepting Ethernet – the use of two mirror end-of-link devices to transceive and transcode the wireline modulation scheme employed).

Meanwhile, wireless trunking to the extender to/from the WAN GW, while facilitating the buy/self-install services remediation paradigm, suffers from “small print description” reliability issues (i.e., extender type may not represent the correct solution to the given coverage problem). This is because the extender can problematically cannibalize channel bandwidth and/or airtime from client device usage for applications in order to sustain the ultimate WAN connectivity across one or more hops of the whole in-home network. As such, extenders with insufficient bandwidth or older MAC technology might actually worsen the whole-home service experience. Triband extenders (meaning 2.4/5Lo/5Hi) and WiFi5 bring some interim relief for this proposition (by differential exploit of the 5Lo and 5Hi bands for front- and backhaul, say), but absent the WiFi6/6 GHz proposition cannot endow the wireless home with the ensemble bitrate support needed to meet coming wireline and 5G fixed wireless WAN speeds.

## **2. Application of the 6 GHz Remedy**

### **2.1. Prospective performance**

As mentioned above, if just a reliable high-capacity wireless trunk over new spectrum can be placed between GW and extender, both legacy and new WiFi devices will experience additional connection capacity (the former via the recovery of spectrum lost to TriBand or other trunking and the latter, via exploit of the newly available BW). In initial deployments, the WiFi6 MAC upgrades associated with MU-MIMO, MU-OFDMA and BSS coloring will not even need to be invoked to see immediate, massive performance gains; even at low power (250 mW EIRP), the constrained dimensions of an average indoor floorplan (~2600 sq ft) suggest accumulated link losses should not force the negotiated MCS to drop below midgrade levels at worst case. The upshot of this observation is that spectral efficiency should remain high over relatively massive WiFi bandwidths (up to 160 MHz – or more if the FCC reserves additional U-NII bands for low power unlicensed service).

To evaluate these expectations for indoor WiFi coverage utilizing the 6 GHz spectrum, the Arris/Commscope WiFi test house (a trilevel “average” US home) was instrumented. Perspective views and floorplans follow:

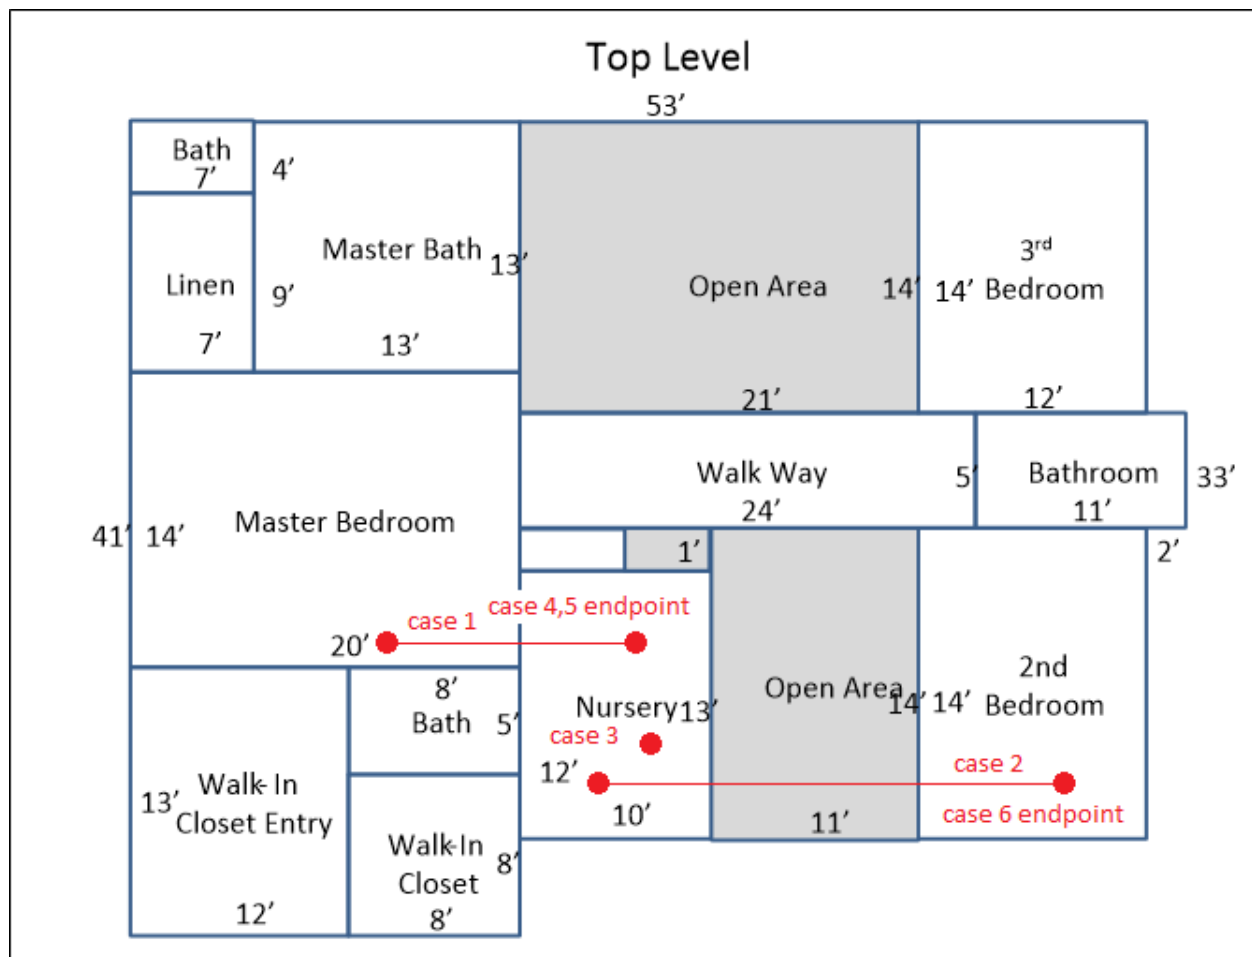




**Figure 4 – WiFi test house, front view**



**Figure 5 – WiFi test house, rear view**



**Figure 6 – WiFi test house, top level floorplan with test cases**

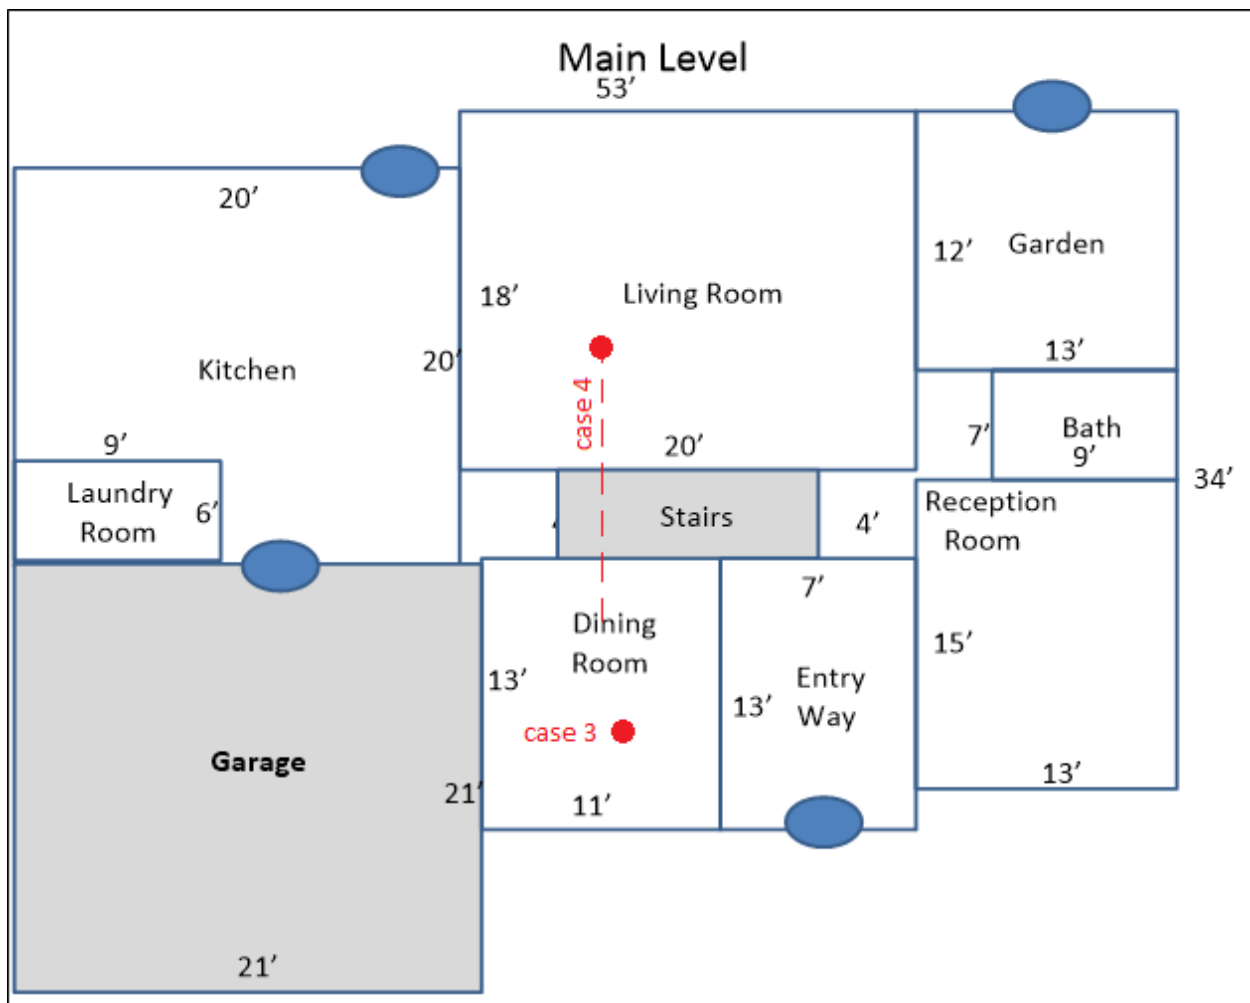
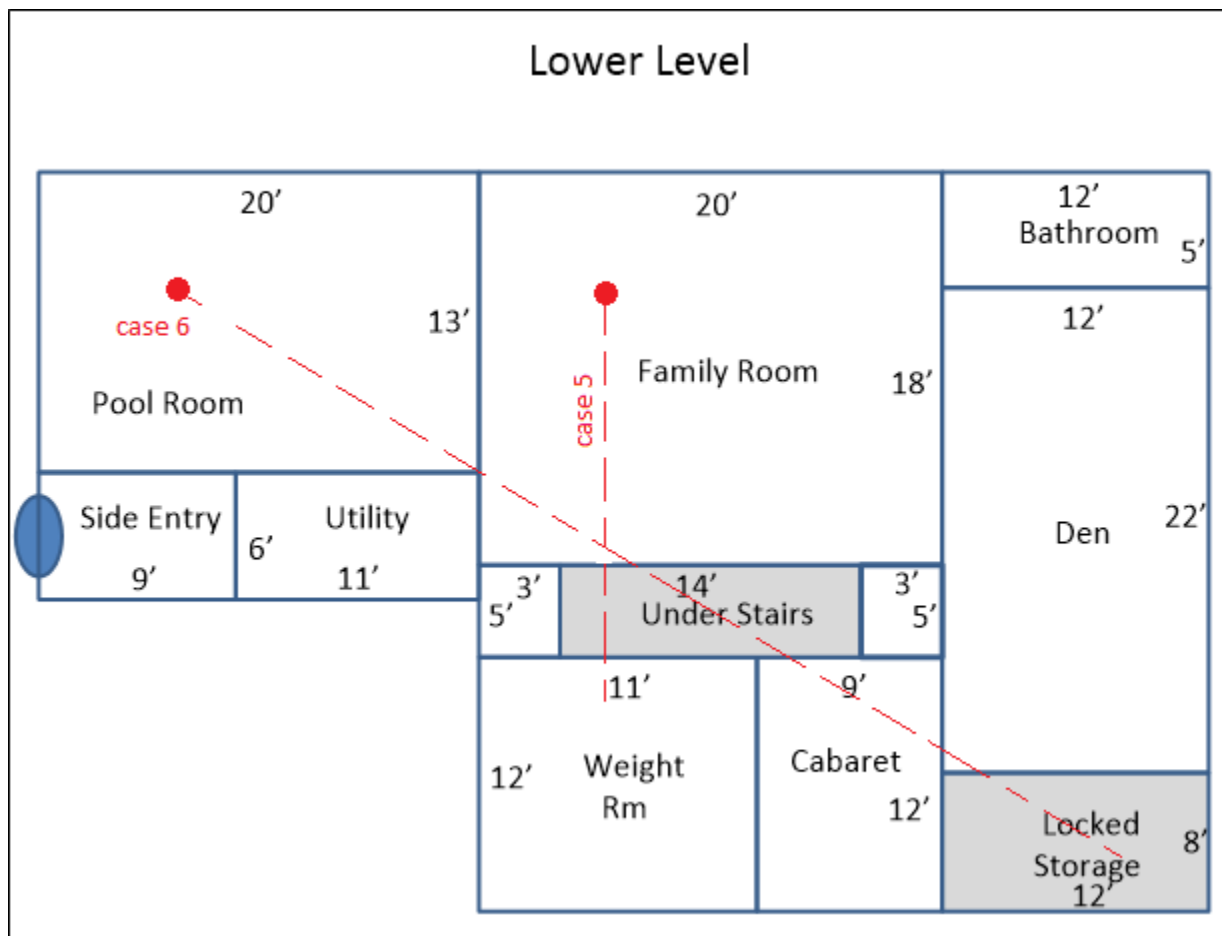


Figure 7 – WiFi test house, main/mid level floorplan with test cases





**Figure 8 – WiFi test house, basement level floorplan with test cases**

To validate expectations for the trunk's performance at low power, this average home's coverage map (as received bitrate performance) was measured across multiple links spanning near-room distances to several lumped element wall/flooring transitions at distances up to 60 feet. Six test cases, each at 100 mW, 200 mW and 400 mW total power and 80 MHz of bonded channel bandwidth were conducted. The link tests were performed with available WiFi5 endpoints set for channel 153 in the U-NII-3 band (so that through-air and lumped transition losses would mimic – albeit slightly optimistically -- conditions in the 6 GHz band). The test results using a standard iperf3 reference (scripted to include TCP messaging overhead) between two 4-chain devices yielded the following results (note that the rates were inclusive of distinctive device radio behavior around MCS selection and AGC setpoints):

TCP Bitrate (Mbps) @ Power (mW)			
	400	200	100
1 (13' + 1 wall)	911	885	856
2 (24' + 2 walls)	850	813	817
3 (11' + 1 floor)	867	835	876
4 (25' + 1 floor)	834	798	776
5 (30' + 2 floors)	575	559	460
6 (60' + 2 floors + ~2 walls)	243	159	118

**Figure 9 – TCP bitrate performance across six test cases in WiFi house**


One cautionary observation: the bitrate best-case asymptote of 911 Mbps, being TCP, can be viewed as just over 1 Gbps UDP (accounting for approximately 10% TCP signaling overhead). This is shy, however, of the expected 1.4 Gbps UDP rate expected at best MCS for 4SS WiFi 5 devices (refer to Figure 13) and appears to be an implementation artifact of the devices used in testing. (Such was confirmed by rescripting the channel BW for 40 MHz and observing a drop in bitrate by a factor of exactly 2.) But it is encouraging that 400 mW manifests the ability to light up a client device to over 200 Mbps service at 80 MHz of channel bonding (as pointed out, measured with nearly 30% implementation overhead and at TCP) across the longest (two-floor breaching) diagonal reach in the study.

Cases 2 and 4 suggest that the loss through flooring approximates that of 2 interior drywall transitions (1 floor  $\sim$  2 walls). Regressing the measured data in case 5 against the analytical expectations in Figure 13 (below) to produce the floor transition loss goes as follows: 1) Free-air path loss of 5.765 GHz at 30 feet amounts to  $\sim$  67 dB; 2) Figure 13 references loss at one meter (48dB) so move along the x-axis to 67-48 or 19 dB; 3) the reference curve assumes 1W of power but case 5/400 mW means we move an additional 4 dB to the right (to account for the lower power of the test) – so the operating point in free air would be here at 23 dB path loss; 4) the test case showed TCP performance at 575 Mbps with the hardware used so now we adjust for that implementation: multiply the TCP rate by 1.1 to get UDP, then by 1.4 to overcome the implementation loss in the HW used, so  $575 \times 1.1 \times 1.4 = 886$  Mbps; 5) Traverse the bitrate curve in Figure 13 down to the point where it indicates  $\sim$ 890 Mbps or so and drop down vertically to read the operational path loss ( $\sim$ 41 dB); 6) the operational path loss minus the free air path loss indicates what the


transition losses are, so  $41 - 23 = 18$  dB. Since this is accounted for by 2 floor transitions, this yields 9 dB for a floor loss. Floor loss being approximately twice the wall loss (in dB), this implies 4.5 dB for each drywall. Both of these estimates compare well with the extant literature on 2.4 GHz and 5 GHz indoor material transmission losses (typically in the 3-4 dB range for drywall and 6-10 dB across the bands from 2.4 out toward 6 GHz).

A USC “common material building loss” study (refer to reference section) from 2002 provides some additional frequency dependent data across the entire 2.4 – 7 GHz region. In broad summary, interior drywall transmission losses in that study averaged a fairly constant value over the bands in question and typical wooden beam and plywood flooring transmission loss seemed to monotonically increase from around 5 dB at 2.4 GHz to nearly 10 dB at 7 GHz.

To complement the measured data (and provide some calculus for expectations of bitrate performance that 6 GHz and similar power should be much better than these measured results), an analytical expectation for free air service radius of 4-chain WiFi6 endpoints operating in similar fashion as the test cases shows the following UDP bitrate expectations across increasing link losses at two power points:



**Figure 10 – 4 SS, 80 MHz UDP bitrate service radius at 5 frequencies in the 6 GHz band, 250 mW**



**Figure 11 – 4 SS, 80 MHz UDP bitrate service radius at 5 frequencies in the 6 GHz band,1W**

## 2.2. Why low power?

With legacy 802.11 indoor radiated power limits permitting more robust levels, it is reasonable to wonder why one might be interested at indoor performance achievable with a very modest 250 mW EIRP for the 6 GHz band. The answer lies in good neighbor coexistence (especially with extant outdoor 6 GHz infrastructure coupled with a desire to be conservatively biased with respect to interference) and the potential for the FCC to partition up the 6 GHz band into “standard-power” (presumably outdoor or outdoor/indoor uses) and “low-power” (indoor) bands in the first place. Of keen interest at the designated low power bands is the ability for indoor CPE to jettison intervention from cloud (or edge) -based interference arbitration schemes associated with the Automated Frequency Coordination (AFC) function (itself bearing some similarity to the CBRS band’s Spectrum Access System -- SAS). Refer to the NPRM proposal under consideration at the time of this writing:




## 6 GHz Unlicensed Device Classes | NPRM Proposal

Band (MHz)	Primary Allocations	U-NII	Devices	Max Power	AFC
5.925-6.425	Fixed Service FSS	U-NII-5	Standard-Power AP	4W (36 dBm) 30 dbm/6 dBi ant gain (U-NII-1 & 3)	Yes
6.425-6.525	Mobile Service FSS	U-NII-6	Low-Power AP (indoor)	1W (30 dBm) 24 dbm/6 dBi ant gain (U-NII-2a)	No
6.525-6.875	Fixed Service FSS	U-NII-7	Standard-Power AP	4W (36 dBm) 30 dbm/6 dBi ant gain (U-NII-1 & 3)	Yes
6.875-7.125	Fixed Service Mobile Service FSS*	U-NII-8	Low-Power AP (indoor)	1W (30 dBm) 24 dbm/6 dBi ant gain (U-NII-2a)	No

\* There is no FSS allocation in the 7.075-7.125 GHz portion of the band.

**Figure 12 – Pending FCC NPRM showing consideration of non-AFC low power bands**

If one could show impressive indoor service reach with power at (or better still, 6 dB below) FCC considerations for indoor use, it follows that the indoor application of 6 GHz to augment present-day 2.4 and 5 GHz bands should slot in with little concern for macro-scale microwave infrastructure interference (and this, without resort of WiFi6's impressive downstream – and with Wave 2 devices, upstream – spatial directivity represented by MU-MIMO). Switching consideration to MDU structures, the general rule of thumb that “the minimum necessary power required to sustain link throughput is the power level at which one should operate” does its part to minimize housing unit-to-unit interference potentials. (More about this in a subsequent section.) To illustrate WiFi6's inherent throughput advantages over WiFi5 (and perhaps dispel some entrenched legacy concerns regarding how much power is necessary), consider the following chart showing WiFi6 performance radius at 250mW versus WiFi5 at 1W:



**Figure 13 – 250 mW WiFi6 UDP bitrate curve @ U-NII-5 & -8 versus 1W WiFi5 @ U-NII-3**

As can be seen, the available receiver sensitivities and ability to support denser spectral modulation schemes of the WiFi6 links imply that for expected indoor AP spacing and coverage, the 6 dB lower power WiFi6 links do as well – and much better, at close range -- as the higher power WiFi5 reference link. The crossover point shown in the chart occurs at nearly 1.4 Gbps UDP service and at a total path loss of 77 dB or so (representing a throw of 85 feet, absent any floor or wall transition losses). So clearly, indoor RF considerations for WiFi6 need not be inclusive of 1W power levels – which, aside from the interference concerns previously noted, bodes well for HW implementation considerations for overall device dissipation (and ultimately, cost).

### 2.3. Challenging the WAN capacity

A high-bitrate demand scenario with mixed-epoch WiFi clients was crafted to illustrate the raw new capacity represented by only partial exploit of the 6 GHz band. In this exercise, a WAN attachment GW device is wirelessly trunked with a quad-band (2.4, 5Lo, 5Hi, 6 GHz) extender to examine traffic capacity of the trunked link. For purposes of the study, the 6 GHz trunk is deemed to be supported by a 4 x 4 radio scheme at that band (though exploits up to 8 chains are permitted). The near-field meshes associated with the WAN GW are not considered (the assumption being that these data demands are supplied and scavenged directly at the GW/WAN aperture and would not impact performance considerations for the trunked extender and its separately served mesh of clients). The extender is linked to the GW via a 160-MHz wide, bonded 6 GHz channel of “best effort” WiFi priority and this AP sees a mix of traffic it manages with five end devices (three specific 6 GHz/WiFi 6 clients with defined spatial capacity and two “ensemble” devices, representing proxy traffic to multiple 2.4 and 5 GHz clients at the specified bitrates

and link priorities). Three new WiFi6 / 6 GHz clients are shown in the accompanying figure as links L1, L2 and L3. L1 is a connection to a UHD TV which travels 50 feet through 1 floor and 3 wall transitions and exploits a 2 SS radio with 40 MHz of bonded channels to deliver ultra-high resolution and definition video via a 70 Mbps stream. (This is a massive ask for a streaming client.) L2 represents a 2 SS mobile device connection 20 feet from the extender, through one wall. And L3 proxies an HD video stream to a TV or tablet 3 rooms (40 feet, 3 walls) away from the extender supporting a 30 Mbps stream via 2 SS and 80 MHz worth of bonded channels. Link CL1 is a WiFi 5 multiclient proxy which, in ensemble, represents a 400 Mbps demand with VI priority located 30 feet away with 1 wall and 1 floor to transition and commands a 4 SS connection of 80 MHz BW. Finally, CL2 is a single spatial stream of BK priority with the same topographical impediments as CL1 but asks for 50 Mbps support on a single 802.11n, 2.4 GHz WiFi channel. This schema is representative of an existing WiFi-invested single-family detached home which is adding WiFi6 at 6 GHz as a service(s) expansion. A block diagram of the exercise follows:

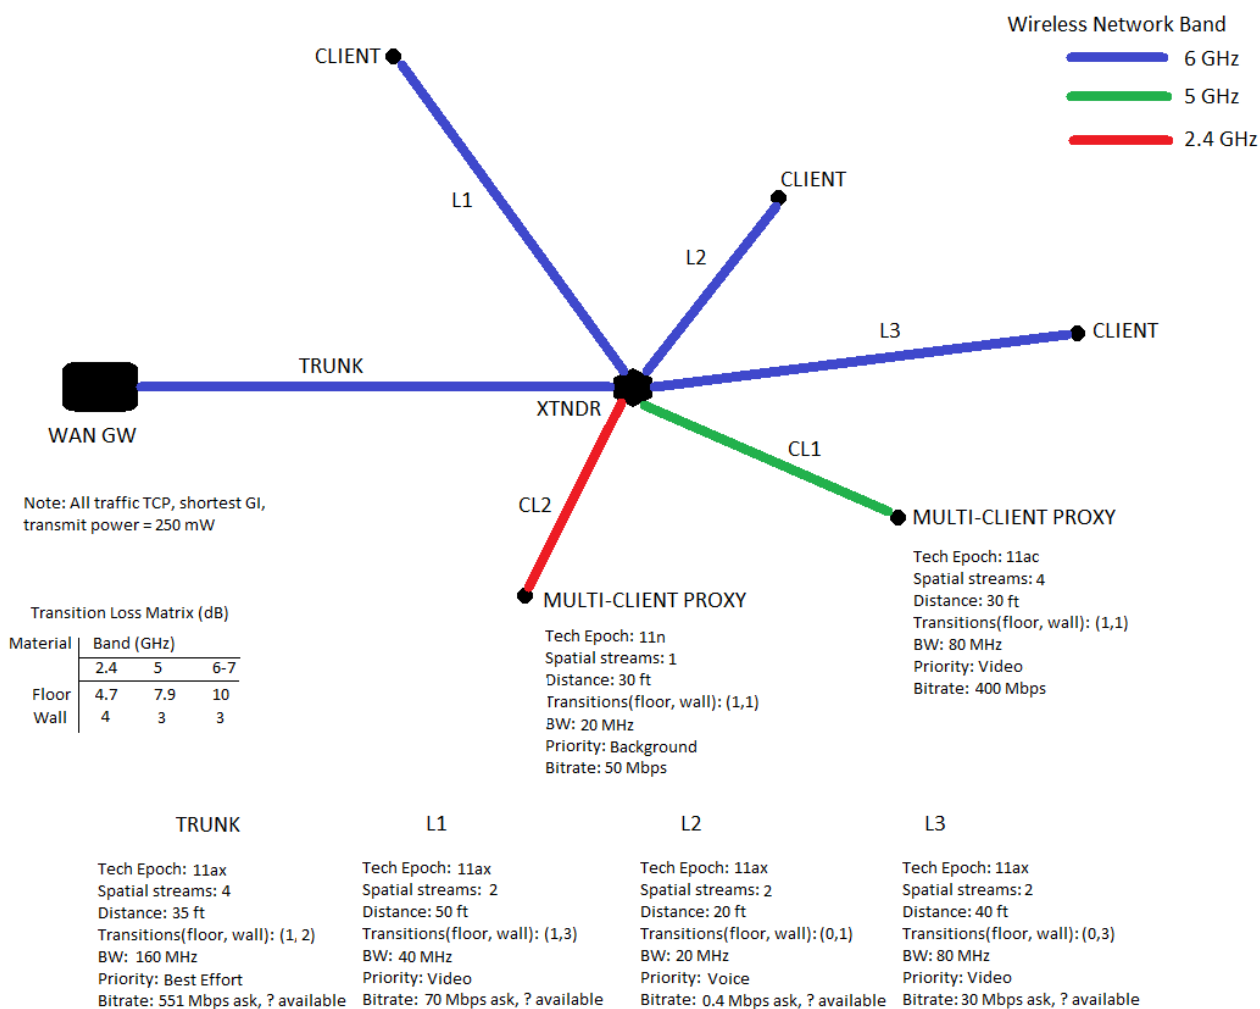


Figure 14 – Model of 6 GHz / 160 MHz BW in-home backbone trunk and service mesh

In addition to the bitrate demands, link losses due to distance and wall/floor transitions are included and link priorities established. The 802.11e priority levels (from highest to lowest) are voice (VO), video (VI), best effort (BE) and background (BK). The resulting whole-home wireless network performance is captured in this summary:

Device	Bandwidth	Radio Type	Tx Pwr (dB)	Spatial Streams	Atten (dB)	Distance	Frequency	Priority	Throughput Requested (Mbps)	Link Rate (TCP) Mbps	%Channel Use !OFDMA
Trunk	160 MHz	802.11ax	24	4	16	10.7m (35 feet)	UNII-6	BE	551	1780.47	
L1	40MHz	802.11ax	24	2	19	15.2m (50 feet)	UNII-6	VI	70	222.05	31.52%
L2	20MHz	802.11ax	24	2	3	6.1m (20 feet)	UNII-6	VO	0.4	236.65	0.17%
L3	80MHz	802.11ax	24	2	9	12.2m (40 feet)	UNII-6	VI	30	596.07	5.03%
CL-2	20MHz	802.11n	27	1	8.7	9.1 (30 feet)	2.4	BK	50	45.17	
CL-1	80MHz	802.11ac	27	4	10.9	9.1 (30 feet)	UNII-3	VI	400	1419.38	
								Aggregate Request	550.4		30.91%
								Total UNII-6 Channel Use			67.64%

**Figure 15 – Expected Performance of the exercise model**

Some significant aspects are immediately apparent: 1) though only fronthaul demand (~550 Mbps) is calculated, if symmetric duplex demand were placed on the trunk (though this is not a likely requirement with the video streaming use cases cited), the total available capacity at the link losses specified would amount to one and a half times the fully symmetric demand (1.1 Gbps) the network would then require; 2) all 6 GHz band traffic (trunk or any of the mesh links) can be distributed without contention (and with surplus BW available); 3) 6 GHz channel capacity on the employed bonding schemes is well below thresholds at which prioritized scheduling need occur (i.e., there are no queuing latencies aside from framing alignment which occur with any of the services mounted – data is dispatched as soon as it is received). Furthermore, as the trunk is a P2P link with management set by the WAN GW, the multi-client (OFDMA) benefits of WiFi6 are not a consideration for this haul – the GW's sole discretionary responsibility is to determine how much bandwidth it needs and where to locate the bonded channels. A note from the legacy support side on this simulation is that the single SS associated with the 2.4 GHz band and the requirement for 50 Mbps service there is that such cannot be supported (at 45 Mbps). This could be remediated by support of a second SS, more BW or an uptick in priorities – but it underscores the motivation for moving away from '11n and onto the MAC support offered by WiFi6.

Digesting this data and simultaneously acknowledging that incorporation of 6 GHz clients into home meshes will consume some fair amount of time, the key inference is that for single-family detached homes, the overwhelming availability of largely interference-free channel BW posed by the 6 GHz band and first scheduling resort of merely distributing trunked traffic in FDM fashion is an obvious source of significant WiFi performance improvement for the home. This further implies that consideration of queuing strategies for opportunistic packaging of the RUs (perhaps against a weighted judgment of total

throughput and accumulated latency on a per-service, per-client basis) – versus merely dispatching data chunks immediately on the assigned channel -- will not be obviously necessary for several years.


MDUs, however, could represent a much different story – bearing in mind that, as pointed out above, initial deployments of band-compatible devices are likely to enter service in scattershot fashion, which will tend to minimize initial CCI emergence. The eventual, thirty thousand foot “qualitative analysis” view is that -- coupled with the potential for even low-power 6 GHz coverage to bleed beyond the intended coverage area of an access point -- MDU’s forced spatial concentration of 6 GHz-compliant APs should test MU-MIMO, OFDMA and BSS coloring attributes of WiFi6 before these ever become necessary alternatives for detached single family dwellings.

And it is worthwhile to note that the 6 GHz spectral pool being manipulated by the APs in an MDU is a necessarily compromised resource. All of the APs compete for leverage of this asset by their clients based upon perceived availability (by either end component of the intended links) of useful spectrum. BSS coloring promotes re-use of spectrum as may be possible, but its very nature also guarantees that the spectrum to be used may be compromised to some degree. This is because with the shifted “channel in use” thresholds associated with discernment of competing traffic from another BSS color, there will inevitably be loss of fade (noise) margin due to ingressing signals on those channels the local AP deems available for use by its mesh. Channels then chosen with nonzero levels of ingressing energy would be forced to reduce spectral efficiency (due to modulation backoff to cope with degrading SINR) and this would necessarily drive down link bitrates (hence, reducing opportunities to pump data, potentially resulting in buffer growth on one end of a given link or the other). A reasonable question might be “How soon does this happen and how dramatic are the effects?”

### **3. Tinkering with expected performance in MDU environments**

#### **3.1. Napkin musings on the scope of the ask**

Granting that the MDU scenario will present the most obvious residential indoor challenge to unlicensed exploit of the 6 GHz band, it seems prudent to attempt to put some numbers to help define the magnitude of any problematic issues. Though not without significant deviation in the data, the “average” MDU in the US comprises 12 units in a two- to three-floor building, each unit of roughly 900 square feet – and most of these one or two bedroom. Floorplans obviously vary, but we might assume 36’ x 25’ units with a centrally located AP and a maximum of 3 wall transitions to service each unit’s palette of 6 GHz clients. To drive AP concentrations up, we can go to 150% of average and conjure an 18-unit MDU with 3 floors of 6 units each.




**Figure 16 – Floor plan of “typical” 900 square foot apartment, AP location in red**

Regarding the units, we might project 2 occupants, each with 3 simultaneously-used client devices (mobile, tablet or PC and a TV) so there are 6 active client devices per unit. To continue to drive the worst case, we can perhaps assume ~240 Mbps per unit (2 x 70 Mbps UHD video streams, 2 x 30 Mbps HD video streams, 2 x 20 Mbps browsing).

### **3.2. Crunching some numbers**

The overwhelming concern in MDUs is the potential for CCI from neighboring units to aggregate towards the middle of the building. A single floor is shown below. The full impact of near neighbor energy would occur in the middle unit(s) of the middle floor, where ensemble distances to either end (AP or client) of all other wireless links in the MDU would be shortest (and hence, CCI strongest):



**Figure 17 – One floor of example MDU (6 units) showing CCI peak location**

With 18 units demanding 4.32 Gbps in ensemble (assuming everything is unicast), the issue becomes how to route the demand with 400 MHz of low power WiFi BW at U-NII-6 and U-NII-8 (10.8 bps/Hz implied). The employ of multiple SS immediately rescues this; in fact using only half of the 8-chain WiFi 6 specification on spatial diversity (4 SS, then), one could theoretically quarter the stated spectral density demand to a very manageable 2.7 bps/Hz (and still not be considering any OFDMA exercise). Put another way: if each (4 SS) AP is considered atomically, 240 Mbps is child's play. Referring to our assumptions on unit topology, a worst-case service radius inside the unit might be 22 feet with 9 dB worth of drywall transition losses. The path loss represented would be 75 dB (66 dB through-air and 9 dB lumped). Recall from Figure 13 that a 250 mW, 4SS WiFi 6 AP with 80 MHz of channel bonding would produce roughly 1.6 Gbps UDP throughput at 75 dB path loss. So even if this device restricts its spectral use to a single 20 MHz channel,  $1600 / 4 = 400$  Mbps (167% of stated demand) should be available to its clients throughout the unit floorplan.

Now, if each AP only requires a single 20 MHz channel to meet its requirements, then there is enough U-NII low-power BW so that each AP in the MDU could exercise its own 20 MHz channel ( $18 \times 20 = 360$  MHz used) without contention for spectrum. One may complain that this is merely fortuitous or somehow a product of a prescient (or perhaps historically aware?) cloud provisioning agent privy to the link data from all of the APs. After all, such peak demand would be the result of arbitrarily originated and continued sessions – which would bloom in aggregate and decay over the course of a day. And there is nothing in the 802.11 specification which restricts an AP to anything other than random pursuit of available channels. But this result could also be an end product of APs whose behavior was such that they routinely mined currently used channel BW for as much client data transport as required before venturing farther afield (spectrally speaking) and were rigorous about initiating connections with the minimum bonding necessary.

There is also a management opportunity here for even more power backoff in the MDU case to improve coexistence implications (within and outside the structure) and still have bursty data overhead in the running links. After all, the service throws per unit are much more constrained than for the detached single family dwelling and there are no floor transitions – and fewer walls – involved. For example in the modeled MDU, the service path loss is estimated at 75 dB and the goal for the unit is 240 Mbps. Referring again to Figure 13, if we drop the power to 100 mW and maintain the same 4 SS AP with 80 MHz channel bonding, we essentially shift our operating point 4 dB to the right; the rate is 1200 Mbps. At a single 20 MHz channel, this still produces  $\sim 1200 / 4$  or 300 Mbps. In fairness, in the game of “WiFi power chicken”, such backoff would need to be practiced by all wireless partners to be most effective, but the end results are significant.

So if available channelization suffices to supply all wireless device connectivity without spectral overlap, as in this particular MDU example, it follows that there are no obvious one-hop latency contributors (due to access contention or CCI, say) past two-way dispatch and recovery mechanisms at each of the two link endpoints; so no backoffs are required. Such latency ought to amount to less than 200 usec then, (roundtrip) for the MAC priorities VO, VI or BE. Relative to WiFi legacy performance in the crowded 2.4 or 5 GHz bands, this pins potential latency in the 6 GHz band to the “extremely responsive” bin.

### 3.3. The Power of BSS Coloring in MDUs

Fortuitous or no, this crafted MDU exercise did not test the most noteworthy of WiFi 6’s MAC benefits, only the bounty which is 6 GHz unlicensed BW. It appears near-future 6 GHz client device populations in indoor residential environments do not look to impose sufficient contention to exercise these MAC mechanisms -- admittedly designed primarily to service large-venue WiFi scenarios. But we can perhaps calculate more Draconian (if unrealistic) indoor plays which would challenge the WiFi 6 MAC if we are willing to significantly up the service loads (as ensemble bitrate and number of separately served clients) past what one would “normally” expect for residences.

If one examines the proposed low power portion of the 6 GHz band for maximal unlicensed bitrate carriage using WiFi 6 mechanics, the asymptotic numbers are impressive: the combined BW of U-NII-6 and U-NII-8 could yield nearly 24 Gbps and 1,440 1 x 1 client devices per BSS served by an 8 SS AP (delivering over 15 Mbps service continuously to each of those clients at close-in service radii). In practice, of course, such is nowhere near achievable; SINR pollution from the OBSS populations and microwave rogue ingresses, differential service radii from AP to clients, non-isochronous session behavior, etc etc all conspire to compromise delivery efficiencies of the MAC and PHY. And on the other side of the coin, there are any number of mobile applications (phone calls representing a classic example) which require much less than 15 Mbps data connectivity -- which WiFi 6’s OFDMA support could interleave into additional client support (over 60,000 simultaneous voice calls, say).

In qualitative terms, the onset of contention issues are guaranteed to occur when simple FDM’ing of the available BW cannot be assured and the bitrate ask on the served channels exceeds asymptotic capacity. But this represents a massive concentration of disparate BSS domains and clients (the former more than the latter, given the implications of lack of control of the OBSS clients). The LPI BW portion of the 6 GHz band being considered at this point in time is 400 MHz. Restricted to a (lucky) distribution of 20 MHz channels, this implies once more than 20 AP’s on different BSS colors are operating such that they (or their client populations) are within  $\sim 10$  dB service radius of all abutting OBSS domains, various backoffs can begin to occur. This effect is accelerated the fewer spatial streams employed by the APs in question and the higher the bitrate (BW) demand – but the onset is anything but precipitous. Reference



Figure 13 again at the 31 dB listed operating point (75 dB path loss @ 100 mW) for an “average” MDU unit and notice that even though MCS backoff is noticeable over the next 10 dB of SINR degradation, the rate is no worse than about half what was sustained before. Unless this attenuation in rate pushes the AP to search for additional BW, the service would simply carry on.

And this type of contention presumes that near-neighbors would not associate interference with close-abutting alternate BSS domains in the first place and fail to seek out better spectral alternatives based upon appropriately shifted channel-in-use detection (they would – that’s the very point of detecting color collision and reassigning colors). That’s the mechanism of OBSS discerned channel re-use – and the threshold of this channel re-use can be shifted to expose the level of interference the AP is willing to entertain.

Given the minimal physical dimensions of the service spaces (call it ~7200 cu ft) and the cubic presentations of the sources of CCI in the 3D model of MDUs, one might suppose that the seven nearest 3D neighbors would pose the greatest threats. (Immediately above and below the unit in question and the five same-floor surrounding units). Note that other units may interfere, but the CCI implications begin to fall off dramatically with increased distance – and surface transitions – that accumulate). If all bitrate consumption in the MDU is high and evenly distributed, then an 8SS AP could still be bonding just a fraction of the available 400 MHz of LPI and, even with huge backoff on MCS (to 10-15 dB of impairment), be capable of delivering over 1 Gbps to its clients due to appropriate repurposing of the “colored” channels available to it.

### **3.4. External Interference by Inside WiFi at 6 GHz**

It is worthwhile to close with a couple of observations on WiFi radiation which escapes residential low power deployments and “enters the wild” (inside/out propagation) since there are some number of legacy outdoor installations which utilize the band (as P2P or P2MP links) and new sources of interference for their operation would not be welcomed with much enthusiasm. To be conservative in estimating impact, one can neglect the selectivity of the link antennas used there (although certainly, the patterns employed are usually very much less than even hemispherical in nature and engaging the antennas on a back pattern would result in ample relative attenuation – front-to-back ratios easily exceeding 40 dB+ in most applications).

Indoor WiFi leakage radiation at 6 GHz needs to overcome the service throw indoors as well as the ultimate transition through exterior walls. In examining outside-in behavior for CBRs, transition loss for wood paneling was pegged at around 6 dB, brick or HardiPlank at 13 dB and stone at around 25 dB. In this case we will reference the loss of wood siding in the calculations, to maintain worst-case parameters for consideration. And based on our previous discussion of detached residences versus MDUs, it appears that the latter can operate with lower WiFi service power with a WAN attachment near the middle of the unit and the former might feature a WAN GW element just inside the exterior wall, operating at higher power. So for the calculations, we will reference 250 mW of WiFi source power with only a 6 dB offset to cite for knockdown of interference radius. Furthermore, we will assume no significant vegetation nor topographical impediments around the candidate interferer.

As to the outdoor equipment, we will assume that the receiver/antenna combination used operates at a convenient 20 MHz BW – but inconveniently on the same frequency as the WiFi in question (and this, at the very bottom of the 6 GHz band), has a 6 dB NF receiver and anticipates worst-case operation down to 8 dB SINR at the detector. (Lots of presumptions here, admittedly). These restrictions set the edge of the operational fade budget for the microwave tower down to a signal level of around -87 dBm. Any

interference would need to be at least 10 dB lower (more preferable) which suggests our leakage radiation needs to be at or below -97 dBm at the antenna to be confirmably non-invasive.

If we hold to our chain of worst-case presumptions, this puts any microwave tower within ~ 2.3 km of the home in question at risk for interference. For perspective, however, if the WiFi AP were operating at 100 mW in the middle of a stone home, itself in a mildly forested area, this radius would collapse to 30 meters (mostly due to internal wall transitions, the huge signal cost of exiting a stone exterior home and tree/vegetation impedances which can easily reach 10 dB or more). A geometric mean of 260 meters is the result of these extreme considerations (though again, this is without allowance for antenna selectivity).

To frame this with radiation patterns in mind, one could assume a tower microwave with a 120 degree main lobe and rationalize the home as having perhaps a hemispherical pattern (referring to the placement of the AP against the exterior wall, the differential internal transitions at minimum would describe a non-isometric pattern of some sort with perhaps 12 dB or so of minimal directive imbalance). If the azimuth offsets of the orientations of the tower and the home were randomly distributed, this implies that perhaps one time out of six when distance to a tower from a home was within the interference radius, the relative directivities of the two radiation patterns would conspire to create CCI.

In generic terms, one would perhaps characterize this as a very slight chance of interference. (And to be sure, the dual of the situation holds – what passes for CCI at the microwave receiver would likely be seen as a “busy channel” by the WiFi – which would eschew – or downrate -- its selection for use in the first place.)

## Conclusion

The 6 GHz band provides a magnificent aperture for unlicensed wireless services to grow in a disciplined and future-proofed direction, providing instant relief for the capacity exhaustion and contention-based loss of efficiency occurring in the legacy 2.4 and 5 GHz bands while also delivering a self-installable, low-latency remedy which grants consumers an indoor-propagation-friendly wireless framework with massive, extensible bitrate support. In detached single-family dwellings, just the insertion of a dedicated 6 GHz / WiFi6 wireless trunk from WAN GW to middle-of-the-home extender promises multi-Gbps, interference-free whole-home WiFi coverage. MDU applications at 6 GHz, while more challenging due to the client and BSS densities implied by multiple overlaid networks operating in close proximity, nonetheless will greatly benefit from the 6 GHz spectrum access while barely invoking MU mechanisms in antenna directivity and the two-axes options for transmission data packing afforded by OFDMA. And the rather more friendly confines of single (one-floor) units in these structures mean additional low power backoff (perhaps to 100 mW, or less) can be implemented without concern for compromised services delivery.

## Abbreviations

AP	access point
AFC	Automated frequency coordination
AGC	Automatic gain control
BE	Best effort
BK	Background

bps	bits per second
BSS	Basic service set
BW	Bandwidth
CBRS	Citizens broadband radio service
CCI	Co-channel interference
CPE	Customer premise equipment
dB	decibel
E2E	End-to-end
EIRP	Equivalent isotropically radiated power
FCC	Federal communications commission
FEC	forward error correction
Gbps	Gigabits per second
GHz	Gigahertz
GW	Gateway
HFC	hybrid fiber-coax
HD	high definition
Hz	hertz
ISBE	International Society of Broadband Experts
ISM	Industrial, scientific and medical
LPI	Low power implementation
MAC	Medium access control
Mbps	Megabits per second
MCS	Modulation and coding scheme
MDU	Multiple dwelling unit
MHz	Megahertz
MIMO	Multiple-in and multiple-out
MU	Multiple User
mW	Milliwatt
NF	Noise Figure
NFC	Near field communications
NPRM	Notice of proposed rulemaking
OFDMA	Orthogonal frequency division multiple access
OOB	Out-of-band
RF	Radio frequency
SCTE	Society of Cable Telecommunications Engineers
SINR	Signal-to-noise-plus-interference ratio
SNR	Signal-to-noise ratio
SS	Spatial Stream
TCP	Transmission control protocol
Tx	Transmit (or transmission)
UDP	User datagram protocol
U-NII	Unlicensed National information infrastructure
VI	Video
VO	Voice
W	Watt
WAN	Wide area network

## Bibliography & References

IEEE P802.11ax/D4.2: *Draft Standard for Information technology – Telecommunications and information exchange between systems Local and metropolitan area networks – Specific requirements*; LAN/MAN Standards Committee of the IEEE Computer Society Copyright 2019

SCTE-ISBE and NCTA: *The New Home as a Hotspot: Wi-Fi Meet CBRS LTE and Meet Your Long Range Brother LoRA*; J.R. Flesch et al, Copyright 2018 SCTE-ISBE and NCTA

*Propagation Losses Through Common Building Materials 2.4 GHz vs 5 GHz*; Robert Wilson, USC 2002

# **DAA, GAP, and Cloud Compute....the Network of the Future**

A Technical Paper prepared for SCTE•ISBE by

**Bill Beesley**

MSO Strategy and Planning Lead  
Fujitsu Network Communications  
2801 Telecom Parkway  
972.442.2653  
Bill.beesley@us.fujitsu.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
Conclusion .....	7
Abbreviations.....	7
Bibliography & References .....	7

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Nielson's Law of Internet Bandwidth showing increases over time.....	3
Figure 2 - Remote PHY Transport Architecture .....	4
Figure 3 - Strand Bloat from multiple devices compared to GAP enclosure concept .....	5
Figure 4 - Fog Computing Architecture .....	6

# Introduction

The coming combination of Distributed Access Architecture, Cloud Compute and the Generic Access Platform offer much promise to the network operator. However, the radical changes in the way distribution networks are built and managed will require new ways of envisioning the network and its capabilities. The new network is ripe with the potential of new revenue generating services and at the same time, fraught with the new engineering and operations issues that operators will have to overcome. This paper will discuss the engineering challenges of pushing equipment and services out to the edge, all the way to the strand, and provide an outline of how manufacturers are working to overcome the space and power constraints that come from moving transmission and compute equipment out of the head end and out into the edge of the network. It will also discuss the operations considerations such as training and tooling field staff to be prepared to manage this increasingly complex network. Lastly, it will offer some forward-looking proposals of future revenue generating services that the converged, intelligent, edge network of the future will support

## Content

After nearly 30 years of working in the telecommunications industry, one thing that I have found to be consistently true is you can never have enough bandwidth. No matter how much is enough to support today's use cases, there will be a new use case or application tomorrow that will require more bandwidth. In 1998 Jakob Nielsen created Nielsen's Law of Internet Bandwidth that states; "A high-end user's connection speed grows by 50% per year." [1] Nielsen was able to demonstrate at the time that bandwidth had grown from 300 bits per second modems to ISDN speeds between 1983 and 1998. He later amended his blog post on the subject to demonstrate that in the ten years since he had written Nielsen's law that his Internet bandwidth had continued to increase to 300 megabits per second.


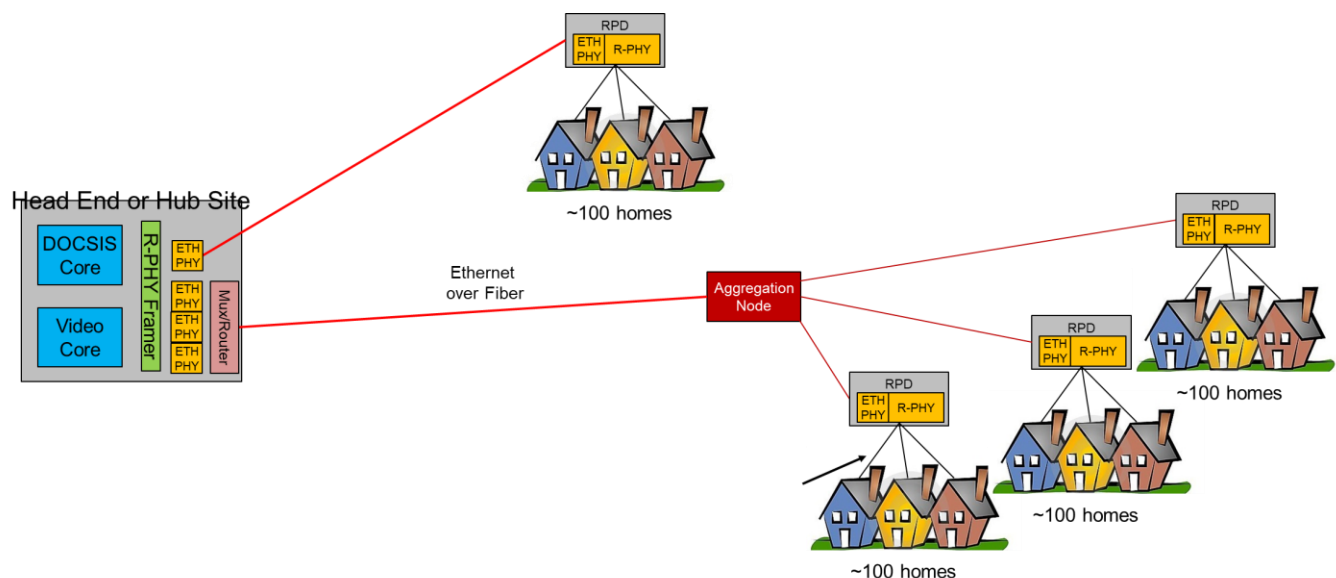


Figure 1 - Nielsen's Law of Internet Bandwidth showing increases over time

This double-digit annual growth rate and the equipment additions to support it has nearly every operator running low on the space, power and cooling capacity they need to maintain and operate traditional transmission equipment. Additionally, as service providers evolve their networks to satisfy their customers with those all-important new services—IP/4K video, business services, IoT, 5G, and cloud/fog computing—transport networks are pushed to breaking point in the struggle to deliver more bandwidth and lower latency that customers demand to support the applications they use..

Providers in the MSO space have begun responding to these bandwidth, space and power challenges by evolving to what has been termed the "Distributed Access Architecture" or DAA, whereby the operator pushes the digital portion of the headend equipment closer to the edge of the network. One of the first use cases for DAA is "Remote PHY" which replaces traditional analog lasers and hybrid-fiber-coax (HFC) nodes with digital optics, usually 10 gigabit Ethernet, that allow placement of the CMTS or CCAP physical interfaces at the edge of the network and closer to the customer. This not only helps increase the amount of bandwidth going to each subscriber, but by moving the physical interfaces out to the edge, frees up precious space and power in the head ends and hub sites.




**Figure 2 - Remote PHY Transport Architecture**

While the primary driver for DAA today is to support Remote PHY, long term this architecture will evolve into a multi-service architecture. The other services supported will include remote CMTS and PON, Business Services transport, Wi-Fi backhaul, 4G/5G xHaul and Edge or Fog computing.

While it has much promise in increasing the amount of bandwidth and lowering latency for subscriber services, this new network architecture introduces its own set of unique operational challenges that now need their own unique solutions. One of the side effects of pushing so many physical interfaces out to the edge is a problem that has been termed "strand bloat." As remote PHY devices, wireless radio heads, small cells, business services aggregation and other new technologies join existing HFC nodes and



amplifiers, the number of devices mounted on the strand begins to explode exponentially. This occurs in part because there is currently no standard definition for a strand-mount housing, requiring equipment manufacturers to build custom housing for each new technology.




**Figure 3 - Strand Bloat from multiple devices compared to GAP enclosure concept**

In response to this issue, Charter, Cox and the SCTE last year created a new SCTE working group to develop a standard known as Generic Access Platform or GAP. The goal of the working group has been to develop a standard definition of size, power, and thermals for an outdoor housing that can be strand-mounted like a traditional HFC node, and that can accept modules to support services such as remote PHY or wireless radios. This will allow manufacturers to focus on building standardized modules to provide specific functionality.

The team of MSO and industry representatives working on the GAP specification started their efforts by attempting to envision the use cases that could a standard installation delivering multi services from the same enclosure could support. Some of the use cases envisioned for delivery via a GAP module would support current transport services such as HFC nodes, amps, Remote PHY and MAC/PHY devices, Remote PON, and Commercial Ethernet Aggregation. In addition, contributors to the GAP specification are considering the ability to support future use cases such as CBRS or 5G small cells and Edge compute.

Strand mounted edge compute is an interesting new use case that will provide not only the capabilities to support compute resources for customers for use in such solutions such as autonomous automobiles, IoT, content delivery and gaming but it will also open up the opportunity to drive network applications out from the head end to the edge. As an example, a virtual CCAP instance could be moved to compute in the GAP node to serve a remote PHY device. Another potential use would be the virtualization of the virtual radio access network and/or virtual baseband unit compute portions of 5G at the edge next to the small cell remote radio head. Distributing the compute and storage closer to the edge is commonly called “Fog Computing” The concept is that, like the cloud, the compute is distributed but as fog is a cloud close to the ground, fog computing is closer to the user.



**Figure 4 - Fog Computing Architecture**

Much like cloud computing created new service and solution opportunities because of its ability to increase the velocity and flexibility with which applications are delivered, fog computing will add additional opportunities due to its decreased latency from close proximity to the user.

As DAA continues to push interfaces ever closer to the customer and GAP attempts to consolidate multiple edge systems into as few housings as possible, there is the opportunity to consolidate both multiple services but also the management of the services provided.

Another change in the way networks are built and managed that will come from the use of these new architectures is the collapsing of disparate networks used today to deliver multiple services. Traditionally networks were purpose built based on the specific requirements of their product use cases. Today we have HFC networks for residential and best effort commercial voice, video and data. We have fiber networks for high value commercial services and carrier xHaul. There are also separate fiber networks for providing video transport between head end and hub site locations. Commonly these networks are built and managed independently of each other. Most MSOs have a group managing their metro networks, another managing the backbone and yet another managing the core aggregation infrastructure for commercial services. There is also likely yet another separate group managing the CMTS, CCAP and other HFC based infrastructure.

Having separately managed and segregated networks was ideal when the services had no method nor reason to co-exist. Thirty years ago, there was not a compelling reason to merge the HFC video network with the business services network delivering a T1 to a commercial customer. However today, voice video and data services are being delivered using Internet Protocols as the common transport. Additionally, as customers continue to use these services ubiquitously across a variety of devices such as

computers, smart TVs, tablets, etc., there is no longer a compelling reason to expend the capital and operating expense to manage the networks separately. In fact, there are disadvantages of segregating networks, especially as traffic begins to traverse between those networks in order to support customer use cases.

Lower operational effort means not only reduced cost but improved customer experience. By consolidating the management of multiple networks, the ability to have a comprehensive end-to-end view becomes possible. This will allow technicians to more quickly resolve issues, as they no longer have to mentally “glue” the flow of information across multiple networks.

DAA and GAP will drive not only network convergence, but also will create new opportunities as the industry begins to introduce new capabilities to the edge access network. By adding wireless to traditional fixed access and by providing edge compute capabilities use cases that require mobility, high bandwidth and low latency such as autonomous automobiles, IoT, and AR/VR gaming, new service revenue streams inevitably will be created.

## Conclusion

As DAA continues to push digital interfaces used to deliver services further to the edge and as the SCTE works to create physical commonality through the GAP standard it is inevitable that both the physical networks and how they are managed will converge. While we are still a long way away from “one network to rule them all” we are starting to see engineering and operations consider how they can deliver a consolidated service delivery network. This will ultimately create opportunities to reduce capex and opex spend but more importantly create new service capabilities that will drive product innovation and revenue.

## Abbreviations

IoT	Internet of things
HFC	hybrid fiber-coax
CMTS	cable modem termination system
CCAP	converged cable access platform
PON	passive optical network
MAC/PHY	media access control layer/physical layer
CBRS	citizens broadband radio service
T1	Transmission System 1
AR/VR	augmented reality/virtual reality
capex	capital expense
opex	operating expense

## Bibliography & References

1. Nielson’s Law of Internet Bandwidth <https://www.nngroup.com/articles/law-of-bandwidth/>

# How to deliver QAM video in a DAA world

A Technical Paper prepared for SCTE•ISBE by

**Karthik Krishna**

Product Marketing Manager

Nokia

200 South Mathilda Avenue Sunnyvale CA 94086

5103812942

karthik.krishna@nokia.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
1. A quick recap on DAA .....	3
2. Challenges of delivering QAM video in DAA .....	3
3. Important considerations .....	4
3.1. How is the broadcast video network laid out? .....	4
3.2. How is the narrowcast video network laid out? .....	4
3.3. How is the video encryption set up? .....	5
3.4. How does the STB-control flow through the network? .....	5
3.5. Is there analog video in the network? .....	6
4. Options for delivering QAM video .....	6
4.1. Option 1: Use Analog Overlay .....	6
4.2. Option 2: Add a Video Adapter in front of the video QAMs .....	7
4.3. Option 3: Add DEPI to existing video QAMs .....	8
4.4. Option 4: Deliver QAM video through a Video Core .....	9
4.5. What does the FMA standard say? .....	11
4.6. Comparing the options .....	11
4.7. Example DAA QAM video deployment scenarios .....	12
5. Evolution to IP video .....	13
Conclusion .....	14
Abbreviations .....	14
Bibliography & References .....	15

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Option1: Analog Overlay .....	7
Figure 2 - Option2: DAA Video Adapter .....	8
Figure 3 - Option 3: Add DEPI to existing video QAMs .....	9
Figure 4 - Option 4: Video Core .....	10
Figure 5 - Example deployment – Scenario 1: DEPI EQAM and Video Adapter .....	13
Figure 6 - Example deployment – Scenario 2: Video Cores .....	13

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Comparing the options for DAA QAM video delivery .....	11

# Introduction

Cable operators are transitioning to Distributed Access Architectures (DAAs) in response to the dramatic increase in High-Speed Data (HSD) traffic on their networks. As cable operators begin planning the introduction of DAA into their networks, one of the critical decisions they need to make is how they will deliver video.

If an operator is not immediately migrating to an all-IP video implementation, they must continue to support Quadrature Amplitude Modulation (QAM) video, which also consumes significant capacity and spectrum. The good news is that a Flexible MAC Architecture (FMA) implementation of DAA enables a variety of options for carrying QAM video regardless of whether the architecture used is Remote PHY (R-PHY), Remote MACPHY (R-MACPHY), or a combination of both.

In this paper, we will address the factors an operator should consider in making the DAA video delivery decision and outline several flexible DAA QAM video delivery options, and discuss their pros and cons.

## Content

### 1. A quick recap on DAA

A DAA distributes specific functions of the Converged Cable Access Platform (CCAP) to other parts of the network. There are several different approaches to DAA. However, with the standardization of FMA, the two predominant approaches are:

R-PHY: Pushes certain DOCSIS MAC functions as well as the DOCSIS PHY and Video PHY to the optical node, creating a Remote PHY Device (RPD). Doing so digitizes the hub-to-node connection and significantly improves signal quality to the end customer.

R-MACPHY: Instantiates an RPD with the DOCSIS Core/virtual CMTS (vCMTS) running locally in the node instead of in the headend, creating a Remote MACPHY Device (RMD). R-MACPHY delivers all the benefits of R-PHY, but also enables huge headend space and overall system power advantages by eliminating the need for a DOCSIS Core /vCMTS in the headend. Additionally, by keeping the DOCSIS components together, R-MACPHY avoids the tight coupling and timing challenges of R-PHY and minimizes latency and jitter concerns for efficient support of capabilities like Low Latency DOCSIS (LLD) and Mobile Backhaul (MBH) over DOCSIS.

### 2. Challenges of delivering QAM video in DAA

In a traditional cable network, broadcast and narrowcast digital video streams are converted into QAM-modulated RF signals either in the headend or hub. This process is done in either a standalone Edge QAM (EQAM) or a CCAP. The resultant video QAM/RF signals are combined with DOCSIS HSD QAM/RF signals and transmitted over analog fiber to Hybrid Fiber-coaxial (HFC) nodes for delivery to subscribers.

With DAA, the QAM modulation and RF upconversion are moved to an RPD or RMD in the outside plant (OSP), digitizing the connectivity to the node. That means an EQAM or CCAP with RF output cannot connect over the now-digital link to the node. Thus, an operator must make adjustments to their existing video network as they adopt DAA. This challenge of QAM video delivery is the same regardless of whether the DAA approach used is R-PHY or R-MACPHY.

### 3. Important considerations

As cable architects design their DAA networks, they must consider how their existing video network is laid out, what video services they currently offer, and how they envision evolving their video network. In this section, we discuss the leading factors to which cable architects must pay attention including:

- ⇒ How is the broadcast video network laid out?
- ⇒ How is the narrowcast video network laid out?
- ⇒ How is the video encryption set up?
- ⇒ How does the set-top box (STB)-control flow through the network?
- ⇒ Is there analog video in the network?

#### 3.1. How is the broadcast video network laid out?

Most broadcast functions are centralized in a traditional cable network, but there are variations in the broadcast network based on the location of the broadcast QAMs and presence of local broadcast channels. The three most common scenarios are:

1. Broadcast QAMs in the hub: This is a common scenario where most broadcast functions are centralized, but the broadcast QAMs are located in the hub.
2. Centralized broadcast: In this scenario, all the broadcast functions, including the QAMs, are centralized with analog super-trunks connecting the broadcast QAMs in the central location all the way to the node.
3. Local broadcast: Here the broadcast QAMs for local channels are located in the hub while all other broadcast QAMs are centralized.

DAA provides options to support all of the broadcast scenarios discussed above. With DAA, the operators can choose to keep their broadcast QAMs as-is in the hub. However, with the QAM function moved to an RPD or RMD in the OSP, it is most natural to centralize the broadcast services.

#### 3.2. How is the narrowcast video network laid out?

Narrowcast services are designed to serve a group of homes in a neighborhood, and hence the narrowcast QAMs are naturally suited to be placed in hubs close to the neighborhood. Video on Demand (VoD) and Switched Digital Video (SDV) are the two predominant narrowcast services.

Before moving to DAA, operators must consider how their narrowcast Single Program Transport Streams (SPTS) are multiplexed and mapped to their QAMs. Today, such a mapping is done in one of three ways:

1. ERM-managed – ISA: In this mode, the narrowcast SPTS sessions are multiplexed and mapped to video QAMs by an Interactive Services Architecture (ISA)-based Edge Resource Manager. ISA architecture is implemented by multiple operators in North America.
2. ERM-managed – NGOD: In this mode, the narrowcast SPTS sessions are multiplexed and mapped to video QAMs by an NGOD-based Edge Resource Manager. Comcast is one of the primary operators using NGOD architecture.
3. Table-based mapping: In this mode, the VOD SPTS sessions are multiplexed and mapped to video QAMs using a pre-defined IP/UDP scheme. This scheme is used only for VOD and is prevalent in regions outside North America.

With DAA, operators need to ensure that the option they select for delivering QAM video supports the existing method for multiplexing the narrowcast services.

Furthermore, with the QAM function moved to the RPD or RMD under DAA, the operators have the option to centralize their narrowcast video network. In order to do so, the operators must ensure their network meets the following conditions:

- The ad-splicing system used for narrowcast services must be centralized along with the narrowcast video.
- There must be sufficient capacity in the fiber connecting the headend to the hubs to accommodate full line-rate narrowcast MPTS traffic.

### **3.3. How is the video encryption set up?**

In North America, the two primary forms of encryption used for QAM video are Cisco PowerKEY and Motorola DigiCipher. In Europe and other parts of the world, DVB CAS is the predominant encryption method. Today, operators use one of three architectures for their video encryption:

1. Proprietary edge encryption: Proprietary encryption of the video services occurs in the video QAM device, either EQAM or CCAP. Proprietary encryption limits the choice of vendor.
2. Standards-based edge encryption: Standards-based encryption of the video services is done in the video QAM device. Standards-based encryption enables an interoperable eco-system.
3. Bulk encryption: The video is bulk-encrypted, or pre-encrypted, before flowing toward the video QAM device, whether EQAM or CCAP. Bulk-encryption provides operators more network flexibility and vendor choice.

DAA provides options to support any of these approaches. However, as operators adopt DAA, they are encouraged to implement either bulk-encryption or standards-based encryption as it enables an interoperable eco-system, provides greater network stability, and enables the centralization and consolidation of the video services.

### **3.4. How does the STB-control flow through the network?**

Cable operator networks are used to transmit the interactive guide and other control signals for video QAM STB operation. Operators use either a dedicated out-of-band (OOB) network for STB-control traffic, or they deliver the STB-control messages using in-band video or via DOCSIS.

1. Separate OOB network:

In this design, out-of-band protocols are used to control STB operation. Vendor-specific headend equipment uses the OOB path to manage the STBs in the field. An OOB modulator modulates the control signals originating from various headend equipment into QPSK RF signals. The modulated OOB RF signals are combined with the DOCSIS and video QAM RF services in the RF combiner. OOB consumes very little bandwidth, but is essential for the operations on the QAM-based STB.

The OOB concept is most prevalent in North America. There are two central OOB systems deployed by operators in America:



**SCTE 55-1:** This is the legacy Motorola/CommScope system that multiplexes and modulates MPEG over IP/UDP streams received from the legacy Motorola headend in the downstream direction, and in the upstream direction, demodulates signals received from STB and packetizes into ATM-like cells over IP/UDP for further transmission (Aloha system).

**SCTE 55-2:** This is the Scientific Atlanta/Cisco system that performs a similar function as SCTE 55-1, but between a Cisco headend and QAM STB using an ATM transport.

The challenge with OOB support is that existing OOB gear in the headend relies on an RF combining network. In DAA, the RF combining network is eliminated, and hence, an alternative solution is required. Furthermore, when operators design their DAA network, they need to accommodate the fact that OOB typically uses two-way communication. DAAs offer various options for delivering OOB over the same digital transport infrastructure used to provide HSD and video services.

## 2. In-band video or via DOCSIS:

In this design, the STB-control traffic is delivered either via in-band video or via DOCSIS. The STB-control in a broadcast setup flows in-band through the broadcast video. The STB-control in a narrowcast setup flows via the DOCSIS path, except for the encryption control which flows through the narrowcast video path.

Since there are no OOB signals used for STB-control in this scheme, there are no special requirements to handle STB messaging under DAA. The same solution used to deliver video services in DAA can also handle the STB-control traffic.

## 3.5. Is there analog video in the network?

Operators are phasing out the analog video from their networks, but it still exists in many parts of Latin America, Asia, and Europe. Some operators have a limited amount of analog video content which they are planning to phase out in the next couple of years. Other operators have a significant amount of analog video content which they may retain for a longer duration.

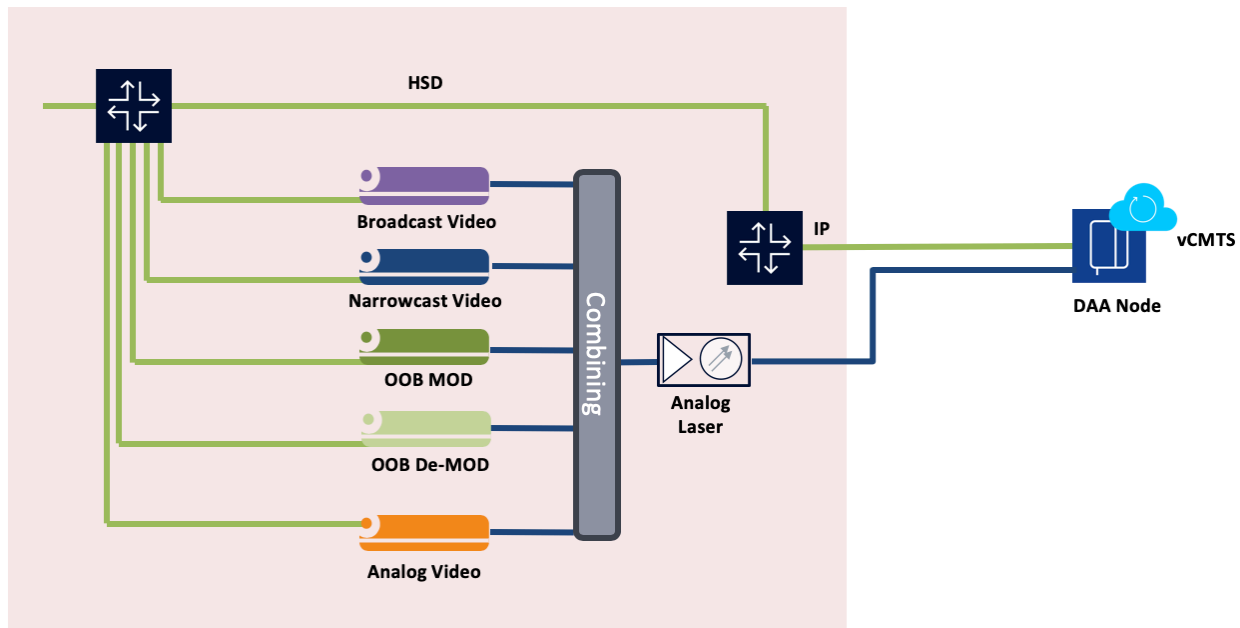
Today, analog video is combined with narrowcast and broadcast QAM services in the hub. As operators transition to DAA, it is an excellent opportunity for them phase out analog video, but it is not compulsory. DAA provides both short- and long-term solutions to handle analog video depending on the number of analog video channels, and thereby provides operators the flexibility and freedom to phase out analog video when they prefer.

## 4. Options for delivering QAM video

DAA provides different flexible options to deliver QAM video.

### 4.1. Option 1: Use Analog Overlay

The simplest option for delivering QAM video in DAA is to use an Analog Overlay. This approach delivers all the existing video-related services, including broadcast QAM video, narrowcast QAM video, analog video, and STB-control, using existing analog fiber infrastructure overlaid on top of the digital link carrying HSD services. The video services can be delivered on a separate analog fiber, or on a dedicated wavelength using wavelength-division multiplexing (WDM), all the way to the DAA node where the optical-to-electrical conversion occurs.



**Figure 1 - Option1: Analog Overlay**


The Analog Overlay is the simplest option as there is no change to the existing video and OOB infrastructure. This option supports the entire line up of analog video and supports most “corner case” scenarios. However, this option provides the least benefits as it:

- Limits how far the nodes can be placed from the hub
- Does not eliminate RF in the headend-to-node connection
- Might require two fiber connections to the node

By fundamentally maintaining two networks, operators incur the highest operating cost by using this option.

#### **4.2. Option 2: Add a Video Adapter in front of the video QAMs**

The Video Adapter alternative requires minimal changes to the network for video delivery, enabling an operator to primarily focus on data delivery as they transition to a DAA. The Video Adapter allows the operator to keep their existing video QAM assets as they are in the headend or hub; the operator simply adds a Video Adapter downstream of the EQAM or CCAP video QAMs. The Video Adapter demodulates the Multi-Program Transport Stream (MPTS) over RF from the video QAMs and outputs digital MPTS either in RTP or DEPI (L2TP signaling) format. The video can then be delivered over the same digital link as the HSD to the node, where the original QAM/RF signal is re-created for transmission to the STB. The Video Adapter supports 1588 PTP to synchronize with the node where the QAM/RF upconversion occurs.



**Figure 2 - Option2: DAA Video Adapter**

To deliver OOB in DAA, the Video Adapter applies the Narrowband Digital Forward (NDF) and Narrowband Digital Return (NDR) concept to handle both SCTE 55-1 and SCTE 55-2 signals. A Video Adapter with NDF samples and digitizes the QPSK/RF signals coming from the OOB modulator for transmission over the converged interconnect network (CIN) to the remote node. The node then re-creates the original QPSK/RF downstream signal for transmission to the STB. In the reverse direction, the NDR on the node samples and digitizes QPSK/RF signals coming from the STB for transmission over the CIN to the Video Adapter, which then re-creates the original QPSK/RF upstream signal. The NDR/NDF approach is defined by the CableLabs R-PHY OOB (R-OOB) architecture.

The Video Adapter applies the Wideband Digital Forward (WDF) technique to deliver analog video in-band using the same digital link to the node used for other services. WDF functionality must also be enabled in the DAA node to support analog video.

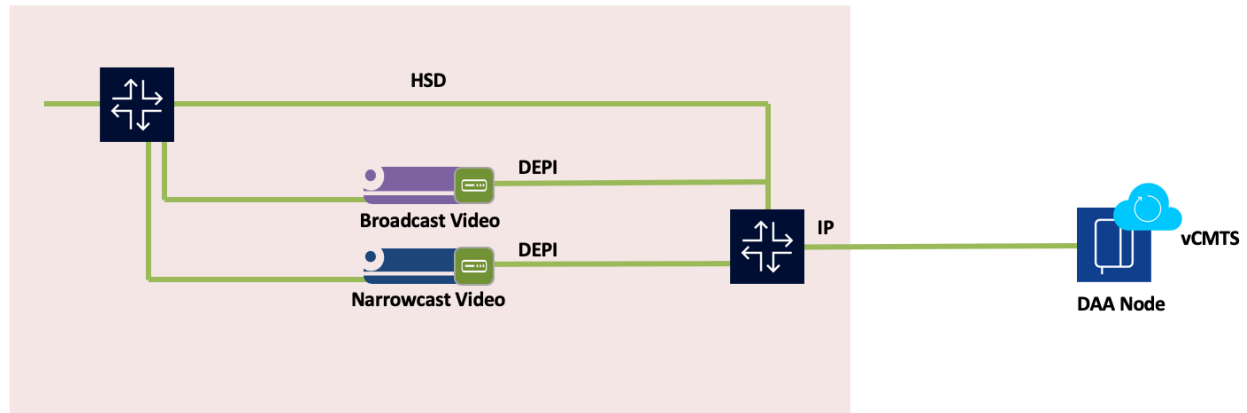
By using a Video Adapter, operators need not change their existing video and OOB infrastructure as they transition to DAA. The Video Adapter can be used for all services including broadcast, narrowcast, and OOB, and can work with any vendors' equipment. The Video Adapter also supports most analog video and "corner case" scenarios. Using a Video Adapter can be a transitional step before adopting longer-term video options explained later in this paper. However, this option could also serve as a longer-term solution.

The downside of the Video Adapter option is that it requires additional equipment in the headend or hub. Also, for OOB application, the distance limitation between the STBs in the field and the OOB gear as specified in SCTE 55-2 still applies.

### 4.3. Option 3: Add DEPI to existing video QAMs

Another option to evolve video QAMs to DAA is to modify the existing EQAM or CCAP video QAMs to send out MPTS via DEPI using L2TP signaling instead of QAM/RF signals. This option ensures an all-digital path to the node in the outside plant. It would require the video QAM vendor to add new hardware

and/or software that supports the delivery of legacy video via DEPI. The QAM vendor may also need to add 1588 PTP functionality for synchronization with the node. The synchronization is essential to ensure the video chassis outputs MPTS packets at the same rate as the QAM modulation inside the node. In some markets, the QAM vendor may offer the option of sending MPTS over IP/RTP in addition to the DEPI.



**Figure 3 - Option 3: Add DEPI to existing video QAMs**

Modifying existing QAMs to add DEPI is an option that does not require additional separate equipment in the headend. However, it does require modifications to the existing video QAM gear to adapt to DAA. Thus the viability of this option depends on the specific QAM equipment in use and support from that equipment vendor.

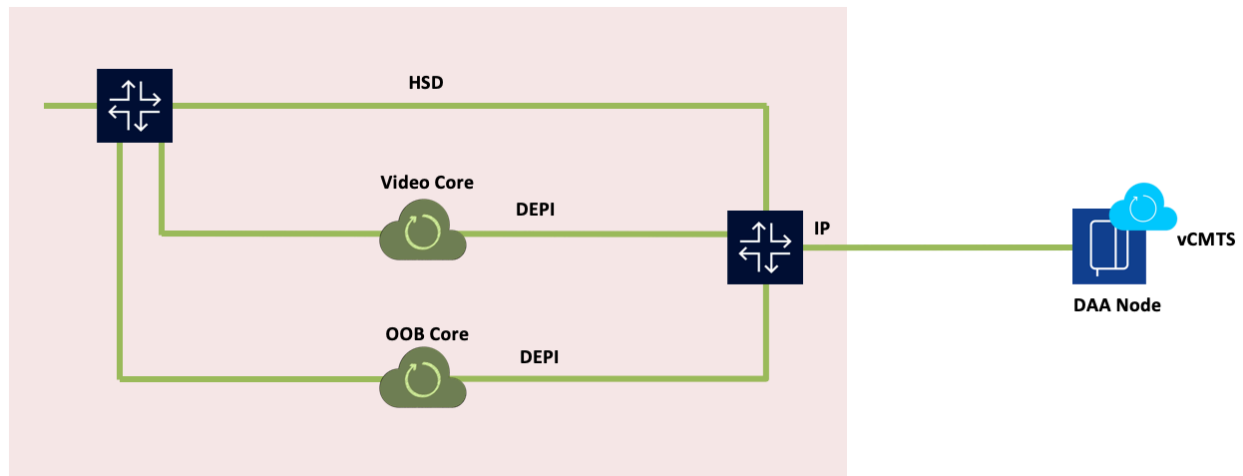
#### **4.4. Option 4: Deliver QAM video through a Video Core**

As discussed earlier, a DAA moves the QAM/RF functions to the node in the outside plant. The remaining core functions for video are described as a Video Core, which is an auxiliary core as defined by the CableLabs FMA standard. In essence, a Video Core performs all the functions of an EQAM except for the modulation and PHY, which are done in the DAA node. The broadcast and narrowcast SPTS that were previously flowing into the EQAM now stream into a purpose-built Video Core instead. The Video Core multiplexes the streams and transports MPTS streams over IP/DEPI to the DAA node in the neighborhood.

The essential functions of a Video Core include:

- MPTS re-multiplexing
- PCR re-stamping/de-jitter
- Program (PSI) manipulation
- DEPI encapsulation
- 1588 timing
- Dynamic setup/teardown of narrowcast video services
- Edge encryption

The video QAMs in an EQAM or CCAP perform all of these functions today with the exception of the DEPI encapsulation and 1588 timing, which are requirements specific to a Video Core.



**Figure 4 - Option 4: Video Core**

### Virtual Video Core

An essential aspect of a Video Core is that it is purpose-built. Thus, there is an excellent opportunity to maximize efficiency by constructing it as a software-based or virtualized Video Core. The distributed nature of a software-based Video Core makes it a perfect candidate to be placed in containers or virtual machines and launched as a cloud-native solution. Virtualizing the Video Core has multiple benefits, including:

- Runs on off-the-shelf servers rather than more expensive, purpose-built equipment
- Reduces the video footprint in the headend and hub, delivering significant space and power savings
- Enables centralization of broadcast and narrowcast video services
- Allows hub consolidation by eliminating the need for any video gear in the hub
- Enables orchestration of the virtual Video Core to simplify network configuration, increasing network agility and service velocity
- Can run alongside other virtualized cores in any location

A Video Core allows operators to phase out their current video assets and sets the stage for evolution to a future-proof, all-IP, all-digital network. Virtualization really makes it easy for the operator to make the transition to an all-IP solution, making a Video Core the most optimal, longer-term option.

### OOB Core

An OOB Core is another auxiliary core used to handle the OOB application. The OOB Core functions are defined by the CableLabs R-PHY OOB (R-OOB) architecture. In R-OOB, the existing OOB modulator is

either replaced or modified to send the OOB downstream signal as MPEG transport streams or ATM cells encapsulated over DEPI, and the existing demodulator is either replaced or modified to decapsulate upstream ARPD datagrams or ATM cells carried over UEPI. The modulation and demodulation functions move to the DAA node in the outside plant. This approach is applicable for both SCTE 55-1 and SCTE-55-2 with a caveat. The SCTE 55-2 OOB system relies on ATM transport and has stringent, low latency timing requirements which, if not met, can cause the STBs to go offline. Hence, some OOB MAC implementation is required in the DAA node to prevent the STBs from going offline. The R-OOB approach resolves the distance limitation challenges and is a better suited long-term approach for handling OOB in DAA.

#### 4.5. What does the FMA standard say?

The FMA standard, currently in development at CableLabs, provides a common framework for video regardless of the selected DAA implementation. The FMA defines auxiliary cores such as the Video Core and the OOB Core to handle video and OOB applications. These auxiliary cores function the same way in both the R-PHY and R-MACPHY FMA alternatives.

The FMA reuses the OOB and video requirements defined in the R-PHY standard. The FMA includes a Video Core that utilizes the Generic Control Protocol (GCP) to manage the resources on both the RPD and RMD. The FMA also supports the concept of a Video Engine, where the FMA MAC Manager is primarily responsible for managing the video QAMs in any DAA implementation, and thereby limits the role of the Video Cores to simple traffic engines with no significant management responsibility. A Video Engine is essentially a Video Core without GCP.

#### 4.6. Comparing the options

In summary, there are a number of options available to the cable operators to deliver QAM video in a DAA network. The approaches include short-term alternatives requiring minimal network alteration as well as long-term more efficient options. An operator's selection will be influenced by their current QAM video network set up, their network centralization plans, and their all-IP video evolution strategy.

**Table 1 - Comparing the options for DAA QAM video delivery**

DAA options for QAM video	When to use it?	Pros	Cons
Analog Overlay	<ul style="list-style-type: none"> <li>Support analog video</li> <li>“Corner case” scenarios</li> </ul>	<ul style="list-style-type: none"> <li>No change to existing video infrastructure</li> <li>No change to existing OOB infrastructure</li> <li>Supports all analog video and “corner case” scenarios</li> </ul>	<ul style="list-style-type: none"> <li>Highest operating cost – essentially maintaining two networks</li> <li>Least benefit – limits distance, does not eliminate RF in headend-to-node connection, requires two fiber connections to node</li> </ul>
Video Adapter	<ul style="list-style-type: none"> <li>Keep existing video and OOB infrastructure as-is</li> </ul>	<ul style="list-style-type: none"> <li>No change to existing video infrastructure</li> </ul>	<ul style="list-style-type: none"> <li>Additional equipment required in the headend/hub</li> </ul>

<b>DAA options for QAM video</b>	<b>When to use it?</b>	<b>Pros</b>	<b>Cons</b>
	<ul style="list-style-type: none"> <li>Near-term IP video migration</li> </ul>	<ul style="list-style-type: none"> <li>No change to existing OOB infrastructure</li> <li>Supports most analog video and “corner case” scenarios</li> <li>Single IP (digital) link to the node</li> </ul>	<ul style="list-style-type: none"> <li>Need to make sure Video Adapter densities align with EQAM locations</li> </ul>
Add DEPI to existing video QAMs	<ul style="list-style-type: none"> <li>Mainstream video QAMs support DEPI</li> </ul>	<ul style="list-style-type: none"> <li>Leverage existing video infrastructure</li> <li>No additional boxes (e.g., video adapter) required</li> <li>Single IP (digital) link to the node</li> <li>Likely “Video Core” ready</li> </ul>	<ul style="list-style-type: none"> <li>Does not support OOB</li> <li>Existing QAM equipment requires modification</li> </ul>
Video Core	<ul style="list-style-type: none"> <li>“Cap &amp; grow” existing video infrastructure</li> <li>Centralize and virtualize the video infrastructure</li> </ul>	<ul style="list-style-type: none"> <li>Replace or “cap &amp; grow” existing video infrastructure</li> <li>Purpose-built appliance or virtualized solution</li> </ul>	<ul style="list-style-type: none"> <li>Requires replacing existing video QAM assets</li> <li>Requires building separate video and OOB core</li> </ul>

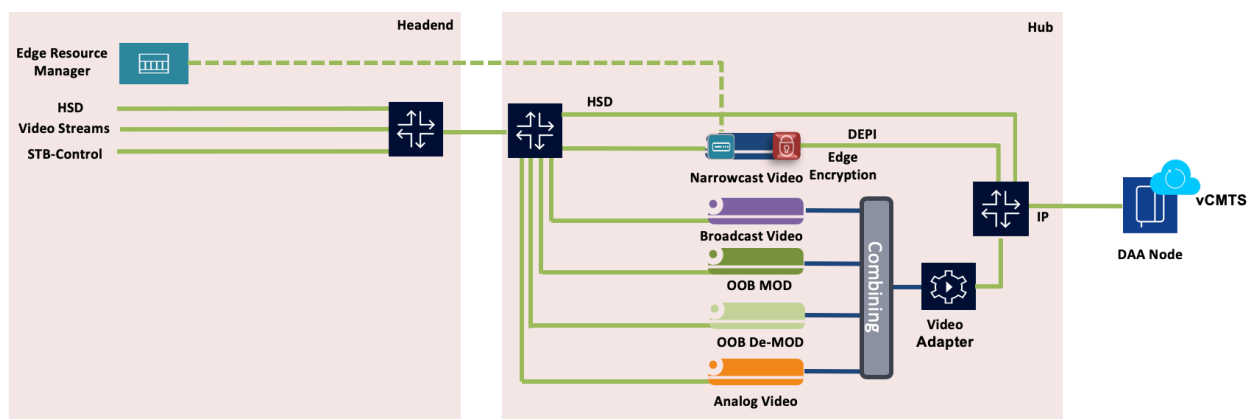
#### 4.7. Example DAA QAM video deployment scenarios

We have discussed some of the network considerations to which cable operators must pay attention when designing their DAA network. We have also reviewed the various QAM video delivery options for DAA. We now present a couple of example deployment scenarios.

##### 1. Example deployment – Scenario 1: DEPI EQAM and Video Adapter

In this example, narrowcast video is delivered by modifying the existing video QAMs to add DEPI functionality. The encryption of the video services is done inside the video QAM device, which is managed by an ERM. Broadcast video, OOB and in some cases analog video are delivered using a Video Adapter. All the QAM video equipment resides in the hub.

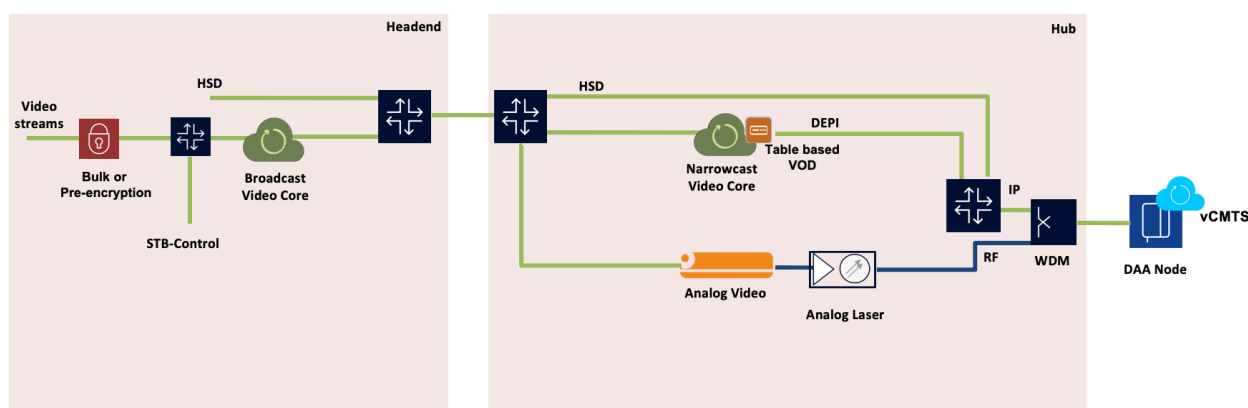
In this scenario, the existing video and OOB infrastructure is either leveraged or retained as-is. This is a possible scenario for most/many deployments in North America.



**Figure 5 - Example deployment – Scenario 1: DEPI EQAM and Video Adapter**

## 2. Example deployment – Scenario 2: Video Cores

In this example, a Video Core replaces the existing broadcast and narrowcast video QAMs. The video streams are either bulk-encrypted or pre-encrypted before reaching the Video Core, and a table-based mapping scheme is used by the Video Core to multiplex the incoming VOD SPTS streams. The STB-Control traffic is delivered in-band through video and via DOCSIS. Finally, an Analog Overlay is used to deliver the analog video. This is a possible scenario for DAA deployments in Europe.



**Figure 6 - Example deployment – Scenario 2: Video Cores**

## 5. Evolution to IP video

Most operators began their migration to IP video sometime back. The following three questions will help an operator gauge how prepared they are to move to IP video:

1. How much of the broadcast and narrowcast content is available on IP video?
2. How much of the IP video content is available on the “big screen”?
3. Where are you in the roll-out of IP-based STBs?

Broadly speaking, cable operators aspire to deliver all of their video services over IP. However, very few are there yet. The reality is that most operators say “QAM video will be in their network for at least the next 5 to 10 years.”



There are different paths operators can follow in their evolution to all-IP video:

- Option 1: Operators can go to all-IP video in conjunction with their DAA transition.
- Option 2: Operators can go to DAA with a simple interim QAM video delivery option (e.g., Video Adapter) before a near- or mid-term evolution to all-IP video.
- Option 3: Operators can go to DAA using a Video Core as a mid- or long-term stepping stone to all-IP video.

Most operators will pursue option 2 or 3, continuing to support some legacy QAM video in DAA as they gradually transition to all-IP video.

## Conclusion

Consumers around the globe are increasingly consuming video via the Internet. This has resulted in an exponential growth of high-speed data traffic, presenting a capacity challenge to cable operators. Most operators will ultimately deliver all of their video over IP. However, this migration will be gradual. In the meantime, operators must increase their networks' data capacity.

DAA is widely accepted as the best solution for ever-greater capacity requirements. Not only does it efficiently accommodate the increased HSD traffic, but it also provides multiple flexible options for operators to accommodate the QAM video in their network as they migrate to all-IP. Operators can deliver QAM video in a DAA by adopting one of four options:

1. Keep their existing video QAM assets as they are and use an Analog Overlay approach;
2. Keep their existing video QAM assets as they are and add a Video Adapter in front of the video QAM;
3. Modify their existing video QAM assets by adding DEPI functionality; or
4. Retire existing video QAM assets and deliver QAM video through a Video Core.

The question of how to distribute QAM video in a DAA network has been thoroughly addressed and should not cause concern or delay an operator's decision to move to DAA.

By applying a distributed architecture and digitizing the network all the way from the headend to the node in the neighborhood, a DAA lays the perfect foundation for evolving to all-IP video. With IP transport as a baseline, operators have a clear path to an all-IP video solution by making minimal changes in the headend and swapping out the QAM-based STBs with IP-based STBs. DAA provides a clear stepping-stone to an all-IP video solution.

## Abbreviations

CCAP	Converged Cable Access Platform
DAA	Distributed Access Architectures
DOCSIS	Data-Over-Cable Service Interface Specifications
EQAM	Edge QAM
FMA	Flexible MAC Architecture
HSD	high-speed data
ISA	Interactive Services Architecture
LLD	Low Latency DOCSIS

MBH	mobile backhaul
NGOD	Next Generation On Demand
OSP	outside plant
QAM	quadrature amplitude modulation
R-PHY	Remote PHY
R-MACPHY	Remote MACPHY
RPD	Remote PHY Device
RMD	Remote MACPHY Device
SDV	switched digital video
VoD	video on demand

## Bibliography & References

*CM-SP-FMA-SYS-D02-190724      DOCSIS Flexible MAC Architecture (FMA) System Specification*

*CM-SP-R-PHY-112-190307      Remote PHY Specification*

# **RDK All Access (Networks): DOCSIS, DSL, PON and Beyond**

## **How Industry Consolidation Created a Need to Support Multiple “Last Mile” Topologies**

A Technical Paper prepared for SCTE•ISBE by

**Marcin Godlewski**  
Director, Product Management  
Technicolor  
5030 Sugerloaf Parkway  
Lawrenceville, GA 30044  
1.404.210.1860  
Marcin.Godlewski@Technicolor.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
1. The RDK Community .....	4
2. RDK and Open Source.....	4
3. Why Global Service Providers Favor RDK.....	5
4. The RDK Deployment Process.....	5
5. The RDK Broadband Architecture .....	6
6. Access Network Evolution.....	7
6.1. Multi-Access Networks Require Unified Software.....	8
6.2. Modular Access Network Architecture Details.....	9
6.3. Message Buses and Component Registrars .....	9
Conclusion .....	10
Abbreviations.....	11

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1: RDK Feature and Deployment Process.....	6
Figure 2: The RDK-B Stack .....	7
Figure 3: Access Network variations in SoCs.....	8
Figure 4: An EPON module plugged into the RDK broadband stack.....	9
Figure 5: “Before and after” messaging configurations within the RDK broadband profile .....	10

# Introduction

Global service providers are no longer tied to a single type of access network – DOCSIS, DSL, Ethernet, or EPON/GPON. Decades of consolidation resulted in a patchwork of last mile access, which necessitated adaptability – in the plant and in CPE (Customer Premises Equipment.) CPE built upon the Reference Design Kit (RDK) began with access networks linked via DOCSIS (Data Over Cable Service Interface Specification), and is evolving to be able to connect over a wide and growing fabric of access network types, from DSL (Digital Subscriber Line), typically used in telco plant, to Passive Optical Networks (PONs) used in enterprise environments.

As multiple access networks become the norm, a need arose for a common and unified software environment, so as to simplify network operations and optimize feature utilization, regardless of underlying network specifics. An advanced suite of WiFi security features, for instance, should be able to be added into a service suite, simultaneously and seamlessly across DOCSIS, DSL, and PON topologies. That common and unified software environment is the RDK (Reference Design Kit) broadband profile, designated in this paper as “RDK-B.”

With a common and extensible broadband software stack, operators can design their product roadmaps beyond the “speed wars,” historically relevant but arguably moot now, with the expansion of Gigabit-grade connections. Feature development can focus on how to bring additional value to devices connected to the access network (again, regardless of last mile type), to differentiate the customer experience.

This paper overviews access network types and related abstraction layers within the RDK/Broadband stack to support them.

## 1. The RDK Community

The RDK community now includes 400+ contributing members, ranging from silicon providers to OEMs (Original Equipment Manufacturers) to systems integrators, and service providers.

From the perspective of an active OEM in the RDK community, and reflective of the consolidating nature of the cable and telecommunications sector, our involvement requires a bit of background. Briefly: Cisco, through its 2003 acquisition of Linksys, inherited and subsequently developed the broadband router software stack, internally called “CCSP,” which became and remains a core component of the RDK profile used in broadband gateways. Technicolor acquired Cisco’s Connected Devices division in 2015, and continues to “vote with code,” in terms of ongoing, open source contributions to the core RDK stacks for video, broadband, and IoT/connected devices. Specific to the CCSP stack, we open sourced it and added operator-requested components. In parallel, we built the RDK stack into all gateways capable of running in DOCSIS (3.0) environments; by the middle of 2016, RDK was running on close to an estimated 7 million broadband gateways, worldwide.

The latest DOCSIS version, 3.1, was fielded onto RDK broadband gateways in mid-2017, and that code base was open sourced to the RDK community last year. We estimate that RDK’s broadband profile is running on more than 70% of the world’s D3.1-based gateways, and continues to advance with operator-relevant features and updates, distributed in a far more agile cadence than previously possible.

## 2. RDK and Open Source

The RDK code base and community are shepherded by RDK Management LLC, headquartered in Philadelphia with core engineering resources in Silicon Valley, Europe (Ireland and the Netherlands), and India. RDKM manages a growing community comprised of more than 400 technology companies, worldwide, with thousands of engineers working on RDK contributions on a daily basis.

RDK operates as an open source project. Its core activities include strategy, roadmap and code releases; tools development and testing; training/technical collaboration; and community events, like this session at the 2019 SCTE Cable-Tec Expo.

Members within the RDK operate along a core set of principles:

- It’s free to join / membership is via a royalty-free/\$0 license
- Development happens in the open
- Leverage community tools and testing wherever possible
- Expect a regular cadence of useful releases
- Contribute and pull from a globally-developed code base
- Operate at a pace that defines product velocity.

The RDK stack used in broadband gateways and routers is 100% open-source, via a royalty-free Apache license. It is designed to be extensible, in that members can bring in other open source software components as desired.

### 3. Why Global Service Providers Favor RDK

At its core, and from its onset, RDK is designed by operators, for operators. As such, it provides an unprecedented amount of technological and strategic freedom, in that operators have direct control over the source code, and can determine the directions they want to go with software and the features important to video and broadband services. The risk of “getting bricked” because of a code or other change, higher up in the chain of command, is nil. Because it was designed by and for operators, RDK stacks are precision-tuned for advanced device management, real-time feature control, device telemetry, and secure code downloads, across all profiles, for all (RDK-based) devices in a home.

One element unique to RDK and the operator community is its tight alignment with System on a Chip (SoC) providers<sup>1</sup>. Prior to RDK, for instance, operators rarely collaborated with silicon providers, working instead with set-top and gateway OEMs who would, in turn, liaise with the chip providers. Those “waterfall”-styled development timeframes were, in fact, a major driver for RDK’s existence in the first place: It used to take as long as three years (and longer), from the time an operator had a concept for a new piece of CPE, to the time that set-top or gateway went into consuming homes. By aligning with SoC providers, operators gain device velocity, as well as full visibility into the RDK stack, which is vital to troubleshooting, triage and optimization.

While RDK began as a video-focused initiative, it has since expanded to become a whole home software platform to manage all devices – video, broadband, and the IoT / connected devices. That’s in large measure because of the driving competitive need to move quickly in both existing and new service domains: Video was first, quickly followed by broadband, and after that the landscape of IP-connected devices from the IoT and elsewhere. One of the comments we hear routinely, as a provider of RDK-based CPE, is that RDK enables service/feature agility for the whole home, with a depth and range of telemetry data that enables related business segments to move more quickly and adeptly, from customer care to field operations and proactive network management.

### 4. The RDK Deployment Process

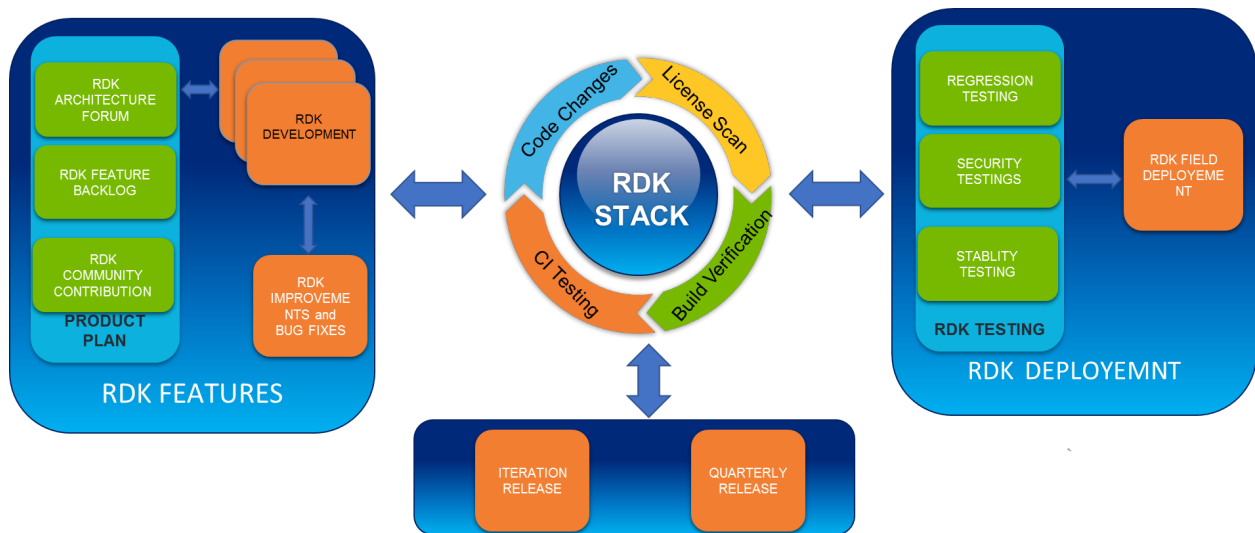
The current process for the RDK stack, in terms of how features and fixes are identified and released, is depicted in Figure 1. From left to right -- from concept to deployment -- operators and vendors develop, test, field-trial, and launch, around the clock, and across more than 400 partners. On the left, a product plan is identified, informed in part by an architectural forum, and RDK feature backlog list, and community contributions. The plan-of-record is developed by the architectural forum, which decides the inclusions for the next release. This is somewhat of a fine art – putting in too many elements can cause delays, yet there’s always more than enough to go into a release.

Once the selected enhancements or fixes are selected, development occurs – operator A may take one element, and vendor B the next, in terms of coding for features and improvements. Field deployment typically consist of regression, security and stability testing.

RDKM manages the process, from license scans to code contributions, continuous integration testing, and build verification. RDK releases are dropped both iteratively (as needed) and quarterly (recurring). A typical quarterly release for RDK-based broadband devices contains roughly 50 “user stories,” which is the lingo of feature enhancements, and 100 fixes/improvements.

---

<sup>1</sup> Broadcom, Intel, Qualcomm and Quantenna, among others.



**Figure 1: RDK Feature and Deployment Process**

## 5. The RDK Broadband Architecture

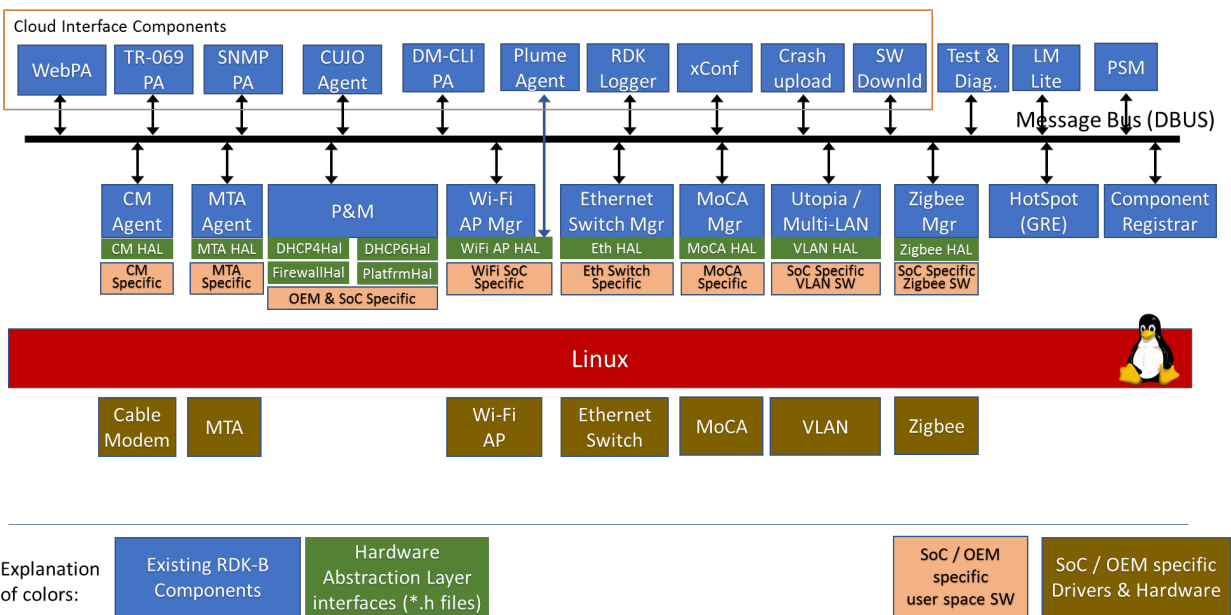
Figure 2 depicts the architecture for the RDK Broadband profile. As can be seen, the stack is entirely based on Linux, with a common message bus (DBUS/RBUS), important for registering attached components, like cable modems, EMTAs (Embedded Multimedia Terminal Adaptors, WiFi Access Points/APs, and Ethernet, MoCA (Multimedia over Coax), and IoT devices, such as through Zigbee.

Hardware Abstraction Layers, or HALs, accompany each RDK device type, and specify the functionality that the SoC needs to provide in its SDK. This means that RDK defines the HAL or header files, and the silicon providers implement the HAL in their SDK. RDK’s broadband profile defines the functions to speak to each agent. For instance, if a cable modem agent wants to know its current channel designation, it “asks” the HAL, which returns the desired answer. This means that RDK, together with silicon providers, define the HALs, or “function files,” designated as “H” files, to define the functionality of each hardware device.

The portfolio of broadband devices communicates via the DBUS/RBUS message bus, for testing and diagnostics and related functions, as well as cloud interface components, such as WebPA, TR-069, SNMP, and various other cloud-based agents.

The RDK stack is expandable, via agents and HALs, to accommodate different device and network types modularly and as needed.





**Figure 2: The RDK-B Stack**

## 6. Access Network Evolution

The “last mile” or access network infrastructure is clearly evolving, from single access (DOCSIS or DSL or PON) to multi-access (DOCSIS and DSL and PON). As mentioned, this is in large part attributable to marketplace consolidation. An operator may have started out with traditional hybrid fiber-coax (HFC) plant, with IP-based services running on broadband DOCSIS infrastructure, then grew by acquisition to support systems with additional access network configurations: Telco DSL, as is often used as a backchannel in satellite-based distribution, or fiber and PON topologies common in enterprise networks. Presently, it is the exception, not the norm, to conjure a major network operator with “only” one access network topology; hence the need for RDK to extend its applicability into all access network types.

This is an evolution that is happening progressively and worldwide. In Switzerland, Sunrise, a mobile telecom provider with assets in fixed-line telephony, TV and broadband, is acquiring UPC. Vodafone, traditionally a mobile carrier, acquired Spanish broadband/telecom provider Ono, and is acquiring portions of UPC, as well as Kabel Deutschland. In Hungary, Magyar Telecom merged with mobile carrier T-Mobile. In Norway, Telia merged with GET-TDC; in the U.K., Comcast merged with Sky; in Argentina, Telecom Argentina acquired Cablevision.

The bottom line: All of them are becoming (or instantly became) multi-access networks, far from their single access network origins. When becoming a multi-access network provider, build options span greenfield (usually fiber), brown field, and “other,” described below:

*Greenfield:* A build environment in a new neighborhood, development, or area previously unserved. Greenfield builds also apply to service areas damaged or obliterated by natural disasters, such as hurricanes, tornadoes, fires and earthquakes. In general, green field builds tend to favor fiber installations.

*Brownfield:* Brown field builds generally favor designs that preserve existing coaxial cable, because of the steep equipment and labor costs associated digging and pulling fiber. Brown field builds are generally based on existing IP connectivity standards, like DOCSIS.

*Other:* If there's one thing that's common in access network topologies, it is variances. In the "other" category, some designs call for direct Ethernet connections, although this is generally considered cost-prohibitive.

## 6.1. Multi-Access Networks Require Unified Software

When transitioning from a single access network to multiple access networks, a need arises for unified software. RDK's broadband profile is a unified software environment that exists to provide a common method to manage various broadband functions on gateways and routers, such as home networking interfaces, device management, and diagnostics.

The rationale for unified software is perhaps best observed by considering the alternative: Different core stacks for each component means that all common functions – testing, debugging, optimizing, releasing, etc. – are handled multiple times, from multiple sources. The permutations can be daunting, involving multiple SoCs, OEMs, integrators, and operators.

A unified software environment, by contrast, enables speed and ongoing optimization. Speed, in terms of more quickly resolving problems, adding features, and deploying applications. Likewise for optimizations, such as pre-emptively identifying and resolving network, stack or device issues before they impact customers, and via machine-level telemetry data and analytics.

A unified software framework is depicted in Figure 3, below. From bottom to top: Access network entities are embedded into SoCs, which become individual, pluggable modules that can be addressed from a unified RDK software layer via each entity's respective HAL. The modular approach enables operators to plug required access network types into the RDK. On the top, applications and services, from home security to WiFi and network controls (both internal and customer-facing) are unaffected, other than their expanded ability to run over additional types of access networks to end users/devices.



**Figure 3: Access Network variations in SoCs**

## 6.2. Modular Access Network Architecture Details

Figure 4 illustrates the architecture of the RDK broadband stack, optimized for extensibility across access network variations. An operator tasked with onboarding or extending its infrastructure to support EPON, for instance, would specify an ONU (Optical Network Unit) with an RDK SoC, and its HAL would communicate through its agent over a common message bus (DBUS, in Figure 4) to cloud interface components, such as webPA, SNMP, testing and diagnostics, and devices, like WiFi extenders and security platforms.

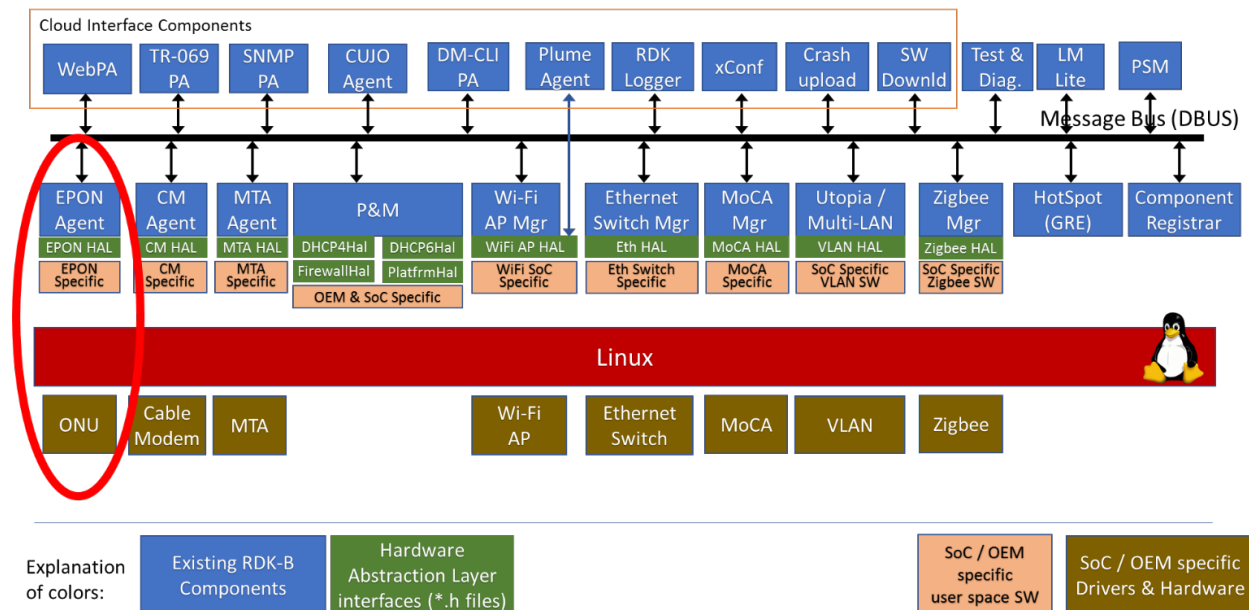
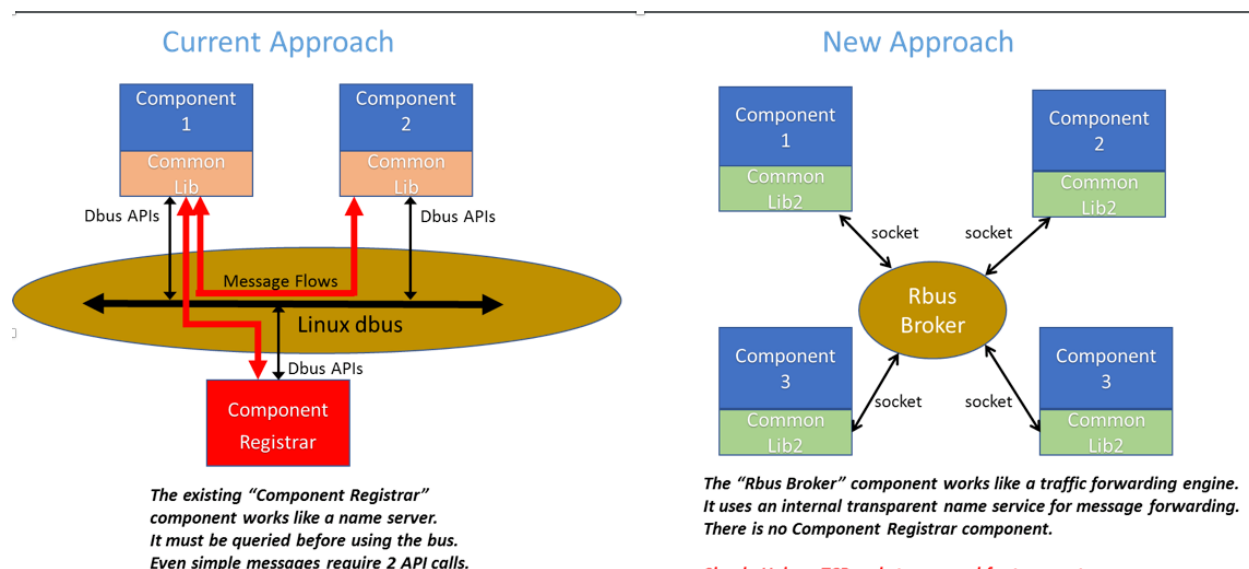


Figure 4: An EPON module plugged into the RDK broadband stack

## 6.3. Message Buses and Component Registrars

A key component in any modular software architecture is a Message Bus, necessary, in this case, to relay information between RDK-based devices, and the cloud-based services they enable. Figure 5 depicts a “now vs. next” treatment of message buses and component registrars within RDK’s broadband profile. In the “now” approach, on the left, a standalone component registrar (Linux DBUS) manages the flow of device-level data to and from connected devices. In the “next” treatment, on the right, individual devices communicate over an RBUS broker. The component registrar recedes, with the RBUS broker forwarding component information much like a traffic forwarding function; Unix or TCP sockets are used for the transfers.

Generally speaking, while DBUS is a stable, standard and relatively common Linux tool for moving messages from one module to another, it is somewhat heavy and slow. By contrast, RBUS is a new tool/utility, which performs the same functions, but in a faster and lighter way. It is an example of ongoing code optimization within the RDK community, to more efficiently communicate between components in a manner analogous to web or TCP sockets – where “socket” is analogous with the port, and the code is vastly optimized, in terms of size and processing power.



**Figure 5: “Before and after” messaging configurations within the RDK broadband profile**

## Conclusion

Global service providers are adapting their last-mile / access network infrastructure in step with marketplace developments, and in particular, industry consolidation. Industry consolidation created an environment in which known, single-access topologies – whether DOCSIS, DSL, or the PONs – are rapidly ceding to multi-access environments, requiring support for different connectivity fabrics. Recognizing this, and at the request of its service provider members, the RDK community responded with a modular, plug-and-play method to augment access network types within RDK broadband profile. Operators who started out with a DOCSIS-based access network can now offer the same devices and services within network topologies like Digital Subscriber Line and E/G-PON.

Such modularity happens in a manner that is non-impactful to existing broadband services and devices. An RDK-based gateway, running a customer- or internal-facing UI to manage WiFi health, for instance, can be adapted in the background to run wherever it lands, be it a DOCSIS, DSL or PON environment. This represents but another chapter in the RDK mindset of continuous improvement and delivery, backed by a community of 400, and serving more than 50 million devices, worldwide. This matters as a mechanism to continue to increase the reach and scope of RDK’s advantages, into multiple network types.

# Abbreviations

CCSP	Cisco Core Service Provider
DOCSIS	Data Over Cable Service Interface Specification
DSL	Digital Subscriber Line
EMTA	Embedded Multimedia Terminal Adaptor
EPON	Ethernet Passive Optical Network
GPON	Gigabit Passive Optical Network
HAL	Hardware Abstraction Layer
IoT	Internet of Things
MoCa	Multimedia Over Coax Alliance
OEM	Original Equipment Manufacturer
ONU	Optical Network Unit
RDK	Reference Design Kit
SNMP	Simple Network Management Protocol
SoC	System on a Chip
TCP	Telecommunications Control Protocol

# **What Can Your CPE Tell You?**

## **Use Cases That Matter When Deriving ML Data from RDK-based Set-tops and Gateways**

A Technical Paper prepared for SCTE•ISBE by

**Bryan Kelly**

Executive Director, Customer Experience Personalization  
Comcast Cable

183 Inverness Drive West, Englewood, CO

+1 (303) 658-7561

Bryan\_Kelly@Comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
<b>Table of Contents .....</b>	<b>2</b>
<b>Introduction .....</b>	<b>3</b>
<b>1. The Timeline Story .....</b>	<b>4</b>
1.1. CX Data Platform-as-a-Service .....	5
1.2. Visualization .....	5
1.3. Messaging .....	6
1.4. Proactive Care .....	6
1.5. Conversations .....	7
1.6. Analytics .....	7
<b>2. The RDK Story .....</b>	<b>7</b>
<b>3. Machine Data Types .....</b>	<b>9</b>
<b>4. Use Case 1: Broadband Gateway Health .....</b>	<b>9</b>
<b>5. Use Case 2: WiFi &amp; Connected Device Health .....</b>	<b>11</b>
<b>6. Use Case 3: Digital Set-top Health .....</b>	<b>13</b>
<b>Conclusion .....</b>	<b>15</b>
<b>Abbreviations .....</b>	<b>15</b>

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1: The Timeline Platform .....	4
Figure 2: A descriptive data screen from Timeline .....	5
Figure 3: How data is visualized in Timeline .....	6
Figure 4: An analytics dashboard in Timeline .....	7
Figure 5: The RDK Stack .....	8
Figure 6: Machine diagnostic data used to inform CXELs .....	10
Figure 7: A depiction of STB-related telemetry data .....	13

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1: Partial list of analytic processes fueled by machine-level RDK data. ....	9
Table 2: Client device types used for WiFi experience telemetry, monitoring and analysis .....	12

# Introduction

As the landscape of “big data” resolves into tangible progress using machine-level (ML) data and artificial intelligence (AI) to proactively optimize the network, service providers using gateways, digital set-tops and related cloud components that are based on the Reference Design Kit (RDK) are reaping beneficial returns. This paper highlights multiple use cases involving RDK-derived data to provide context about, predict and fix problems before they impact customers, recovery with alacrity when problems occur, and apply anomaly detection to resolve “edge cases” and related challenges that occur across mixed cloud environments.

Operators around the globe are making use of machine-level (ML) data, derived from RDK devices, to feed artificial intelligence (AI) engines and algorithms. Those algorithms, in turn, fuel business- and care-facing dashboards designed to continuously monitor and proactively address events that impact customers. That machine-level data is increasingly mined from in-home devices, including digital set-top boxes and broadband gateways, and based on the RDK’s open-source software stack inside system-on-chip (SoC) silicon, both in set-tops and broadband gateways.

For the past few years, representatives from the RDK community have addressed Cable-Tec Expo audiences, detailing progress in business and technical operations and involving Artificial Intelligence (AI) algorithms fed by RDK-sourced machine data. This year, Comcast, as one of the RDK’s founders, will detail how it developed a customer experience optimization data platform, called Timeline, which builds a continuous and linear view of each interaction a customer has with the company – be it a series of nonscheduled equipment reboots indicating service problems, or an actual call.

This paper describes how the Timeline software platform, developed by Comcast, creates an empathetic and linear view of its customers’ experiences in terms of their interactions with the company including but not limited to signal quality, care-related incidents, product usage and assistance. It was built with APIs that are leveraged by many departments within the company, from product, to care, to technical operations.


Within, we describe three use cases, some in-production and some in-development. These illustrate how RDK-based data derived from set-tops and gateways is being used within the Timeline platform: 1) Determinations of whether broadband gateways are healthy; 2) Measurements of the aggregate and individual health of WiFi-connected devices; and 3) Proactive analysis of error conditions within set-tops to obviate truck rolls and improve the customer experience.



## 1. The Timeline Story

The Timeline platform began as a simple quest for empathy. With data about customers' experiences invariably scattered across a company, it was difficult (understatement) – both for consumers and service providers -- to really know the whole story about how different products and services were doing.

Consider: There are dozens of touch points with consumers, and depending on the nature of the journey, a customer may have questions about or interactions with their existing products, marketing touchpoints, product constructs, online forums, frontline agents (chat, phone) and technicians.



**Figure 1: The Timeline Platform**

What was missing was a “beginning-to-now” chronological platform, capable of connecting products, workforce tools and any customer interaction, and serving as the foundation for personalized, consistent customer experiences. It needed to be visual, like a dashboard, so that care, dispatch, and any organization focused on making the customer experience our best product could see the whole experience, in one place – both past history and interactions in real-time. With the addition of APIs, any system would be able to leverage that data to provide improved experience context, informing agents why customers may be calling, send proactive communications to customers, offer powerful journey analytics to the business to understand the impact of key CX initiatives and many other data driven initiatives designed to provide differentiated treatment.

That’s what triggered what became the Timeline platform, in 2015; since then, it has captured over 40 billion interactions from over 50 million customers, and fields over 25 billion consumer inquiries per month. Internally, it is used by more than 35,000 employees monthly, and has been used by more than 130,000 employees across the Comcast organization. Timeline is integrated with dozens of internal tools, including Einstein 360, Tech360, Xfinity.com and many more. Operationally the platform has been integrated with multiple systems including Slack, Salesforce, DataDog and Pagerduty, providing a seamless, self-service based operational environment across data producers, platform engineers, software

developers, product owners and consumers of the Timeline products. In 2016, it was voted “most impactful CX tool,” in an internal recognition program; since its inception, it has saved the company millions of dollars by improving first call resolution, reducing truck rolls and providing improvements across the customer experience.

The Timeline Platform has grown to include six product subcategories, detailed below.

## 1.1. CX Data Platform-as-a-Service

“Timeline Elements” is a data ingest tool that provides (internal) users a simple and self-service means to connect their data. The foundation of any customer experience exchange platform is its data; Elements is envisioned to augment the applicability of an enterprise data lake management tool with data ingest, prep, presentation management, metadata and messaging management, governance, security and APIs. Internal design goals included making it easy and frictionless for various entities within Comcast to submit configured data, so that Elements can automatically build the schemas and trigger the processes necessary to start ingesting, validating, monitoring and visualizing the data (in a non-production environment.)

**Describe Your Data**

**What:** For each field, choose one of the available options from each section and provide the purpose of the field.  
**Why:** The details you provide will be used to operationalize your data feed and populate the data dictionary for other users.

**ENTITY TYPE**  
 Set up the entity for this event  
 Customer

**EVENT CLASSIFICATION**  
 The intended usage of this event  
 Default

**CUSTOMER ID**  
 The field that represents customer ID  
 customerId

**TIME**  
 Time of the event  
 timestamp

**TIME ZONE**  
 Where it occurred  
 UTC

**TIME FORMAT**  
 How it is recorded  
 yyyy-MM-ddTHH:mm:ss.SSSZ

**accountNumber**  
 string  
 CAN THIS FIELD BE NULL?  
 Yes Data may or may not be NULL.  
 No Data cannot be NULL.

**FIELD LABEL**  
 Set the category of the field label  
 Person

**FIELD NAME**  
 Set name in Customer Timeline  
 Account Number

**FIELD VALUE**  
 Set the category of the field value  
 IDs & Codes

**FIELD DESCRIPTION**  
 Description that will show in legend  
 Account Number

**ENCRYPT**  
 Encrypt this field when stored. Consider for PII

**passwordReset**  
 Password Reset  
 string  
 CAN THIS FIELD BE NULL?  
 Yes Data may or may not be NULL.  
 No Data cannot be NULL.

**FIELD LABEL**  
 Set the category of the field label  
 Interaction

**FIELD NAME**  
 Set name in Customer Timeline  
 Reset Status

**FIELD VALUE**  
 Set the category of the field value  
 Labels, Attributes, Status

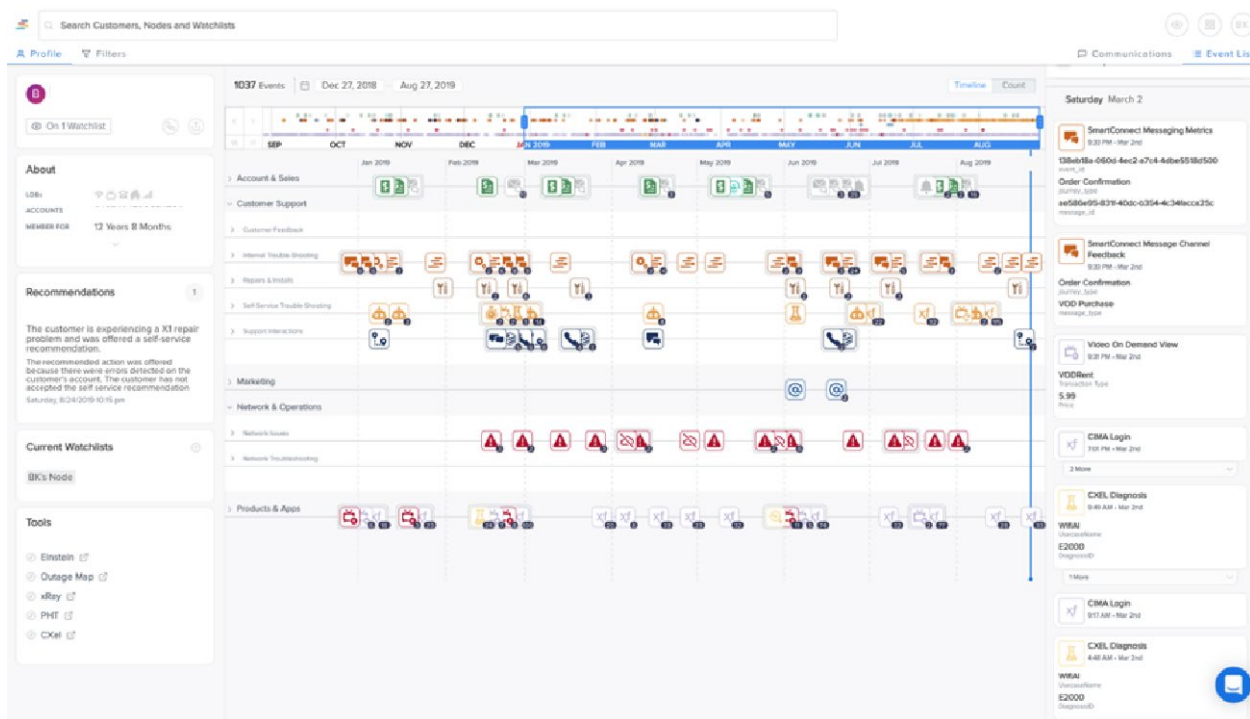
**FIELD DESCRIPTION**  
 Description that will show in legend  
 The result of the reset attempt.

**ENCRYPT**  
 Encrypt this field when stored. Consider for PII

Figure 2: A descriptive data screen from Timeline

## 1.2. Visualization

“Customer Timeline” is a journey visualization user interface (UI) that sequences customer interactions, including their device(s)/service(s) health, so as to better understand and handle each customer experience. It is an intuitive, 360-degree and lifetime customer view, for employees to see historical and real-time status information, in one place. The intent is to be able to see the big picture, target trouble areas, and dive into specific events. Event details include information about timing, who was involved, and what was done, with an eye toward solving problems faster.



**Figure 3: How data is visualized in Timeline**

### 1.3. Messaging

“Smart Connect” is a real-time, automated messaging and notification subsystem, designed to engage the workforce and customers with messaging and conversations supported by machine learning (ML) and artificial intelligence (AI)-based business rules. The reasoning: reactive communications do not don’t cut it anymore. Customers are where they are, and are smartphone-trained to start a conversation within an app or other digital experience. The Smart Connect system leverages real-time event data and provides the curated best-next-step to any consuming product, channel or tool via an API. From “Tech ETA” messaging, that provides visual and near-real-time arrival information during service calls, to an automated appointment waitlist, when working from home and available for a sooner service window, to frontline agent, next best action recommendations.

### 1.4. Proactive Care

“Timeline Watchlist” is a customizable, real-time monitoring tool designed to proactively organize, track and support specific segments of customers enabling differentiated treatment. For context, the feature arose out of a constant frustration amongst field and care representatives, who wanted better ways to stay in touch with customers having problems, to see the matter through. Watchlist is driven by real-time data feeds, so that employees are updated the moment a customer “on the Watchlist” is having a problem. With mainstream messaging and email integrations, employees are empowered to take action on their customers’ behavies. The Watchlist dashboard enables live lists, that automatically refresh, to track a static list of customers in care situations, or build lists that dynamically changes, based on set preferences.

## 1.5. Conversations

“Timeline Talk” will connect conversations across all communications channels (web, phone, SMS, chat, email) to reduce customer effort and personalize experiences. It puts all customer conversations in one place, no matter who talks with a customer, or why. Anyone can initiate and review calls, SMS and email interactions, in one thread. Talk also lets team members to communicate, and update the customer, without switching windows. The reasoning: With cross channel shared context, everyone is on the same page, conversations are continuous and problem solving is simplified. Talk will leverage an integrated and AI-powered virtual agent to triage customer issues and quickly connect them with the correct resolution team. Calls and written correspondence are stored in one place, with advanced search functionality, so that prior interactions and current status data can be retrieved in seconds.

## 1.6. Analytics

Timeline’s analytics platform provides a dedicated analytics data environment with a standardized structuring, enabling partnerships with leading journey analytics providers. Data flows through the Timeline platform and into the analytics environment in real time. This empowers the business to understand the capability of their processes designed to improve the customer experience, the impact of those initiatives on broader business objectives, including simplifying customer effort through digital transformation and reducing churn. Most importantly, the analytics platform combined with the advanced journey analytics tools, enables the business to leverage data as a guide to understanding the next opportunities for improving the customer experience. For example, an “outlier management” tool highlights customers based on their experiences, and illuminates those experiencing the most frustration. It provides fast insights into specific cohorts of customers – those connected to a downed node, those with chronic problems related to a specific equipment type – to prioritize and take corrective action.

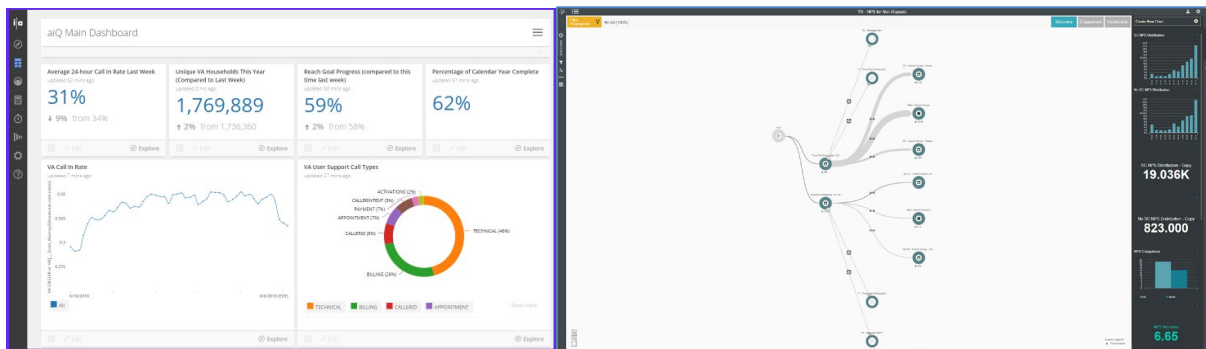


Figure 4: An analytics dashboard in Timeline


## 2. The RDK Story

The first formal gathering of the RDK community happened in December 2012, in New York, and was co-hosted by Comcast and what was then Time Warner Cable; Liberty Global joined soon after. Its intent, then and now, was to hasten service agility and gain more visibility into the inner workings of digital set-tops, to begin, and later, broadband gateways. Characterized as a royalty-free distribution of shared source components, RDK is intended to be a template and software distro for building CPE software. For system operators, RDK is designed to enable them to control their own destinies, using software and data.

The first well-known outcome of the RDK was Comcast’s X1 video experience, at the time a complete break from what had been a blue grid-styled guide. Since then, Liberty’s Horizon, Cox’s “Contour,” Rogers’ “Ignite TV” and Shaw’s “BlueSky TV” have all been powered by the RDK.

Subsequently, a broadband stack was developed by the RDK, as well as additional profiles to cover additional connected devices. In that sense, RDK is designed to be a modular, portable and customizable open source software suite for the connected home. It standardizes core functions used in video, broadband and IoT devices, providing a common data model and enabling operators to manage those devices and easily customize their UIs and apps. By mid-2019, more than 50 million RDK-based devices had been installed in homes throughout the world.

The RDK software stack is depicted in Figure 5, below.



**Figure 5: The RDK Stack**

For the purposes of this paper, the RDK’s contributions to analytics and telemetry, to improve the customer experience and drive business results, will be the focus. Most of that data is derived from the lowest layers of the stack shown in Figure 1 – the RDK hardware abstraction layer (HAL) and device management layer.

Within Comcast’s RDK team, some 30 engineers and data scientists routinely monitor telemetry markers, derived from the double-digit millions of fielded digital set-tops, gateways/routers, and connected devices, like cameras. For broadband devices, as many as 30 criteria are gathered, and classified by rules that indicate whether a device’s health is good (green) or bad (red.) For instance, a telemetry marker for a home may show that the gateway is green, but the router (and therefore the WiFi) is down. That information is then relayed to internal tools used by care agents.

### 3. Machine Data Types

Considerable informative data can be extracted from the RDK-based firmware inside digital set-tops and broadband gateways, which can be correlated with other data sets – customer care, infrastructure metrics, in-home signal strength – to proactively optimize CPE and related business systems.

Among the data that can be gathered and analyzed: Dropped packets, channel change responsiveness, voice recognition timeliness, MPEG/IP layer information about video quality, tune/stream errors, and home network fault segmentation (WiFi or wired LAN).

RDK-based analytics can enable service providers to better manage the functions listed in Table 1:

**Table 1: Partial list of analytic processes fueled by machine-level RDK data.**

Device Trials/Rollouts	Feature Trials/Rollouts
Release Rollouts	Field Issue Triage
Customer Call Reduction	Truck Roll Reduction
Customer Experience Improvements	Functional/Performance Analysis
Trend Analysis	Log/Crash Analysis
Device Optimization	Network Optimization

### 4. Use Case 1: Broadband Gateway Health

Among the anecdotal evidence that triggered the original RDK, then focused entirely on the video experience, is that prior to initiatives like RDK, it was difficult, if not impossible, to “see” into various software stacks vital to the customer experience. Because of then-monolithic and proprietary code, sourced from multiple internal and external providers, conducting service or equipment triage generally involved two things: Time, and finger-pointing. It follows that a core design goal with the original (video-oriented) and consequent RDK profiles is stack visibility, so as to continuously monitor and assess performance, and proactively mitigate device or service issues.

A category of machine level diagnostic data is being developed within Comcast and designated a “CXEL,” pronounced “sixel.” It’s a linguistic mashup of CX, customer experience, and pixel, to mean the baseline unit of a video image. CXELs can come from several sources, and “go red” when operating out of various parameters. Used to indicate overall health, gateway CXELs inform just part of the picture. When combined with CXELs sourced from the layer3 network and the access/last mile network, they make it possible to know where problems are occurring – is the problem with the WiFi, the last mile network, or the “big” Internet?

The machine-level criteria used to define a gateway-related CXEL is depicted in Figure 6. Across RDK-based gateways and WiFi devices, more than 30 criteria are collected every four hours, to ultimately



assess device and overall network health. Markers (left column) that start with “RF” generally indicate DOCSIS-related data. Markers that begin with “SYS” correlate with RDK broadband data. “Range” (middle column) indicates the number of times the markers are reported. The markers and range data is combined to measure the health of the gateway, via the CX-el. If a device meets an aggregate score of 80+, that gateway’s CXEL is flagged as red (problematic). Similarly, all markers that begin with “WiFi” (not pictured) are used to develop WiFi-targeted CXELs. If the combined CXEL score for WiFi markers/range is 30+, the device is marked as red / problematic.

A	B	C
Marker	RANGE	Weight
SYS_ERROR_DNSHostname_Error	27	30
SYS_INFO_Transition_RedLED	15	30
RF_ERROR_WAN_stop	3	30
SYS_INFO_Transition_WhiteLED	8	20
SYS_ERROR_Zero_CID	2	60
RF_ERROR_T4_TIMEOUTS	15	30
RF_ERROR_wan_restart	3	30
RF_ERROR_LAN_stop	1	70
SYS_ERROR_CMTSretry_lock_restart	200	20
RDK-10037	15	30
SYS_ERROR_Webpareconnect_PingMiss	4	30
RF_ERROR_IPV4PingFailed	2	40
SYS_ERROR_erouter0link_not_ready	6	30
RF_INFO_CMSTATUS21	15	20
SYS_SH_RDKB_FIREWALL_RESTART	10	30
SYS_INFO_XI5_detect_enabled	2	40
RF_ERROR_WAN_stopped	4	40
SYS_ERROR_erouterLink_down	3	30
SYS_ERROR_Zombie_dnsmasq	2	40
SYS_ERROR_FPM_Pool_InvalidToken	6	30
SYS_ERROR_LOGUPLOAD_FAILED	5	40
RF_ERROR_IPV6PingFailed	1	20
SYS_ERROR_MemAbove600	10	20
RF_ERROR_MDDLost	5	40
RF_ERROR_Wan_down	4	30
SYS_ERROR_Nvram_spacefull	30	20
RF_INFO_CableCut_event	2	30

**Figure 6: Machine diagnostic data used to inform CXELs**

## 5. Use Case 2: WiFi & Connected Device Health

When service providers first began offering WiFi connectivity, the most frequent customer calls asked “what’s my SSID.” After that: “Why isn’t the WiFi reaching the back bedroom” / the farther reaches of a home. At the time, the machine-level “hooks” didn’t exist in a way that could be correlated, interpreted, and improved. More recently, queries have advanced to cluster around “why isn’t the Nth connected device working”?

A CX phenomenon well known to service technicians is the speed test, almost never conducted to actually assess received speeds, but rather because something is working right – a streaming video is buffering, a download is taking too long.

In all cases, the advancement of behind-the-scenes analytics, based on telemetry data, can vastly improve the service experience, to the point that a reduction in speed tests could likely be identified as a key performance index (KPI).

At the 2018 Cable-Tec Expo, panelists on an RDK session discussed at length the notion of a “WiFi Happiness Index,” developed as a way to create a weighted, realistic view of WiFi behavior in a home. It involves more than 50 parameters, and enormous amounts of machine-derived data – making it an excellent candidate for AI handling, which gets stronger with increased data loads.

Since then, work has advanced to fine-tune the business of interpreting gateway health, in ways that feed apps like, in Comcast’s case, the xFi Assistant, which shows customers things like how many devices are connected, network activity, and how to troubleshoot devices. (It is impossible to discuss xFi Assistant without two bits of color commentary. First: A “pause WiFi” component is known internally as “the child locator,” because they come running when Mom pauses the signal! Second, last summer, when an xFi customer retreated to the Colorado mountains with her husband for the weekend, she quickly clued into the party her teenager was hosting at home – by the volume of “device joins” reported via the app; a quick peek into the connected back door camera, after phoning home, showed a pile of teenagers exiting the premises. The story goes on: Later that night, the xFi app’s camera captured an intruder, and triggered a photo notification: A Colorado black bear was on the back patio.)

Another of the AI/ML advancements, since last Expo, is the collection of additional telemetry data from broadband and WiFi gateways, then correlating it with other gateways on the same (logical) node, to ascertain whether groups are “going red,” indicating localized or systemic problems. When a device or logical node of devices turns red, RDK components are being developed to report descriptive information, which can be exported into ML and AI engines, like Timeline, along with recommended actions – either proactive self-help (along the lines of “we see you’re having a problem, try this”) or to prescribe corrective action to technicians connected to the Timeline platform via other connected enterprise apps.

It is one thing to see a logical node of gateways “go red,” however, and quite another to know if those gateways have truly “gone down.” Significant analysis was performed to correlate any red logical nodes, comprised of RDK-based gateways, with other known “bad events,” such as unscheduled reboots over a period of time, or T3 timeouts, which are associated with the connection between a DOCSIS device to its CMTS. (Over a certain “pain threshold,” too many T3 timeouts signal a larger problem.) Such machine-derived data is being fed into related decision engines, used by other care-oriented engines, to continuously interpret, cross-check and eliminate customer-impacting anomalies.



Work is similarly underway to advance the WiFi Happiness Index (WFHI). As usage increased, it became apparent that viewing the household’s WiFi environment, in aggregate, was sub-optimal. Viewing signal clarity and reach, based on certain mean and median throughput thresholds, was insufficient, such as in terms of mean values for RFFI and channel utilization. In essence, prior versions of the WFHI aggregated measurements along different parameters, assessing an overall household snapshot that was largely inaccurate, from a connected device perspective. In reality, and in constant flux, the needs of embedded WiFi radios differ: An iPhone 5 can only do so much, relative to an iPhone XS; likewise for laptops and all other WiFi-consuming devices, and as bounded by their respective silicon footprints.

Therefore a need existed, to assess the overall health of an in-home or in-enterprise WiFi service from the perspective of the devices connected to it, to proactively “see” and welcome them, from a provisioning perspective. (Along the lines of, “Hi, you’re new here, welcome, let me get you set up for the best you can do.”) Simultaneously, a visual and notification-ready means was developed, to alert technicians and customers of the same, and via the tangents of the Timeline portal.

Behind the scenes, and for a household, the machine-level data available from RDK devices can be applied to an internal “collective pain index.” When someone is high on the CPI, corrective action is necessarily maximal. The CPI under development within the RDK sources machine data from five types of devices, shown in Table 2, and examines the normal and pain thresholds for those devices, as a substandard derivation, per device. Such data is enormously useful in deriving maximal and cumulative device-triggered pain in a home, and goes a long way in proactively correcting our own issues that inadvertently impact our customers.

**Table 2: Client device types used for WiFi experience telemetry, monitoring and analysis**

Broadband Device	Laptop, smartphone, tablet
Media Player	Set-top
Media Source Device	Camera
IoT Device	Thermostat, lightbulb, sensor(s)
Access Point	WiFi extender POD using WiFi-connected backhaul

To summarize, forthcoming versions of the WFHI will include per-device telemetry information, in part to keep customers informed of the realities of their WiFi-connected device capabilities, and in other part to ensure that those per-device needs are being met. This mitigates current customer pain points, like having to tell someone that their IoT thermostat or connected webcam is too far from the gateway. By aggregating the per-device pain thresholds into the larger WFHI, we attain a more realistic and granular view, that can be harvested internally for ML/AI support engines, and, ultimately, by customers, to see current device state.

On the Timeline side, those RDK-sourced CX-els about WiFi and gateway device health are ingested, processed, and correlated with downstream impacts, like call-in rates (CIRs) and truck rolls (TRs). Interested care and field representatives can be proactively alerted, and impacted customers proactively messaged – we see a problem, we’re working on it, thanks for your patience – before any CIRs or TRs. In

many cases, service calls can be obviated with a notification of suggested action; already, this has saved millions in non-rolled trucks, and earned positive NPS points with customers.

## 6. Use Case 3: Digital Set-top Health

A third correlation between RDK-derived telemetry data and the AI engines at work to improve field operations and customer experience is the proactive analysis of error conditions with set-top boxes and the resultant video experience. Comcast currently supports millions of RDK-based digital set-tops throughout its U.S. footprint, with the potential to extend the footprint into Europe via its Sky acquisition (and via the addition of additional access network types discussed separately in this SCTE workshop) in the future.

From a machine learning and AI perspective, that aggregate footprint is a prodigious producer of “big data.” A telemetry team within Comcast, responsible for applying telemetry and analytics to set-top data, relates that it gathers upwards of 113 Terabytes of data per day of technical telemetry data, which is equivalent to roughly 4.5 billion daily events, ranging from single key presses to channel tunes to closed captioning and/or secondary language activation.

The data also enables the team to proactively optimize the video experience, such as through stack state analysis. For instance, RDK data, known internally as “XRE data,” for “Cross-platform Runtime Environment,” can represent the last 30 pages a customer saw, before they executed a tune. If a design goal is to perpetually make it easier and more intuitive for consumers to find and consume the content they desire, then the stack state should be optimized to get smaller: A person went through fewer screens to get to a tune event. Telemetry data is a reliable, accurate, and fast means to empirically identify whether or not a feature or change works, in terms of the customer experience.




Figure 7: A depiction of STB-related telemetry data

From a field technician perspective, XRE error gathering and analysis has already resulted in the identification and correction of set-top tuning anomalies in some in-home networks that use MoCA (Multimedia over Coax). Last year, the XRE telemetry team began seeing an appreciable increase in XRE-10007 QAM tune failures on two types of RDK-based set-tops. When the devices were in IP-only mode, they tuned successfully, and at a latency similar to a QAM-only or hybrid QAM-IP devices. But if the tune request traveled over a MoCA path to consume a QAM channel, the success rate dropped to the 89% range. The RDK code was consequently optimized to immediately tune via IP, if a MoCA path failed a QAM tune. The action produced video tuning events that succeeded 99.2% of the time.

From a Timeline perspective, XRE data represents a regular flow of data, flowing into the data mart that sits behind the Timeline platform and that can provide context to related internal care systems, like Einstein360, Comcast's agent customer management tool. As a direct result, care agents are proactively alerted to XRE events when speaking with or otherwise helping a specific customer. Timeline orchestrates the event data and makes it available to Einstein360's ITG (Interactive Troubleshooting Guide) function. The ITG guides the agent through the resolution of the problem. It's a useful and fast way to more quickly understand a situation, from problem to resolution, simplifying the conversation. It essentially lets care agents know, with high likelihood, why a customer is in touch, and how to fix the problem.

# Conclusion

Data analytics are an invaluable tool to proactively optimize device and network health. The analytical data derived from RDK-based set-tops and broadband gateways/routers is vital to the service provider's quest to make the customer experience its best product. This paper outlines three use cases, fueled by RDK-sourced data markers, and spanning broadband gateways, WiFi, and video, via set-top boxes. In all cases, the telemetry data available is vast and informative, capable of device and network diagnostics not previously applicable.

When populated into Comcast's internal Timeline platform, these data inform a linear view of each customer's experiences, whether related to signal quality, care-related incidents, and proactive assistance. The Timeline platform began as an internal, employee-generated platform, initially shown at Comcast's thrice-annual "Lab Week" program, and because involved care employees sought a way to keep in touch with customers they'd been helping, to assure that an incident was satisfactorily resolved.

The landscape of "big data" is showing tangible progress in the use of machine-level (ML) data and artificial intelligence (AI) to proactively optimize the network. Service providers using gateways, digital set-tops and related cloud components that are based on the Reference Design Kit (RDK) are reaping beneficial returns. This paper highlighted multiple use cases involving RDK-derived data to provide context about, predict and fix problems before they impact customers, recovery with alacrity when problems occur, and apply anomaly detection to resolve "edge cases" and related challenges that occur across mixed cloud environments.

# Abbreviations

AI	Artificial Intelligence
API	Application Program Interface
BPS	Bits Per Second
CIR	Call in Rate
CMTS	Cable Modem Termination System
CPI	Customer Pain Index
CX	Customer Experience
CXEL	Customer Experience marker
DOCSIS	Data Over Cable Service Interface Specification
ITG	Interactive Troubleshooting Guide
FEC	Forward Error Correction
HAL	Hardware Abstraction Layer
HFC	Hybrid Fiber-Coax
HD	high definition
Hz	hertz
IP	Internet Protocol
ISBE	International Society of Broadband Experts
KPI	Key Performance Index
LAN	Local Area Network
ML	Machine Learning

MPEG	Moving Pictures Expert Group
QAM	Quadrature Amplitude Modulation
RDK	Reference Design Kit
SoC	System on Chip
SSID	Service Set Identifier
SCTE	Society of Cable Telecommunications Engineers
TR	Truck Roll
UX	User Experience
XRE	Cross Platform Runtime Environment
WFHI	WiFi Happiness Index

# **Cable 10G Vs. Wireless 5G – Foe Or Friend? A Survey Of Next Gen Network Directions**

A Technical Paper prepared for SCTE•ISBE by

**John Ulm**  
Engineering Fellow  
CommScope  
978-609-6028  
john.ulm@CommScope.com

**Zoran Maricevic**  
Engineering Fellow  
CommScope  
203-303-6547  
zoran.maricevic@CommScope.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Content .....	5
1. Cable 10G – A journey, not a destination.....	5
1.1. Cable 10G Overview.....	5
1.2. Cable 10G Key Attributes .....	6
1.3. Network Capacity Planning for 10G .....	7
1.4. Migrating HFC from DS only 10G to More symmetrical 10G .....	8
1.4.1. Full-Duplex DOCSIS (FDX).....	8
1.4.1. Ultra-High Splits and Soft-FDX.....	9
1.5. Cable 10G - Summary .....	9
2. Gigabit Wireless Technologies – 5G, CBRS, WiFi-6 .....	11
2.1. The Outside Inside Discussion.....	12
2.2. The Inside Outside strategy .....	14
2.3. 5G as a Fixed Wireless Access solution.....	20
2.3.1. FWA Economics .....	20
2.3.2. FWA Summary.....	21
2.4. Guiding Wireless Trends for Tomorrow's Smart Home.....	21
3. Using Outside HFC Plant for 5G – considerations and logistics .....	23
3.1. Case Study – HFC Power for Wireless Cell sites .....	23
3.2. Mapping 5G Small Cells to Existing HFC Plant .....	26
3.3. HFC Adaptability in a changing World .....	33
Conclusion .....	34
4. Cable 10G vs. Wireless 5G – Friend or Foe?.....	34
Bibliography & References .....	35
Abbreviations.....	36

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Full Duplex DOCSIS (FDX) Spectrum Band Options .....	8
Figure 2 – Automated Driving, Mobility Services, and Augmented Reality .....	11
Figure 3 – 28 GHz vs. 70 GHz Spectral Efficiency.....	12
Figure 4 – 37 GHz System Performance & Spectral Efficiency .....	12
Figure 5 – Sub 6 GHz with Outside the home antenna .....	13
Figure 6 – mmWave Success Requires Roof, Window, and Wall Installations .....	13
Figure 7 – Schema of Inside/Out Propagation Studies for 3.5 GHz CBRS .....	15
Figure 8 – Sub 6 GHz Connections to the Home.....	15
Figure 9 – Networks to the Home Cell.....	16
Figure 10 – Home Cell Coverage.....	16
Figure 11 – LTE Range versus Performance for Inside Out Coverage.....	17
Figure 12 – 1W HaaT Coverage for 25/5 Service (1 x 10-16 Homes) .....	18

Figure 13 – 4W HaaT Coverage for 25/5 Service (1 x 20-40 Homes) .....	18
Figure 14 – Service Mounting Potential for Outdoor 4W CBRS Mesh .....	19
Figure 15 – In Home/Office 60 GHz Coverage .....	19
Figure 16 – Defining the Roaming (Home Exterior) Proposition for CBRS/LTE.....	23
Figure 17 – A N+3 HFC plant example with Power Supplies, Voltages and current draw .....	24
Figure 18 – N+3 HFC plant upgraded to FTTLA, with Voltages and current draws.....	25
Figure 19 – Median Lot Size, new single-family detached homes started in 2015 .....	26
Figure 20 - New England suburb, illustration for variability of lot sizes within .....	27
Figure 21 – Higher density (~190 HP/mile) HFC overlayed with ~200m/150m radius cells.....	29
Figure 22 – Higher density (~190 HP/mile) HFC overlayed with ~150m/100m cells .....	30
Figure 23 – Higher density (~190 HP/mile) HFC overlayed with 100m cells.....	30
Figure 24 - Lower density (~60 HP/mile) HFC overlayed with ~200m/150m cells.....	31
Figure 25 – Lower density (~60 HP/mile) HFC overlayed with ~150m/100m cells.....	32
Figure 26 – Lower density (~60 HP/mile) HFC overlayed with 100m cells .....	32



# Introduction

There has been much hype around 5G in the wireless world for several years now. At the start of this year, the cable industry announced its own 10G vision. Cable 10G and Wireless 5G initiatives offer disruptive, revolutionary technologies that at first glance seem to be at odds. But when combined, they offer an evolutionary strategy with much synergy.

The cable industry has undergone unprecedented technology changes over the last several years. The introduction of DOCSIS 3.1<sup>®</sup> started this and was quickly followed with development of Distributed Architectures such as Remote-PHY (R-PHY); Full-Duplex (FDX); Low Latency DOCSIS initiative; and now Extended Spectrum efforts. These developments are all building blocks that are a part of reaching 10G goals. The paper will review these goals and what it means for the operator to achieve these.

Gigabit Fixed Wireless technologies are emerging to rival wired services. 5G will be offered across different spectrums with each band having unique capacity/distance tradeoffs. Wi-Fi 6<sup>®</sup> and CBRS have also entered the fray. Doesn't this position Wireless as foe to Cable 10G? The technologies actually need each other to make a better system. The future high bandwidth, high frequency wireless systems need small cells with many access points requiring a low latency wired backhaul; and APs positioned inside the home/MDU and outside in every neighborhood for optimum coverage. Cable is ideally suited to support this backhaul and powering infrastructure. Meanwhile, Cable 10G can provide multi-gigabit capacity to the home's entry point but needs a robust high capacity wireless connection for that final 100 meters inside and around the home to every mobile device.

In the end, we describe how Cable 10G and Wireless 5G/CBRS are much stronger together and are at the core of a next gen network evolution.

# Content

## 1. Cable 10G – A journey, not a destination

On January 7, 2019 at the Consumer Electronics Show (CES), NCTA – The Internet & Television Association®, CableLabs®, and Cable Europe® introduced the cable industry’s vision for delivering 10 gigabit networks, or 10G™ – a powerful, capital-efficient technology platform that will ramp up from the 1 gigabit offerings of today to speeds of 10 gigabits per second (Gbps) and beyond – to consumers across the globe in the coming years.

10G is a goal, a lighthouse in the distance towards which all MSOs and vendors can steer their boats. It will take some time to get there. It is just a single point in a continuum of improvements that will occur in the future. And it is not the end-point, it is an interim point. We will likely push on beyond that 10G point in the future. The focus in this paper is on the migration needed over the next 7-10 years using existing technologies to achieve the 10G goals. [CLO\_2019] takes a deeper look into the future with some potential new technologies to see where the industry might be in 15-25 years beyond 10G.

But in getting to 10G, what does it take for the DOCSIS/HFC System to deliver on the 10G Service Level Agreement (SLA) promise? Many MSOs do not really know what it takes to do 10 Gbps Downstream, let alone 10 Gbps Symmetrical. To reach 10G goals might require an aggregate of last-hop technologies to accomplish this; HFC, DOCSIS, PON and Wireless are all possible candidate technologies. [ULM\_2019] provides a detailed discussion on 10G network capacity planning along with outside plant considerations and logistics. Some highlights from that paper are provided in this section.

### 1.1. Cable 10G Overview

A quick introduction to 10G can be found on the CableLabs website [www.cablelabs.com/10g](http://www.cablelabs.com/10g):



#### What is 10G?

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 lower latencies, enhanced reliability and better security in a scalable manner.

#### Why do we need the 10G platform?

Our digital future will stall without a platform that can meet our needs. While we don’t know what the next trend will be, we do know the internet will be central to it. By advancing device and network performance to stay ahead of consumer demand, 10G will provide a myriad of new immersive digital experiences and other emerging technologies that will revolutionize the way we live, work, learn and

play. Like the saffron in paella, or the milk in a latte, our industry's networks and innovations are the crucial ingredients in creating a better future for humanity.

In the downstream (DS) direction, 10G will be helpful in providing delivery of immersive video services (virtual reality & augmented reality for Holodeck experiences). It will also be useful for providing more snappy service. For example, downloading:

- A two-hour HD movie in 3-4 seconds (vs. 5-7 min @ 100 Mbps)
- A two-hour 4K UHD movie in 12-15 seconds (vs. 20-25 min @ 100 Mbps)
- A large gaming program such as Call of Duty's Black Ops® (~100 Giga-Bytes) in 90 seconds
  - instead of a 2½ hours @ 100 Mbps.

In the upstream (US) direction, it could be used for providing more snappy service again, but it could also prove to be very useful in enabling low-latency transport. The extra bandwidth (BW) helps enable Predictive Grant Services (PGS) to accelerate US delivery. That lower latency can permit 5G backhaul and mid-haul, and if the latencies drop low enough, it could even permit 5G fronthaul.

## **1.2. Cable 10G Key Attributes**

There are four key attributes to the 10G platform:

1. Speed
2. Latency
3. Security
4. Reliability

10G's promise of faster speeds, more capacity, lower latency and greater security will enable and help fully realize a wide variety of new services and applications that will change the way millions of consumers, educators, businesses and innovators interact with the world. Much of the underlying technology has already been specified or is work in progress.

The speed attribute will leverage technologies such as:

- DOCSIS 3.1
- FDX, Extended Spectrum DOCSIS
- PON
- Coherent optics
- Advanced Wi-Fi including Wi-Fi 6 (a.k.a. 802.11.ax)

Some applications driving the need for Lower Latency DOCSIS include: Gaming; VR/AR (avoiding nausea); CoMP; and autonomous navigation. Latency is a function of packet processing times, queuing times, transmission times, and propagation times. We can improve all areas. CableLab's Low Latency DOCSIS (LLD) project includes ideas in all these areas. The existence of 10G bandwidth capacity also helps, because higher bandwidth capacity leads to less congestion and permits new techniques like Predictive Grant Service (PGS) to expedite BW Grants in the upstream (US). Work is also being done in low latency mobile X-haul and low latency Wi-Fi.

Security is an integral part of 10G. Work continues at CableLabs in Micronets, secure downloads, & MACsec. This will become part of the new DOCSIS 4.0 specification. Vendors are also working to make more secure systems with separate, isolated processors and memory in chips.

With respect to reliability, new DAA Node designs of the future will likely be adding in redundancy in processing and redundancy in NSI-Side links to Leaf Switches in CIN as Moore's Law improvements in gate density help. Reliability is being addressed by proactive network maintenance (PNM) and dual channel Wi-Fi. This will improve observability and redundancy and allow A/I monitoring (PNM, more data analytics in DAA nodes, redundancy in ring of DAA).

And for all four 10G attributes (i.e. speed, latency, security, reliability), it is equally important that they scale... on all markets.

### **1.3. Network Capacity Planning for 10G**

The "10G" in the announcement is for 10 Gigabits per sec (Gbps). But what exactly does that mean as there are different ways of measuring capacity? For example, Liberty Global's Virgin Media division in the U.K. ran tests earlier this year over a 10G EPON network and demonstrated users getting 8.5 Gbps throughput. It turns out that this is extremely close to the theoretical maximum throughput for 10G PON technology. "10G" PON has a raw physical rate of 10 Gbps but there is ~15% overhead from the PHY and MAC layers. So, the customer nets 8.5 Gbps from a 10G PON.

Our analysis 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.

So, what kind of service tiers will subscribers enjoy in this new high 10G bandwidth era? As previously discussed, the 10G PON provides a net downstream capacity of ~8.5 Gbps to the consumer. This capacity might reasonably support a downstream SLA of 8 Gbps. The service group (SG) utilization (i.e.  $N_{sub} * T_{avg}$ ) for a 64 subscriber PON might grow to a bit over 1 Gbps in the 7-8 year window. That means a consumer with a 8 Gbps SLA will have a QoE coefficient of  $K=0.9$  to 1.0 which is reasonable for this relatively small SG size.

Getting to a true 10 Gbps downstream SLA that is equivalent to 10G Ethernet will mean providing slightly greater than 10 Gbps network capacity. This will push the PON networks into next generation PON technology (e.g. 20+ Gbps). Because HFC can incrementally add capacity with additional spectrum, there are certain downstream scenarios that will be discussed where existing 1218 MHz HFC might be able to hit this target. In general, future technologies, such as 1.8 and 3.0 GHz HFC plants, are out of scope for this paper and are discussed further in [CLO\_2019].

Choosing the upstream SLA is a more complicated matter. As can be seen with the DS:US consumption ratio, there might be a 20:1 difference between the two. However, in the new 10G era, there may be a need for gigabit US SLA tiers with high burst rates, even if the US consumption might be much lower than downstream.

Looking at PON systems, they offer both symmetric and asymmetric data rates. GPON provides 2.5 Gbps downstream data rates with 1.2 Gbps upstream data rates for a 2:1 ratio. The IEEE 10G EPON downstream might be paired with either a 1G or 10G upstream for 10:1 or 1:1 ratios. In the ITU world, XG-PON pairs 10 Gbps downstream with 2.5 Gbps upstream (i.e. 4:1 ratio) while XGS-PON provides a symmetric 10 Gbps in both directions for 1:1 ratio.

HFC systems have traditionally been extremely asymmetric, but these trends are changing. In the following sections, a range of upstream SLAs are considered to pair with the 8 Gbps DS SLA with a discussion on the technology trade-offs needed for each.

## 1.4. Migrating HFC from DS only 10G to More symmetrical 10G

HFC systems have traditionally been extremely asymmetric with only 42 to 65 MHz of upstream spectrum compared to downstream spectrum on the order of 700 to 1000 MHz. But these trends are changing. [ULM\_2019] and [ALB\_2019] look at technologies that cable operators can use to reach the 10G goals.

An 85 MHz upstream split provides the minimum upstream capacity needed to pair with a 10G downstream. This might support DS:US ratios in the 15:1 to 20:1 range. DOCSIS 3.1 also introduced a 204 MHz upstream split option that enables 1 to 1.5 Gbps upstream SLAs, but it is still around a 5:1 ratio. Getting beyond a 1 to 1.5 Gbps US tier will require additional technologies besides the 204 MHz upstream split.

### 1.4.1. Full-Duplex DOCSIS (FDX)

Some recent work at CableLabs has focused on a new technology called Full Duplex DOCSIS (FDX). FDX leverages echo canceller technology to allow simultaneous upstream and downstream operation in the FDX band. FDX is targeted at a fiber deep Node+0 DAA environment. FDX is now becoming part of the new DOCSIS 4.0 specification [FDX\_PHY].

The FDX capability offers a fundamental benefit that permits upstream spectrum expansions to occur without causing reductions in downstream spectrum. FDX proposes to have downstream and upstream transmissions occurring in the same frequency band at the same time. In the FDX specification, the overlapping frequency bands can be in any of the following ranges: 108-204 MHz, 108-300 MHz, 108-396 MHz, 108-492 MHz, 108-588 MHz, or 108-684 MHz as shown in Figure 1. These FDX bands are in addition to the standard 5-85 MHz upstream that can be utilized as well.




Figure 1 – Full Duplex DOCSIS (FDX) Spectrum Band Options

On a fiber deep Node+0 plant, the upstream OFDMA channel might net capacity of as much as 8-10 Mbps per MHz. This means that a 108-300 MHz FDX system might support 2.0-2.5 Gbps US SLA while

the full spectrum 108-684 MHz FDX system might support 5-6 Gbps US SLA. Using the full FDX band would enable the operator to offer such DS/US service tiers as 8 Gbps x 4 Gbps, 7.5 Gbps x 2.5 Gbps or a fully symmetric 5 Gbps x 5 Gbps SLA.

With the Node+0 architecture, the 108-1218 MHz of downstream spectrum might net over 10 Gbps of downstream capacity which means the 1218 MHz FDX system could be pushed to a true 10 Gbps DS SLA with 5 Gbps US SLA. The downstream is now on par with 10G Ethernet while supporting a 2:1 DS:US ratio.

Current FDX work is moving along well. Initial field trials were completed in 2018 and continued in 2019 with very promising results. Real-world deployments are targeted to take place in 2020.

#### **1.4.1. Ultra-High Splits and Soft-FDX**

A number of operators are reluctant to jump to N+0 but are still interested in achieving more symmetrical upstream service tiers. Some folks are investigating a simpler approach than FDX by just pushing the upstream split even higher. 300 and 396 MHz upstream splits have been discussed as they add one or two additional 96 MHz OFDMA channels to the upstream capacity. However, this approach eats into the downstream spectrum, so is often considered with extending the downstream higher (e.g. 1.8 or 3.0 GHz).

While FDX should work fine in Node+0 environments, its ability to perform in Node+X environments is still under study. The issues with traditional FDX in this environment and possible alternatives are explored in [ALB\_2019]. Because of the FDX challenges in N+X HFC plant, other technologies are under consideration for those scenarios.

A problem called Interference Group Elongation has been identified in [ALB\_2019] that causes serious issues with this “FDX in Node+X” proposal. In the end, it causes large Interference Groups to be created that span most of the length of each RF Leg on a Node. This forces that RF leg to operate in more of a Time Division Duplex (TDD) manner. This has led to some new proposals called Soft FDX.

Soft FDX is a special mode where the network is operated with non-overlapping US & DS spectra (just like today!). The ‘soft’ adjective refers to the ability to change the location of the US/DS split easily (potentially via software). Soft FDX helps in supporting high US speeds, which are occasionally demanded by users, without permanently locking the spectrum to the US which can severely affect the valuable DS spectrum that is used to offer many services including video and high DS speeds which are demanded more frequently than the US.

Soft FDX can be either static or dynamic. Static soft FDX refers to the case where the US/DS split location does not change frequently (e.g., on the order of months or years). On the other hand, Dynamic soft FDX refers to the case where the US/DS split location changes in real time based on traffic demand (e.g. on the order of milliseconds or seconds). For instance, in the dynamic soft FDX case, when there is a need for more US spectrum to upload a large file to the cloud or run an US speed test as examples, the split changes to accommodate that and when the need for the added US spectrum goes way, the split changes back to reclaim the valuable DS spectrum. Both static and dynamic soft FDX can be implemented using special assignment of the FDX RBA messages.

### **1.5. Cable 10G - Summary**

As can be seen by these options, an operator can choose how symmetric they want their system to be. This will be driven by competitive market forces as well as new yet unknown upstream applications that may appear in the future.

After accounting for all the different overheads (e.g. PHY, MAC, IP layers), the subscriber is getting an 8 Gbps SLA in a 10G world. Table 2 summarizes the various options 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 a range of values.

**Table 1 – Summary of 10G Access Network Options**

<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

These SLAs will be quite adequate for the vast majority of subscribers over the next decade. There may be a small number of innovators and early adopters that want to go beyond these service tiers later in the decade, but that can be handled with a Selective Subscriber Migration strategy that moves this small percentage to a next generation PON (e.g. 20+ Gbps) or to an extended spectrum HFC plant (e.g. 1.8 or 3.0 GHz).

Traffic engineering and network capacity analysis in [ULM\_2019] shows that 1218/204 MHz technology meets the needs through the end of the next decade. While getting to fiber deep N+0 is a good long-term strategic goal, a 500 HP node size, N+X system is still reasonable as long as it can be segmented.

If more symmetric upstream services are needed or desired (i.e. greater than 1.5 Gbps), then a migration to traditional FDX for N+0 or Soft-FDX for N+X is a reasonable path. These also restore some downstream spectrum (i.e. 108-258 MHz) that may enable a true 10 Gbps DS SLA.

The investigations into 1 GHz tap technology show that operators can achieve the 10G goals with the existing installed base of taps. This will buy the operator more time before they need to pull the trigger and replace them. Hopefully the 1.8/3.0 GHz future taps will be cost effective by that time.

Finally, fiber deep and Distributed Access Architectures (DAA) become more important technologies at helping operators to achieve the 10G goals.

## 2. Gigabit Wireless Technologies – 5G, CBRS, WiFi-6

Gigabit Fixed Wireless technologies are emerging to rival wired services. 5G will be offered across different spectrums with each band having unique capacity/distance tradeoffs. WiFi-6 and CBRS have also entered the fray.

5G Wireless Networks are being built to serve new applications. These applications must fund the development of the solutions and the infrastructure. Below are four revenue or new business opportunities that will be key to the deployment of 5G Wireless.

1. **Fixed Wireless Access (FWA)** – The connection of wireless broadband to homes or other fixed location services. The trend for bringing Gbps speeds to consumers and enterprise is driving to smaller and smaller cell sites.
2. **Massive outdoor and indoor Internet of Things (IoT) connectivity** – Where everything gets connected. Narrowband IoT (NB-IOT) and existing Category M1 (CAT M1) LTE (Long Term Evolution) services will fill this growing area initially and leveraging technologies such as Long Range (LoRA) using unlicensed spectrum below 1 GHz and LTE-M.
3. **High bandwidth and capacity mobile wireless** – As silicon technology evolves, there will be increased burst and sustained speed applications to mobile devices.
4. **Connected Car and Connected Augmented Reality (AR)** – Vehicle to Vehicle (V2V) and Vehicle to Anything (V2X) will grow with Self Driving Car technology. It may also drive new applications as diverse as digital signage and hands free, eye glasses powered experiences.

Applications alone will not drive the development and deployment of the networks required to carry the additional bandwidth at the speeds required. There are several technical aspects to be considered from the “Outside House In and Inside House Out” perspectives as well as whether the deployments will use “Line-of-Sight and Non-Line-of-Sight” solutions.




**Figure 2 – Automated Driving, Mobility Services, and Augmented Reality**




## 2.1. The Outside Inside Discussion

[FLE\_2017] discussed some outdoor issues with high band wireless to consider including:

- 100' tall poles in subdivisions struggle to get coverage required for Line of Sight solutions
- The attenuation differences between 29 GHz and 2.9 GHz through various foliage and tree types. As an example, millimeter wave penetration through pine trees with hard and angular leaves, and pinecones make it almost impossible to transmit any distance from the Base Station
- The spectral efficiency of millimeter wave is affected severely by foliage and different window types with real problems at higher frequencies like 70 GHz to penetrate through evergreen tree and e-windows with double glazed gas filled panes (see figure 3)
- Rain and snow both influence propagation and reduce the range




**Figure 3 – 28 GHz vs. 70 GHz Spectral Efficiency**



**Figure 4 – 37 GHz System Performance & Spectral Efficiency**

The physics of sub 6 GHz LTE applied in Fixed Wireless Access (FWA) solutions to the home typically requires an outdoor antenna to be used to mitigate 9-35dB signal loss of outdoor and indoor walls. This increases the overall OPEX and CAPEX of the solution yet is typically still good for sparse and large cell site connections. What it gains in coverage it typically loses in speed and bandwidth capabilities.




**Figure 5 – Sub 6 GHz with Outside the home antenna**

The problem is starker with the use of high capacity millimeter wave spectrum where the cell size typically must shrink because of LOS requirements. Even high 100 to 200-foot tower locations still typically require outdoor transmit & receive (Tx/Rx) antenna solutions to work reliably and be economically viable.

To try and increase the bandwidth to the client, the cell sizes shrink to a point where a common architecture now being considered is to add the cell to light poles or other mounting points less than 50 meters from the CPE device. Some CommScope testing has shown that in Single Family Unit Housing (SFU) Estates to achieve Gbps millimeter wave speeds to the consumer 25 meter or less cell sites may be required serving 8 or fewer homes.

These types of solutions require getting to very small cell size and leveraging light poles closer to the single-family unit home. There's a clear opportunity for MDUs. As well as an opportunity for service providers to use their own wired plants.



**Figure 6 – mmWave Success Requires Roof, Window, and Wall Installations**

## Choose Your Weapon – Sub 6 GHz versus Millimeter Wave

5G deployment is likely to be a combination of **both** sub 6 GHz and millimeter wave to get the reliability required for FWA and mobility services.

**Table 2 – 5G Range or Performance, mm Wave vs. sub 6 GHz**

5G mm Wave	5G Sub 6 GHz
24 GHz – 90 GHz	Wi-Fi, LTE / CBRS, LoRA, and NB-IoT
<ul style="list-style-type: none"><li>• Lots of new spectrum</li><li>• Performance &gt;&gt;Gbps – 128 Gbps in 60 GHz</li><li>• Line of Sight – mostly</li><li>• Range does not like walls, windows, or conifers!</li><li>• Reliability is lower</li></ul>	<ul style="list-style-type: none"><li>• Wi-Fi, LTE and new Shared Spectrum CBRS</li><li>• Performance &gt;Gbps</li><li>• Non- Line of Sight</li><li>• Range – Ranges from multiple Km to Wi-Fi sectors</li><li>• Reliability is high</li></ul>

### 2.2. The Inside Outside strategy

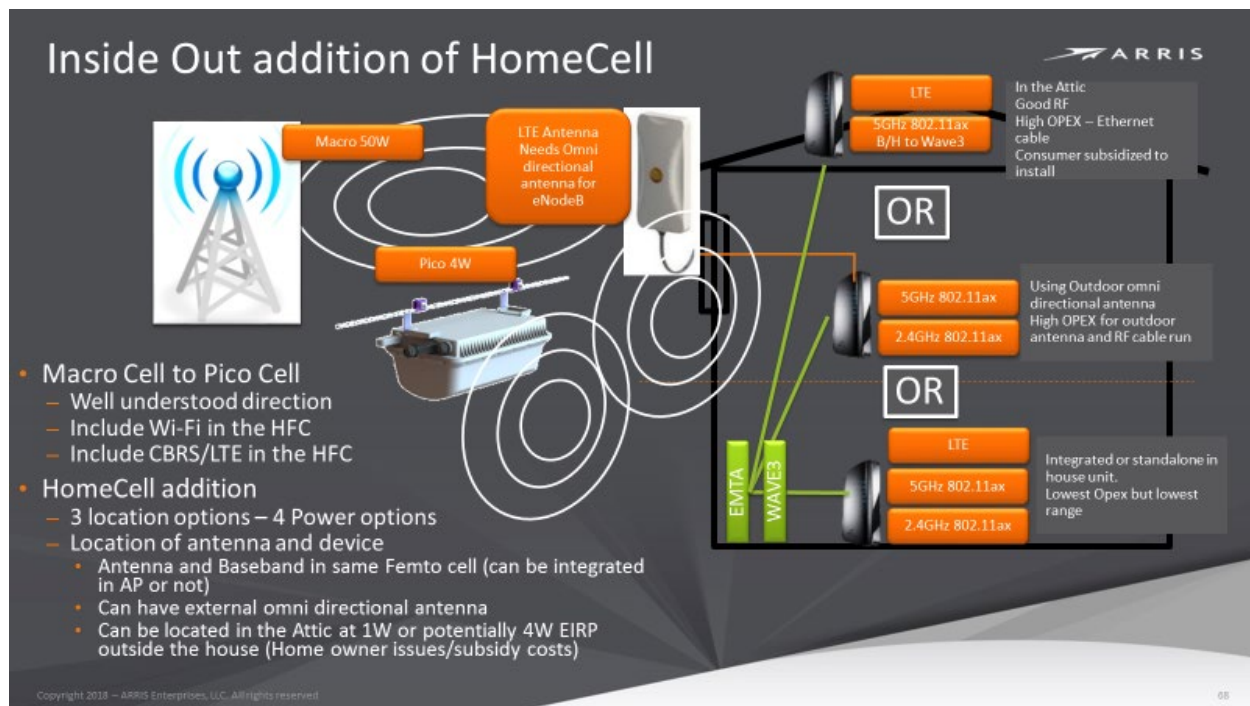
This approach has been discussed recently with the overall direction of smaller cell architectures. The idea is to support a small cell or home cell approach at the edge of the service provider's network. The implementation typically manifests itself as placing a Wireless Home Cell in the consumer's home to support both in and out of home Wireless Connectivity services. It tries to support many convergent applications:

1. In home Wi-Fi and Licensed Spectrum connectivity for best wireless solution
2. Surrogate connectivity for others outside the house leveraging immediate wired backhaul on the consumer gateway device
3. Better HomeSpot/Community Hotspot experience augmenting Wi-Fi with licensed spectrum
4. Easy connect of eSIM or other authenticated and authorized devices from an LTE or 5G Wireless Cell
5. Offload home and close to home bandwidth to lower power home cell from the outside Pico or macro networks
6. To get 5G services 'inside' the home millimeter wave bandwidths do not penetrate deep enough into the home and struggle to connect to one device inside the outer wall/window


Location of the Home Cell is very important. Having it on the outside of the house requires using an outdoor omni-directional antenna. This approach has high OPEX for the outdoor antenna and RF cable run. Putting the Home Cell in the attic provides good RF but may have a high OPEX to run an Ethernet cable to it. Integrating it into the home CPE is the lowest OPEX but also the lowest range.

One short term option that seems to make sense particularly with Shared Spectrum solutions like CBRS in the US is to add a Femto-Cell to the home. This provides the ability to offload the Macro network and provide convenient SIM based authorization for services. The Home Cell could then contain both Wi-Fi and LTE support which can be found today in many LTE deployments.

Using a DOCSIS or Fiber Network to connect Home CBRS LTE cells provides timing and backhaul services. The home gets a CBRS 3.5 GHz LTE Pico Cell. The operator supplied phone connects on a CBRS frequency. There is potential to work with a Mobile Virtual Network Operator (MVNO) partner. And there's future potential to offer Neutral Host Home Cell – and offload the MVNO partner macro network.



**Figure 7 – Schema of Inside/Out Propagation Studies for 3.5 GHz CBRS**



**Figure 8 – Sub 6 GHz Connections to the Home**

## The Inside Outside Discussion – Home Cell

[FLE\_2017] shows there is a choice of power for LTE radios. Two possible examples to consider:

- 125mW EIRP which can cover the home and ~10m outside the home – affording connection to handsets on the street or better performance at range for the consumers LTE enabled devices
- 1W EIRP which can cover the home and 3+ neighboring homes as well as good street coverage

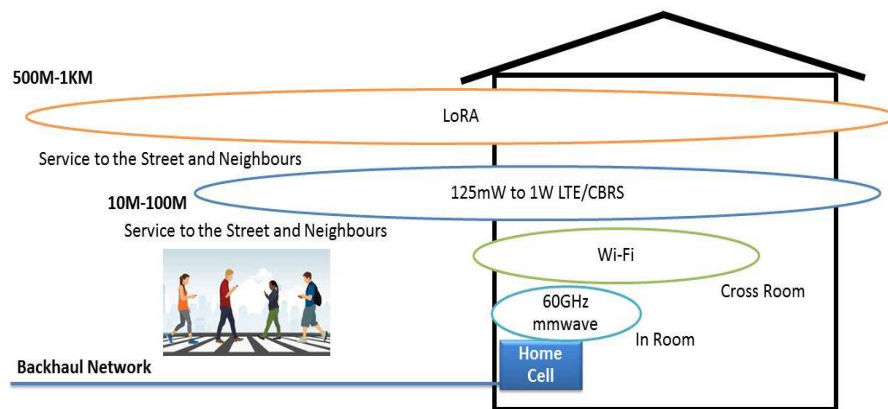
Decisions on what power device to use are determined by the cost, ergonomics of size and complexity of managing the overlapping cells.

Another common technology being explored for Inside Out Networks with limited BW SLAs is LoRA<sup>®</sup>. It runs in the 900 MHz frequency and affords a low-cost Macro Narrowband IoT solution from the 100

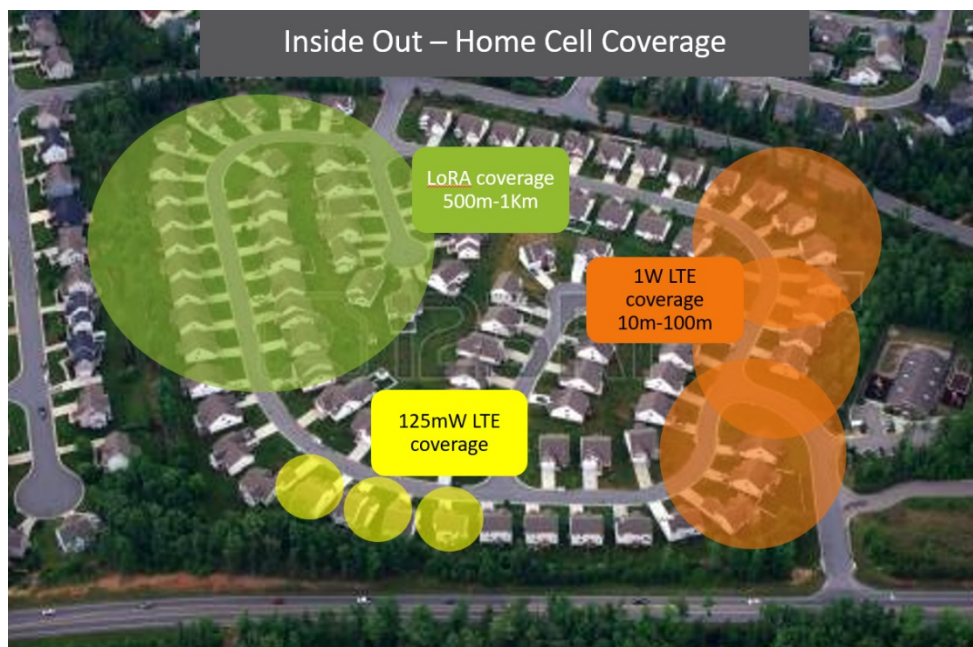
feet + Monopole – and is also capable of being used from the home. A low cost LoRA addition to the home can create a network of connections from 500 meters to 1 kilometer from the home depending on placement of the LoRA device in the home (see figure 9).

Here are some questions to consider regarding Figure 9:

- Can a 5G cell operate inside the home?
- Which Sub 6 GHz opportunities like CBRS 3.5 GHz Cell and LoRA can provide outdoor coverage?
- Millimeter Wave potential, but at what cost/range?
- Can 60 GHz networks in the home develop >>1 Gbps intra room wireless solutions by leveraging 802.11ay<sup>®</sup> technology.



**Figure 9 – Networks to the Home Cell**



**Figure 10 – Home Cell Coverage**

AP Service Radius (feet) 25 Mbps (D) / 5 Mbps (U)					AP Service Radius (feet) 10 Mbps (D) / 2 Mbps (U)					AP Service Radius (feet) 100 kbps (D)/100 kbps (U)							
1W AP	Size (sq ft)	Exterior	Stone	Brick	Wood	1W AP	Size (sq ft)	Exterior	Stone	Brick	Wood	1W AP	Size (sq ft)	Exterior	Stone	Brick	Wood
	1500		81	217	494		1500		135	361	822		1500		267	715	1627
	2500		63	170	386		2500		105	282	643		2500		208	559	1272
	5000		63	170	386		5000		105	282	643		5000		208	559	1272
500mW AP	Size (sq ft)	Exterior	Stone	Brick	Wood	500mW AP	Size (sq ft)	Exterior	Stone	Brick	Wood	500mW AP	Size (sq ft)	Exterior	Stone	Brick	Wood
	1500		63	170	386		1500		105	282	643		1500		208	559	1272
	2500		49	132	302		2500		82	221	502		2500		163	437	994
	5000		49	132	302		5000		82	221	502		5000		163	437	994
250mW AP	Size (sq ft)	Exterior	Stone	Brick	Wood	250mW AP	Size (sq ft)	Exterior	Stone	Brick	Wood	250mW AP	Size (sq ft)	Exterior	Stone	Brick	Wood
	1500		49	132	302		1500		82	221	502		1500		163	437	994
	2500		39	104	236		2500		64	172	392		2500		127	341	776
	5000		39	104	236		5000		64	172	392		5000		127	341	776

**Figure 11 – LTE Range versus Performance for Inside Out Coverage**

Figure 11 shows the inside out coverage of LTE small cells in the home and their potential roaming contribution at different LTE data and voice support levels. A potential visualization of an LTE based Home Cell generated network in a housing estate is given in Figure 10. The potential does exist to get high enough coverage across homes to develop services that don't need complete coverage and can also handoff to macro networks of lower speeds or available capacity.

[CHE\_2019] also provides results on the overall performance in a housing development with 156 homes with approximately 0.75 acre lots where the house construction is mostly brick exterior showing the level of expected 3.5GHz performance with the foliage, terrain and house types. The position of these small cell devices is not inside the home but are assumed to be located on the roof of the homes. This model was run to see if the attenuation of the indoor walls and furniture was removed would a 1W EIRP outdoor home mounted cell have a much better contribution to coverage (which it does).

Equally for something like the CBRS standard in the US – where 4W EIRP is also permissible – running the 4W model also showed good results. This strategy of using the Home as a Tower – is potentially useful when looking at a strand mount Small Cell strategy using aerial Hybrid Fiber Coax or Optical and a substantial portion of the network may be underground.

If the exterior AP power is raised to the CBRS allowed maximum of 4 W, the HaaT RF coverage permits even less density. However, the per-user data coverage now might become limited by the number of roaming customers (data pipe-sharing) as opposed to bitrate throttling (loss of spectral density):

From millimeter wave in the home it is not practical to have a Home Cell support outside home connection. However, with the recent approval of the 802.11ay standard, the 60 GHz spectrum can be utilized in the home to support higher speeds and leverage millimeter wave solutions. 60 GHz can certainly be used in room but with sheetrock (-7db to -9dB) dividing walls it can also be used across rooms. 802.11ad devices (rated to 1.5 Gbps) can generate higher than Wi-Fi capacities in room and even across rooms.

In actual testing performed by ARRIS/CommScope of 60 GHz in-home propagation with 802.11ad, there was higher than Wi-Fi performance in much of the home. Typical issues with millimeter wave were reflections from 'hard' surfaces like TVs and Porcelain in bathrooms. See Figure 15.





Scale: 200 ft/61 pels = 3.28 ft (1m/pel)

With 1W external roof mount and ~10-12 dB of foliage attenuation, expect ~400 ft service radius for 25/5 service

Total area shown: 136.3 acres  
Total homes shown: 158

The 1W wireless service group for this home clustering seems to vary from 10-16 homes

**Figure 12 – 1W HaaT Coverage for 25/5 Service (1 x 10-16 Homes)**



Scale: 200 ft/61 pels = 3.28 ft (1m) / pel

With 4W external roof mount and ~ 10-12 dB worth of foliage attenuation, expect ~ 700 ft service radius for 25/5 service.

Total area shown: 136.3 acres  
Total homes shown: 158

Roof-top mounting easily supplies > 20 homes per mount roaming support

**Figure 13 – 4W HaaT Coverage for 25/5 Service (1 x 20-40 Homes)**

Service Reach of 4W strand  
mount to mobile @ 3.5 GHz

25 / 5 Mbps  
10 / 2 Mbps  
100 / 100 kbps




Figure 14 – Service Mounting Potential for Outdoor 4W CBRS Mesh


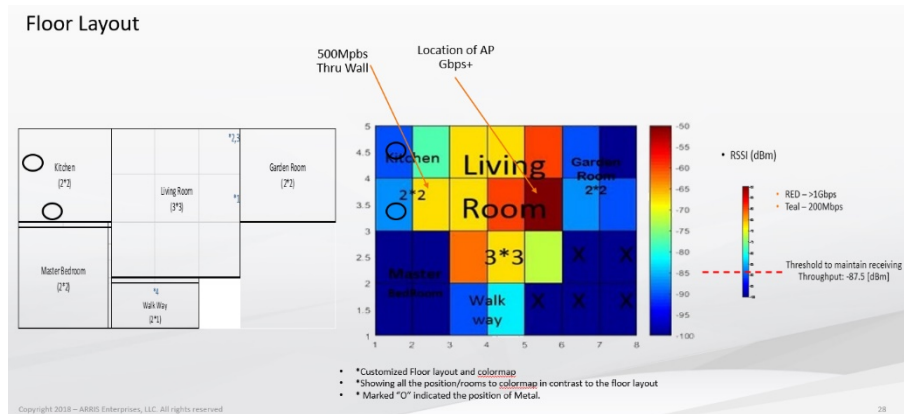


Figure 15 – In Home/Office 60 GHz Coverage



## 2.3. 5G as a Fixed Wireless Access solution

For wireless communications, this is an unprecedented time. More licensed and unlicensed radio spectrum between UHF whitespace and millimeter wave mega-block partitions is being made available for commercial interests to invest in and grow business services than at any single prior point in history. The use of millimeter wave spectrum has sparked many debates about its architecture and economics — given the physics restrictions of primarily requiring “Line of Sight” to deliver the promise of multi Gigabits of wireless delivery. It is this non-determinism of signal propagation that has generated lots of research, innovation, and testing of solutions to create and define a deployable architecture that will support both Fixed Wireless Access and mobility uses.

Can a Fixed Wireless Access solution be developed to compete with or augment the wired broadband solutions today? [FLE\_20017] looked at using 5G as a Fixed Wireless Access (FWA) solution. It considered the available spectrum options for delivery of a reliable, high-bitrate wireless connection over the last few hundred meters of front-haul as an alternative to wired connections. These are the cases where a newcomer wants to overlay incumbent, existing greenfield opportunities, or CAPEX considerations render the latter alternative unsound. It must leverage of the best attributes of near-line-of-sight (nLOS), non-line-of-sight (NLOS), and line-of-sight (LOS) signaling.

### 2.3.1. FWA Economics

The economics of Fixed Wireless access lie somewhere in the following parameters:

- Cost of the spectrum used: There are licensed and unlicensed bands to consider.
- The size of the cell for bandwidth distribution
  - With the migration towards high bandwidth low latency wireless broadband services — the size of the traditional cell size is likely to reduce substantially.
  - This is a function of propagation and distance of technologies like millimeter wave as well as a requirement to provision speeds of Gbps burst levels
- The backhaul distribution and connectivity
  - As the cell size gets smaller it still should deliver multiple Gbps to enough customers to make it economically viable. This makes the backhaul to the cell important for speed and scaling to meet the front haul costs to consumers.
  - The ideal solution is that a fiber connect to every Gbps capacity Small Cell for 5G — likely at least 10 Gbps.
  - DOCSIS and Line of Sight wireless are other backhaul options.
- The cost of the CPE equipment
  - This is one of the main barriers and inertia contributors to using Fixed Wireless Access. It includes costs of outdoor transceivers and external antenna with ergonomic impacts too.

While CBRS 3.5 GHz sub 6 GHz solution is not a solution typically targeted at 5G, there is the likelihood of 5G being a dual PHY or dual standard technology. Millimeter wave is not deterministic in its performance due to the environmental and NLOS issues — and the economics of deploying to the worst condition don’t work, therefore there may be likely solutions that provide Small Cells with both Millimeter wave and sub 6 GHz as well as CPE that support both PHYs. A dual PHY solution will make the solution more expensive, larger in size and higher in power consumption.

Significant roaming coverage gaps for premium data services begin to occur once exterior materials approach the density of brick and become unusable for cases where home placement become spaced by large lots – unless the coverage is augmented by exterior high-power APs.

This leads to a solution where the separation of interior-home and neighborhood roaming coverage by employment of a scaled picocell internal AP in each home (to accept handoff of the mobile from its outside roaming) and a network of outside mast or second-floor mount APs of either 1 or 4 W power (using the acronym “Home as a Tower” or HaaT) every N houses to provide the “outside home” (neighborhood roaming) data coverage.

### **2.3.2. FWA Summary**

The implications of providing a future-proofed wireless bitrate capability to all subscribers beyond the reach of wireline in a cable system requires the analysis of wireless delivery options which include LOS, nLOS and NLOS systems — each of which comprises a mixed bag of capability and compromise. The broader bandwidth of millimeter LOS delivery, with its promise of massive MIMO antenna structures on both base station and client endpoints, unfortunately burdens itself with compromises involving client-side signal recovery costs, short signal throw, aesthetic challenges and perhaps non-deterministic link quality.

nLOS and NLOS sub-6 GHz systems can be made to overcome these challenges. However, the available bandwidth puts considerable pressure on massive MIMO and signal processing upgrades on the base station side to create the scalar benefits which effectively multiply spectral efficiency to levels necessary to anticipate user bitrate consumption a mere 4-5 years in the future. The relentless bitrate consumption growth defined in Cloonan’s Curve suggests that, ultimately, the facility of sub-6 GHz NLOS will be associated with a redundancy role for more LOS-based delivery — or perhaps in an ad hoc augmentation role for temporal housing arrangements.

For MNOs, those that don’t own wired broadband networks, the use of FWA is an opportunity to cherry pick areas for a Fixed Wireless overbuild. Some Wireless ISPs already offer millimeter wave broadband delivery services targeting dense areas with only one incumbent, areas where consumers are deprived of choice of broadband provider. The investment scale which nourishes those shared wireless technical advances applicable to both unlicensed and dynamically licensed space for MSOs (cable, telco, and MNO alike) means that applications of FWA will emerge as we move to mobility on 5G systems.

The economics and the size of the optimum cell is still under debate. What is clear is that the easier direction for FWA is dense MDU environments targeting a single wireless connection to the outside and using other solutions internally, like Ethernet, MoCA and Wi-Fi. The Residential 5G deployments will only emerge driven by the rise of mobile 5G devices which will happen in 2021 at scale and will then see the 5G small cells deploy in ever decreasing cell sizes.

## **2.4. Guiding Wireless Trends for Tomorrow’s Smart Home**

What changes will dominate the home wireless services market in the next five years? Several fundamental trends are emerging that will guide device and solution requirements for next-generation homes. Some driving forces behind these shifts include:

- Higher WAN speeds with Gbps burst modes that will be reflected in home Wi-Fi connections.
- Convergence of outdoor 5G wireless services with the indoor overlap of Wi-Fi connections.
- Wireless extenders ensuring 500+ Mbps potential coverage to all points of the household:
  - Introducing tri-band devices, which use tri-band concurrent (TBC) 2.4 GHz 4×4 and 5 GHz 4×4 in homes that require more premium Wi-Fi services and even higher throughput speeds, ensuring Gbps speeds to all end clients.

- By 2021, the emergence of 6 GHz Wi-Fi applications in the UNII-5 to UNII-8 bands to support high-capacity projects that can be driven by service providers. They include:
  - Employing 6 GHz channels in Wi-Fi backhaul applications
  - 6 GHz STB solutions that enable reliable 8K video applications
  - Utilizing 6 GHz Wi-Fi in 5G femtocells for reliable connectivity to the primary gateway.
  - Superior Wi-Fi performance, typically 2×2 in Wi-Fi 6 and up to 4×4 in Wi-Fi 5.

Service providers can complete the new wireless and smart home by considering other trends like LTE™, CBRS, LoRA Alliance™, and 5G. There are several ways to leverage a home — particularly one at the end of a cable or fiber network — and add these technologies to the backhaul of the wired system. Providers can add femtocells connected to the primary fixed wire gateway or access point to provide in-neighborhood or in-home services. They can also offload LTE and 5G handset usage for the home to Wi-Fi, but the licensed spectrum is better for larger homes.

Current 5G applications are focused more on FWA and outside mobility. However, there have been discussions about trying to extend the range of 5G NR signals or applying 5G femtocells in the home. While sub-6 GHz 5G has the scope to be effective for residential use, millimeter wave 5G will suffer high attenuation and propagation loss that may make 6 GHz Wi-Fi the better home high-speed connection.

One view is that the 60 GHz WiGig (802.11ay) standard offers an unlicensed solution to millimeter wave usage in the home that could drive its use for 5G services in a more universal and standardized fashion than licensed frequency band femtocells.

In the next three-to-five years, wireless smart homes will drive new high-speed services by leveraging Wi-Fi 6, 6 GHz and 5G NR and the standards work being done across IEEE®, 3GPP®, Wi-Fi Alliance®, and others.

As cable operators move Fiber Deeper going to an all passive coax network, the ability to deliver multiple Gbps of capacity to a single home seems an easier path than building out an FWA millimeter wave architecture. However, given that 5G POP/Small Cells require wired backhaul, the potential for the MSO to leverage its network for mobile 5G seems to be a more complimentary investment. In discussions with MSOs, who are also MNOs, they struggle now to see an FWA solution to deep residential deployments. They see some potential to use their network to potentially lead to target MDUs served predominantly by their competitors, and often see value in pulling fiber. However, they do see the value of adding 5G and CBRS POPS to their HFC and growing fiber networks for outside mobility applications.

It will be the Wired Backhaul Network that supports these ever-decreasing 5G small cell architectures. Why? Because...

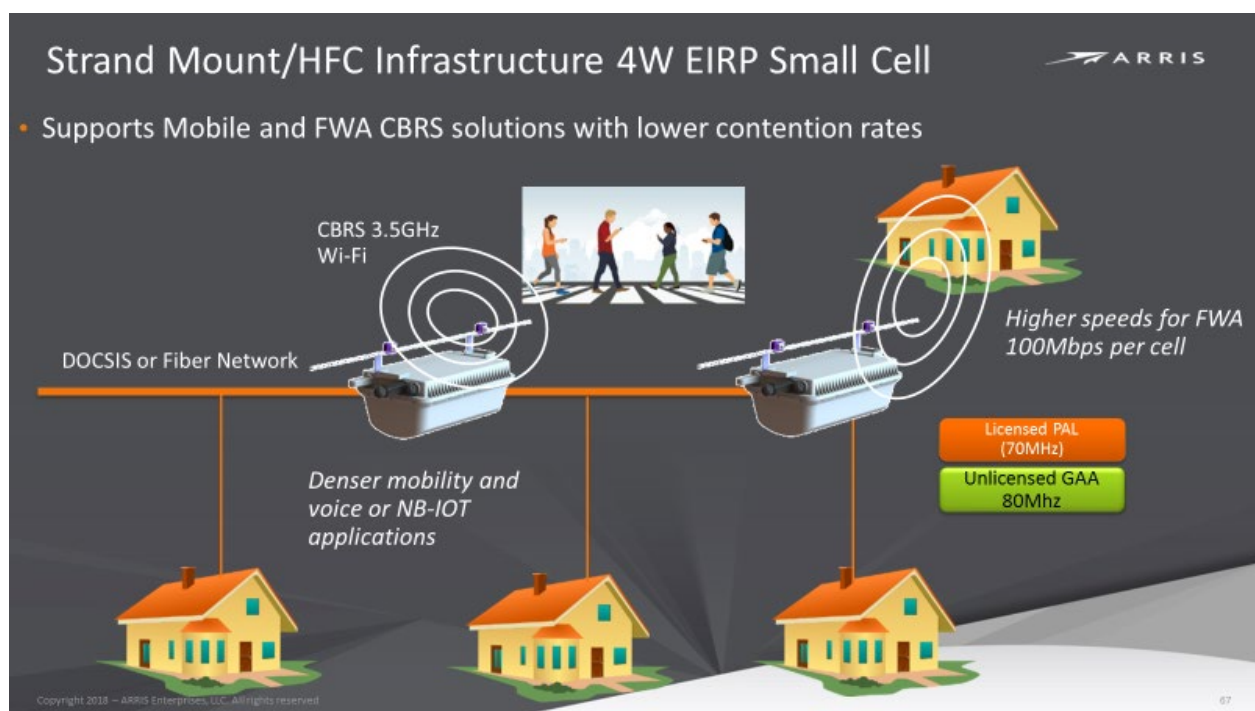
- 5G cannot exist without a Wired Network to backhaul it!
- It's important to go to ground wire as soon as possible to minimize latency
- Industry multiple network operator and service provider consolidation are driven by this requirement and convergence

Because of the latency requirements for 5G Wireless, and the availability of spectrum, it will be increasingly likely that 5G architectures will be many and varied. The likelihood of any sub 6 GHz spectrum being used for 5G services is high and the potential for architectures to have to support both sub 6 GHz for range and millimeter wave for bandwidth is highly probable. There is also scope for adding inside out schemes for leaf Home Cells on wired networks. Whether these deploy as first stage or final stage elements of the new 5G Wireless world, we will see in the coming five years of rollouts.

### 3. Using Outside HFC Plant for 5G – considerations and logistics

Across all the various wireless options just discussed, there is a driving need for much smaller cell sizes. To make this happen requires an infrastructure that supports both *the power and the backhaul* to the small cells. The cable industry Hybrid Fiber Coax (HFC) networks is ideally positioned to support this. The HFC networks might support the addition of attached in-line small cells at various demarcation points on the HFC plant. These cells can be added to the DOCSIS network to support 5G, Wi-Fi and/or CBRS/LTE over the HFC.

In addition to inside/out possibilities discussed above, figure 16 shows a backbone of 4W strand mount POPs used to extend coverage for the case of neighborhood roaming. This lies within the FCC guidelines of 10W maximum for urban areas and 50W for unpopulated rural tracts.




**Figure 16 – Defining the Roaming (Home Exterior) Proposition for CBRS/LTE**

The Coax plant also supports powering for attached devices and has built in backhaul capabilities making it a potential low CAPEX and OPEX small cell host. Possible small cell locations include the Fiber Node, Amplifier locations or even at the Tap distributions to 12 or fewer single unit, stand-alone, houses. Expect to see a mix of high monopole and top of MDU cell sites and smaller cell in fill sites to connect to revenue generating services.

#### 3.1. Case Study – HFC Power for Wireless Cell sites

While backhaul challenges seem to generate the most industry discussion, at the end of the day it is the ability to power these small cells that will become very critical. Today's HFC plant has already put into place an extensive powering infrastructure for enabling active devices out in customer's neighborhoods. This could become one of the cable industry's key assets going forward. An example is reviewed to show how much power might be available for powering these wireless small cells.

Figure 17 shows powering details of an example HFC plant. In the figure, pink line represents the fiber link from the headend to the fiber node; black lines are hardline coax plant, blue triangles are RF amplifiers, blue circles are RF splitters, and the orange diamonds are RF taps, with furthest tap at ~3,000 feet from the node in this N+3 example. Two 90V line power supplies, each capable of delivering 15A of current draw, feed the total of 22 field actives comprising: one fiber node and 21 RF amps. Voltages and the current draws are also shown, with the power consumed totaling ~1.1 kW, with ~35W (~3.2%) of the total due to the Ohmic cable loss.



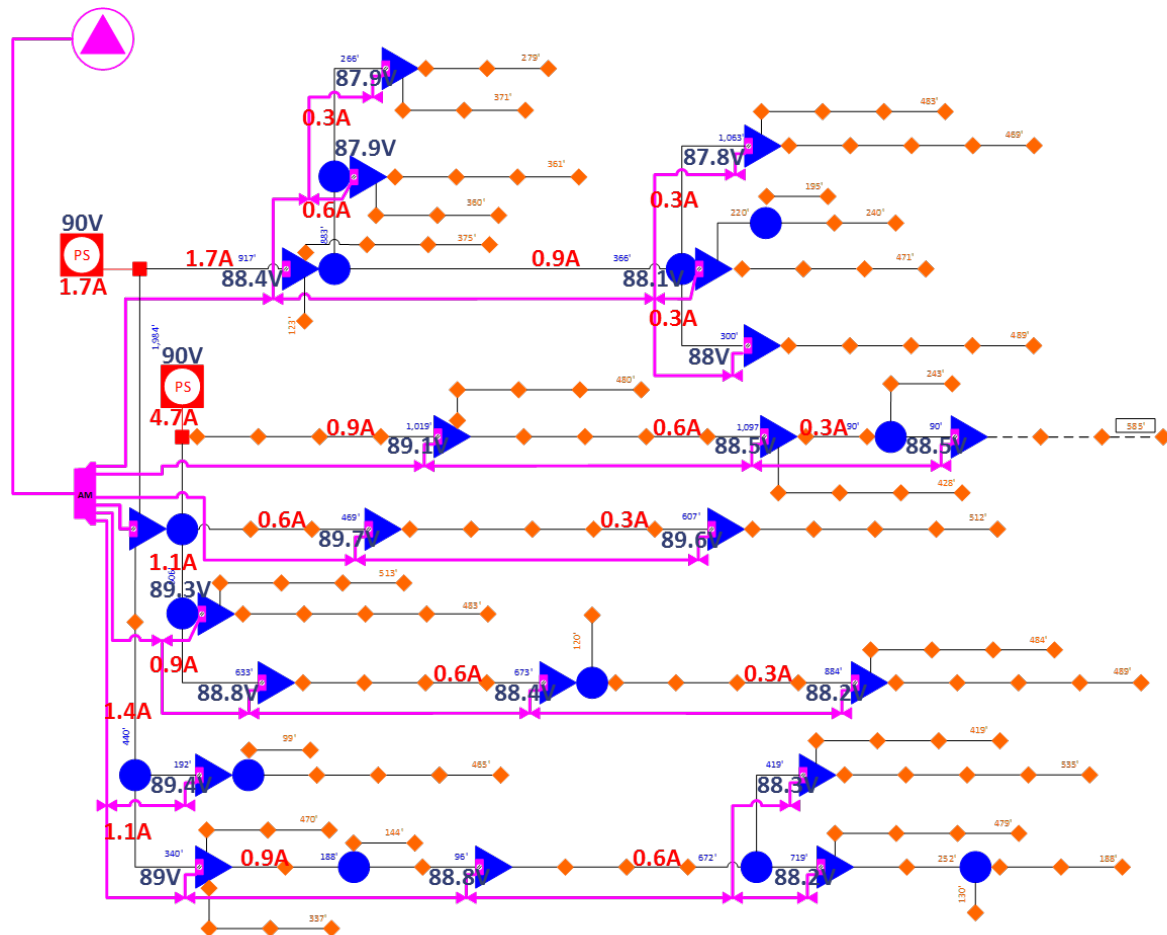
**Figure 17 – A N+3 HFC plant example with Power Supplies, Voltages and current draw**

In this example, there is ample power available to add wireless microcells to the HFC plant. The two 90V Line Power Supplies can typically deliver 15A each and are using much less than that. The power supplies in N+3 plant in Figure 17 are only at ~40% utilization, using 4.0A and 8.5A respectively. There is a significant power budget available for other purposes such as wireless 5G cells.

So, while some HFC plants may be fortunate enough to have an abundance of power, there are many other HFC plants that are close to maxed out. What are they to do? There are other factors at play as well. In a typical USA HFC plant, those RF actives are often either 750 or 860 MHz and were last touched in mid-1990s to early 2000s. It happens that this is when most of the originally CATV video-distribution intended networks were upgraded to the two-way HSD-capable HFC plant. The plant age itself may pose

network reliability issues. That along with an ever-increasing customer bandwidth capacity demand and competitive market mean these aged HFC plants could be due for a refresh real soon. If the HFC plant is going to be refreshed, it should be done with an eye toward improving powering as well.

Many view Fiber to the Premise (FTTP) as the ultimate upgrade, but this can be prohibitively expensive. Nevertheless, due to various tradeoffs, some variants of refurbished actives, node splitting, cascade depth reduction and getting fiber-deeper into the network is what often gets implemented. Per [Ulm\_2016], one of those upgrade paths is a fiber-deep, fiber to the last active (FTTLA) approach that yields an upgraded plant shown in Figure 18.



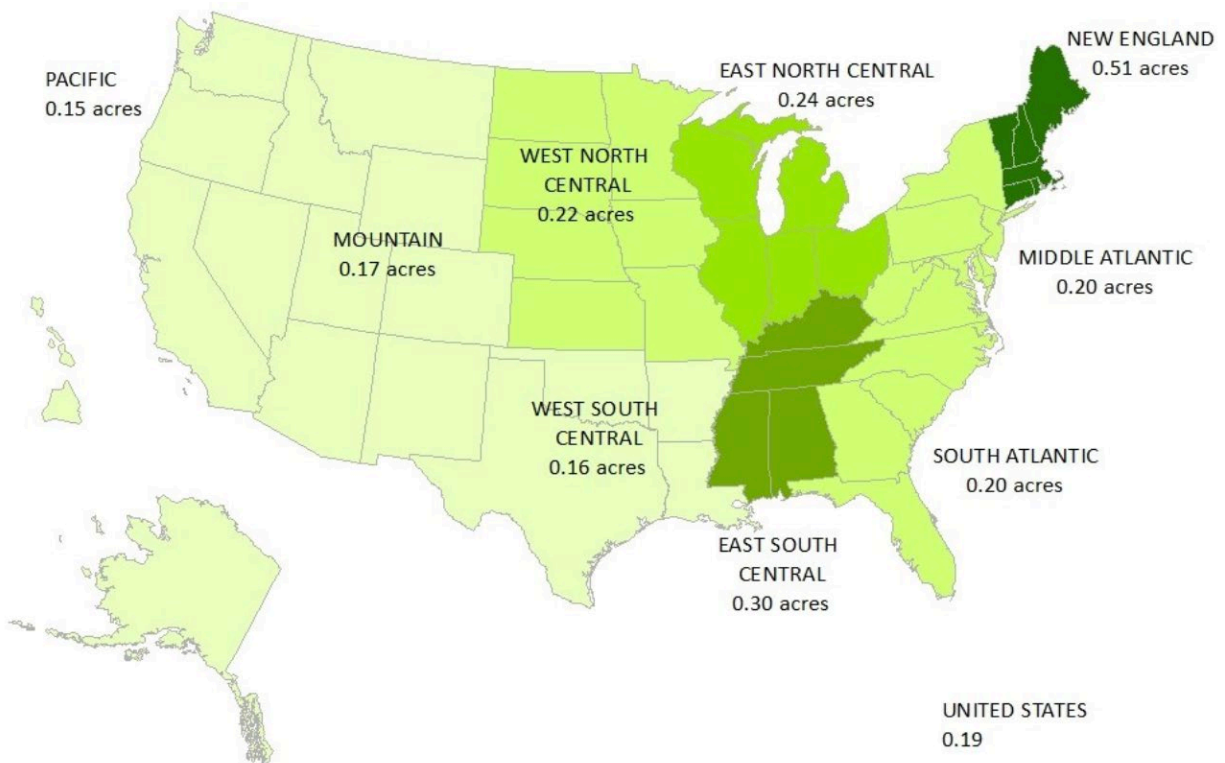
**Figure 18 – N+3 HFC plant upgraded to FTTLA, with Voltages and current draws**

For FTTLA upgrade, the original node location gets converted to a fiber in / fiber out “VHUB aggregator”, all the RF amplifiers that are feeding customer premises get converted to small fiber nodes, with fiber overlashed to each and every node; where the RF line extenders not feeding customers are no longer required. Total power consumption of such upgraded plant is estimated at ~0.57kW, with only 8W (~1.4%) due to the Ohmic cable loss. This is *almost half* the power consumption of the original N+3 plant in Figure 17. The power supplies in Figure 17 are using 4.0A and 8.5A respectively while the FTTLA plant in Figure 18 is using an even smaller fraction of it with 1.7A and 4.7A respectively. So, for HFC plant that is short on power, doing a fiber deep upgrade can help find that extra wattage to power those wireless cells. And pulling fiber deeper will help with the backhaul as well!

### 3.2. Mapping 5G Small Cells to Existing HFC Plant

Densification is a term often used when planning 5G implementations. This is being driven mainly due to the bandwidth demands and the propagation limits of millimeter waves. As seen previously in figures 10 to 14, densification means getting wireless access points into the neighborhoods. Considering existing infrastructure for delivering bandwidth to neighborhoods, what is more appropriate than HFC as the backbone for 5G? Thus, let's look at some HFC network density characteristics next and see how it might map to a 5G deployment.

Homes-Passed (HP) per mile is a metric that is most-often used to describe HFC network density. It is primarily driven by the density of “developed environments” often classified as urban, suburban and rural. Geography plays a significant role. Figure 19 shows how median lot size of new single-family detached homes sold in 2015 varies across the USA regions. Lot size of 0.19 acres across the whole country is the smallest median size on record since 1992, when Census Bureau Survey of Construction started tracking these statistics [NAHB\_2018]. The regional difference can vary by more than three times.




Source: 2015 Survey of Construction, NAHB Estimates

**Figure 19 – Median Lot Size, new single-family detached homes started in 2015**

The same reference [NAHB\_2018] points to a handy analogy in how to envision this size: almost 6 such lots, each ~53 x 160 feet large, would fit in between the goal lines of a 300 x 160 feet football field, itself ~1.1 acre in size. Albeit quite illustrative of regional variability, this turns out to be but a single data point, and does not even include statistics for the custom homes, which tend to have larger lots. Furthermore, older homes are likely to be on even larger lots. From this analysis, one can intuit that across-the-country median size lot gravitates towards a fraction of an acre – 1/5<sup>th</sup> on the lower end and perhaps 2/3<sup>rd</sup> on the higher end, without a precise knowledge of what that fraction exactly is.



Figure 20 shows that, even for a particular type of settlement and for a particular region there is an additional level of variability to consider. This particular example is from a suburban city in New England. In figure 20, lots on the left-hand side are approximately 1 acre in size, lots in the middle ½ acre and lots on the right-hand side are ¼ of acre or smaller. This clearly demonstrates that an analysis cannot just focus on the average or median lot size, rather it must consider a range of sizes. Our analysis looks across the variance of different lot sizes to understand what happens in particular sized developments.



**Figure 20 - New England suburb, illustration for variability of lot sizes within**

The analysis starts with some theoretical calculations on HFC and wireless densities to see how they align. Some HFC densities are shown in table 3. In the table, residential lot size is represented by a range: from ⅛<sup>th</sup> of an acre all the way to 2 acres. Under assumption that a typical lot is rectangular in shape with roughly 1 x 2 ratio of its sides, the frontage of such lots, then, varies from ~50 feet for the smallest ⅛ acre lot, all the way to ~210 feet, for the 2 acre lot.

The HFC world typically uses linear measure of homes-passed per mile (HP/mile) of hardline coaxial cable to measure and quantify network density. The “lot frontage” length is used to calculate HFC network density by finding how many lots are passed by a hypothetical 1-mile length of a plant. Column 3 of table 3 shows those results. Interestingly, the mid-point of 100 HP/mile roughly represents the average density of all US based HFC plants.

**Table 3 – Linking lot size of Homes-Passed to the HFC and Wireless plant density stats**

Lot Size, Acre	~Lot frontage, feet	~# of Homes-Passed per plant mile	~# of H-P per square mile	~# of HFC miles per square mile
0.125	50	210	3,800	18
0.250	75	140	2,100	15
<b>0.500</b>	<b>105</b>	<b>100</b>	<b>1,100</b>	<b>11</b>
1.000	150	70	600	9
2.000	210	50	300	6



In these examples, from low to high density, HFC plant is mapped along the streets in a linear / one-dimensional way – that is, a coaxial hardline passes along the streets and serves homes on both sides of the street (this element is used in calculation of HP/mile, in that the 1 mile = 5,280 feet long plant passes by ~50 lots 105 ft wide lots on one side of the street, and as many on the other side – for the total of ~100 HP/mile – as in the middle row of the Table 3).

Wireless plant density metrics, however, are two-dimensional, as is the wireless reach of a cell radius (i.e. circle) around a wireless access point, shown previously in figures 10 to 14. The 4<sup>th</sup> column of Table 3 links the two-dimensional density, expressed in Homes-Passed per square mile, to the lot size, lot frontage length, and the linear density metrics of HP/mile. An assumption made here is that town blocks are formed by 4x2 = 8 lots, surrounded by ~33 feet wide streets. For ½ acre lots, with ~105 feet frontage, the area of such 4x2 town block is:

$$(33 + 4 \times 105) \times (33 + 2 \times 207) = 202,491 \text{ square feet, or } 1/138 \text{ of a square mile}$$

Since there are 4x2 = 8 HP per each block, area density calculates as 8 x 138 = ~ 1,100 HP/ mile square, as in the middle row, column 4 of the Table 3. Column 5 of table 3 shows HFC plant length per square mile, required to pass by each lot of a given size.

Table 4 now links HFC node area size, in number of Homes-Passed and in number of nodes per square mile to the number of wireless cells per that same area, with variations for plant density, from 50 HP/mile to 210 HP/mile, as well as for the wireless cell reach radius of 100, 150 and 200 meters.

An assumption of node size, in number of Homes-Passed per node, is made and shown in column 2 of Table 4. Another assumption, of ~7.5 RF actives per plant-mile for the high-density case and linearly down to ~5 RF actives per plant-mile for the low-density case yields typical number of RF amplifiers per node area, as shown in column 3 of Table 4. Column 4 shows how many such nodes fit in a square mile. The last 3 columns calculate how many wireless cells are needed per node area for 100% coverage for each of the plant densities, with radius of coverage of R=~100m, ~150m and ~200m respectively.

Looking at an average-density HFC plant with 100 Homes-Passed per mile (i.e. row 3 in table 4), it has a node size of ~350 HP per node with ~20 RF amplifiers per node. A cell radius of ~150m will require ~18 wireless cells for the 100% coverage of the node area. This is slightly less than the ~20 RF amps but raises the question as to whether the amp location maps well for the wireless cell sites. Note that if the wireless cells can only manage a 100m radius, the number of wireless cells is now double the numbers of RF amplifiers.

**Table 4 – Linking HFC node area size to the number of required wireless cells per node**

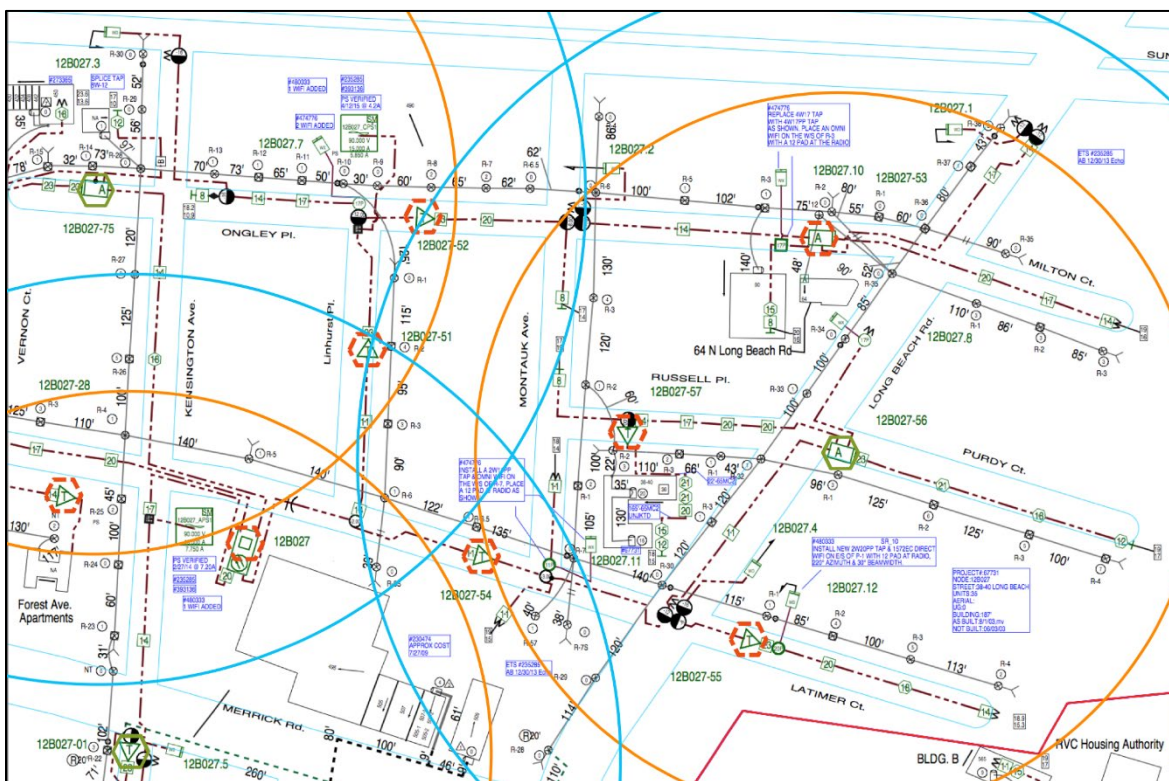
~# of Homes-Passed per plant mile	~HP/node	~# of RF amps/node	~# of nodes/ square mile	~cells/node if R=~100 m	~cells/node if R=~150 m	~cells/node if R=~200 m
210	550	20	7	18	8	5
140	450	20	5	27	12	7
<b>100</b>	<b>350</b>	20	3	41	<b>18</b>	11
70	250	20	3	52	23	13
50	150	15	2	65	29	17

For the high density, 210 HP/mile example in row 1, there are more RF amplifiers than wireless cells, even at 100m cell radius. But at the low density, 50HP/mile example in row 5, there are more wireless cells than RF amplifiers even with 200m radius. So, this theoretical analysis clearly shows that the mapping of wireless cells to the HFC will be extremely dependent on the housing density.

Our next step is to look at some actual HFC plant examples and to map some wireless cells to them to see how the real world holds up. Figures 21-23 show a higher density HFC plant example that is ~190 HP/mile. In figure 21, wireless cells with 200m radius are overlaid and placed next to an active amplifier location as to achieve as close to 100% coverage as possible. These RF amp + cell locations are highlighted with a green hexagonal. Other amplifier locations without a wireless cell are shown with a red hexagonal. Since wireless capacity is a function of distance, figure 21 also shows the 150m radius where homes might achieve higher capacities.

Note that figure 21 represents approximately 2 total cell sites (i.e. one full and two partial cells shown for this particular neighborhood), but that there are also 8 RF amplifier locations without a cell. Also note that most homes will be inside the 150m radius and get higher wireless capacities. This result is also reasonably close to our theoretical analysis in table 4.

Figure 22 looks at the same high-density neighborhood, but now with 150m cell radius overlaid for close to 100% coverage. The figure also shows the 100m cell radius where higher wireless capacities can be expected. Figure 22 shows almost 3 full cell sites co-located with RF amps, while 7 RF amplifier locations are without a cell. Again, reasonably close to our theoretical analysis in table 4.



**Figure 21 – Higher density (~190 HP/mile) HFC overlaid with ~200m/150m radius cells**



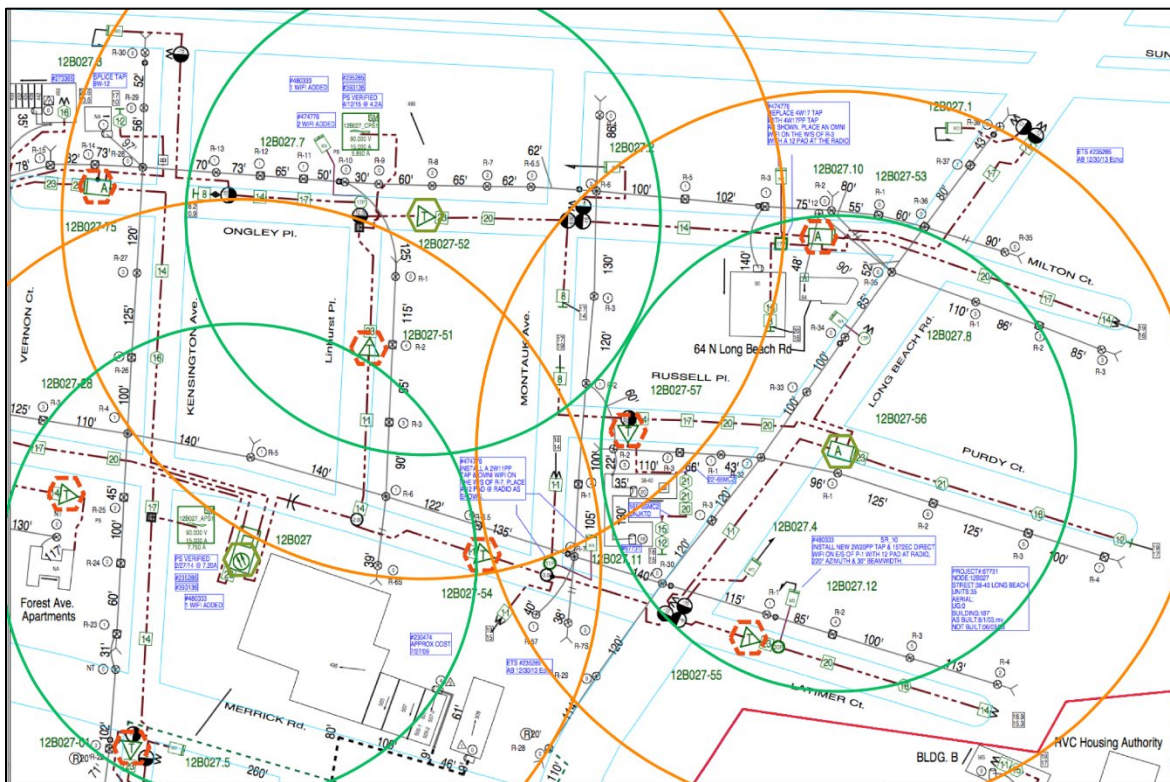


Figure 22 – Higher density (~190 HP/mile) HFC overlaid with ~150m/100m cells

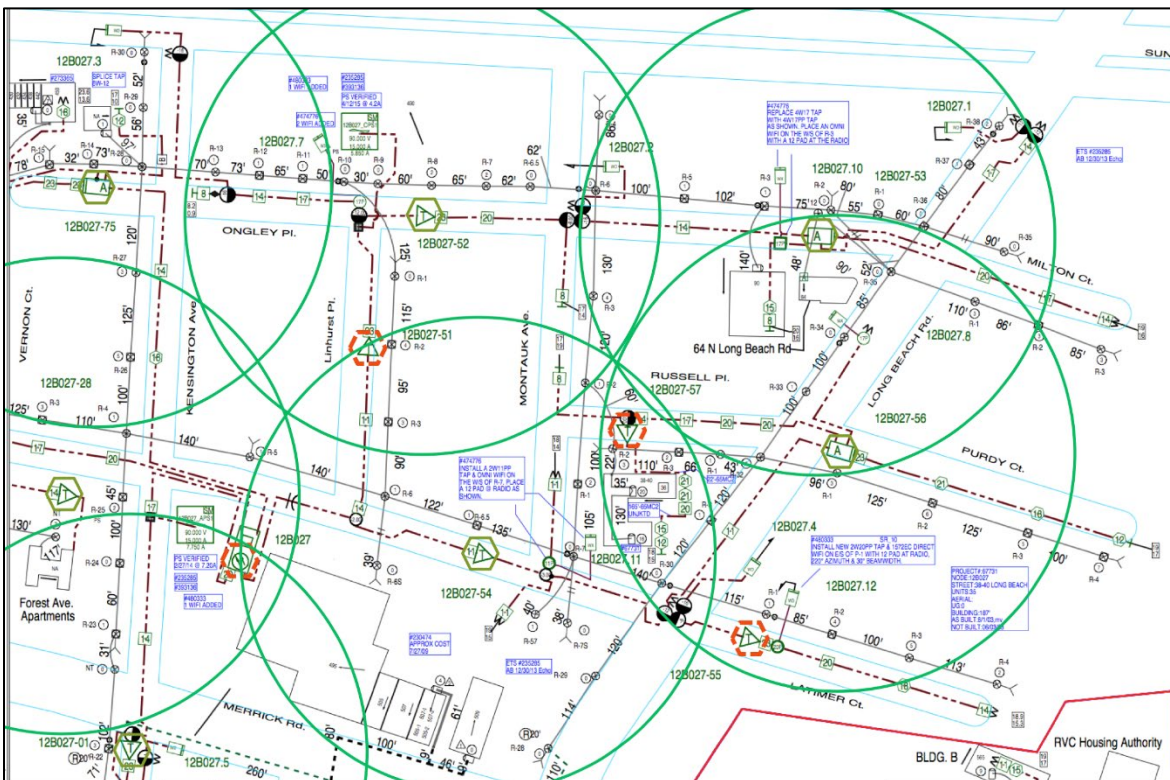


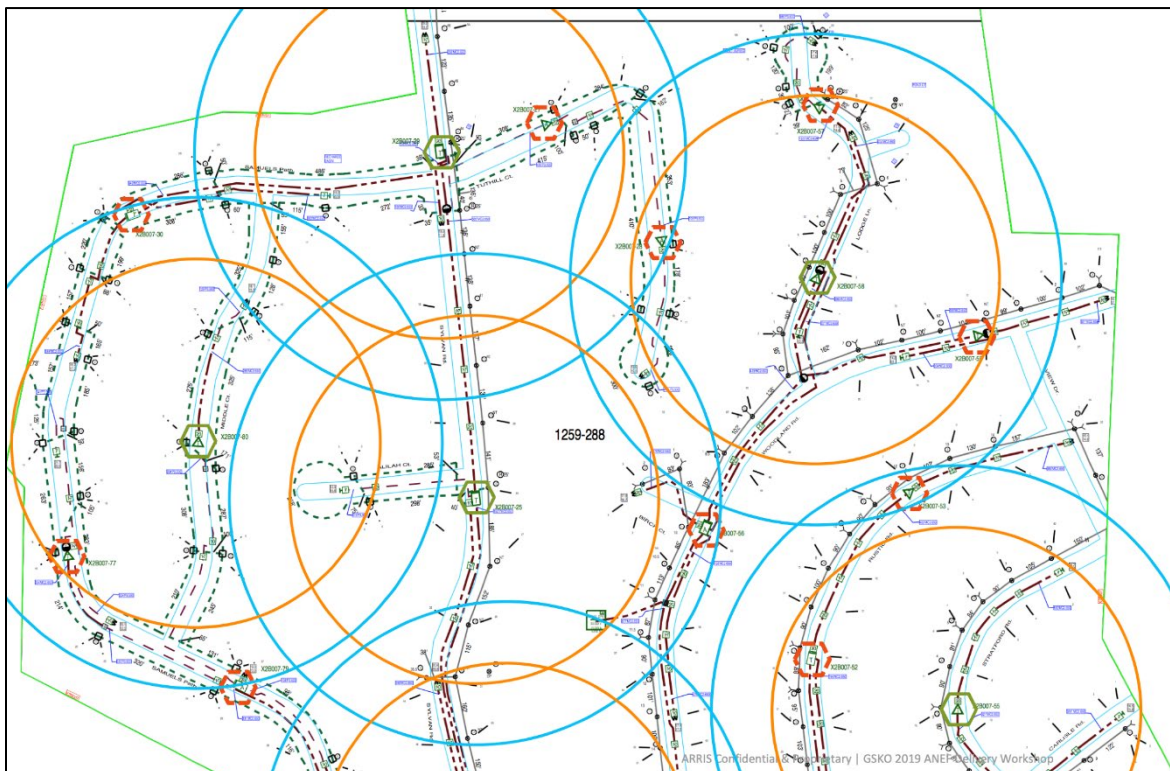
Figure 23 – Higher density (~190 HP/mile) HFC overlaid with 100m cells

Finally, figure 23 looks at the same high-density neighborhood, but now with 100m cell radius overlaid for close to 100% coverage. Note that this smaller cell radius might be needed if there are dense foliage and/or buildings that impede the wireless millimeter wave signals. Note that figure 23 is showing approximately six wireless cells but there are still 4 RF amplifier locations that are without a cell. This is much better than our analysis in table 4 that predicted most RF amplifiers would have a wireless cell. Overall for the high-density example, co-locating the wireless cells with the HFC RF amplifiers is not a problem at any of the three-cell radii considered.

Figures 24-26 now look at a low-density example that is ~60 HP/mile. Figure 24 starts with the 200m cell radius to try and achieve 100% coverage with the 150m radius inside for higher capacities. At low density, there are some homes that fall outside the coverage area. For this particular neighborhood, there are about 5 wireless cells that are co-located with RF amplifiers (i.e. green hexagonal). But there are still 10 RF amplifier locations (i.e. red hexagonal) that are without a cell. This is much better than the 1:1 ratio predicted in table 4.

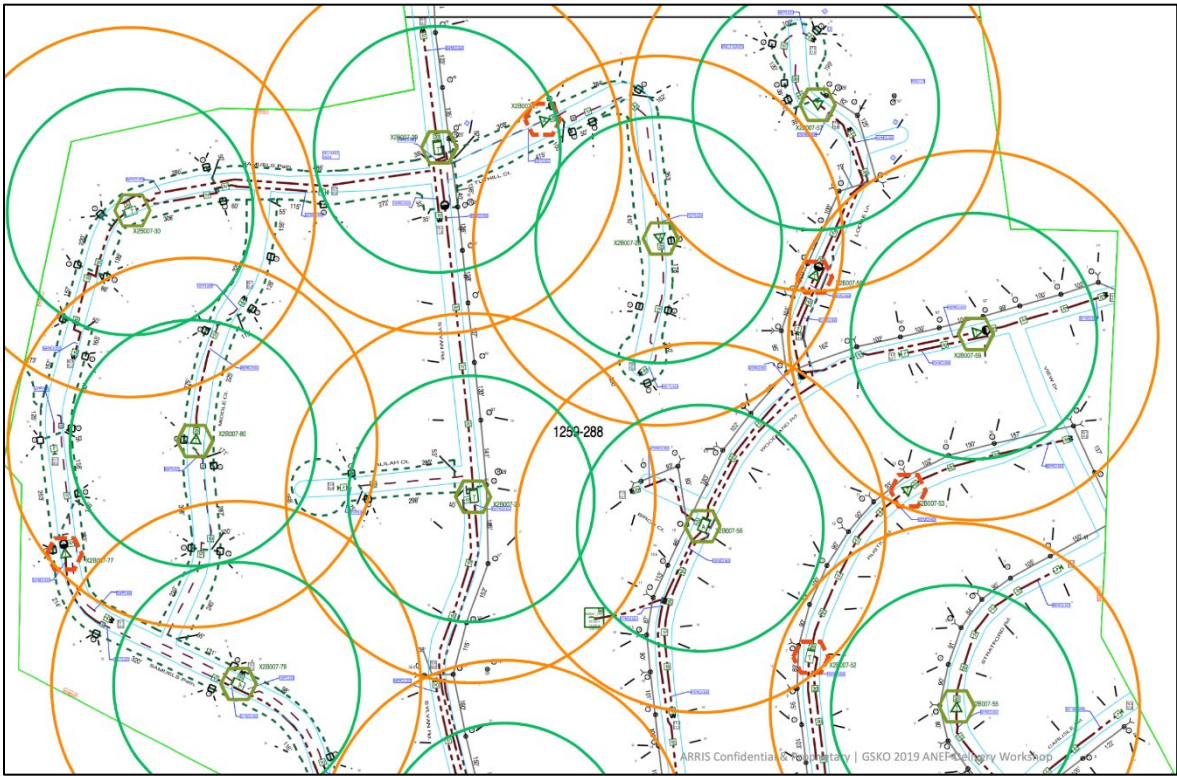
Figure 25 looks at the 150m cell radius to approach 100% coverage with the 100m radius inside for higher capacities. There are about twice as many wireless cells in this example that are co-located with RF amplifiers, which also results with close to 100% home coverage too. But there are still 5 RF amplifier locations without a cell. Our theoretical analysis in table 4 predicted there should be twice as many wireless cells as RF amplifiers.

Finally, in figure 26, 100m cells are placed at every RF amplifier. While overall coverage is still decent, there are some noticeable gaps without coverage. The analysis in table 4 predicted a need of four cells for every RF amplifier, so this example shows that the low density HFC plant might be in much better shape than the paper analysis.

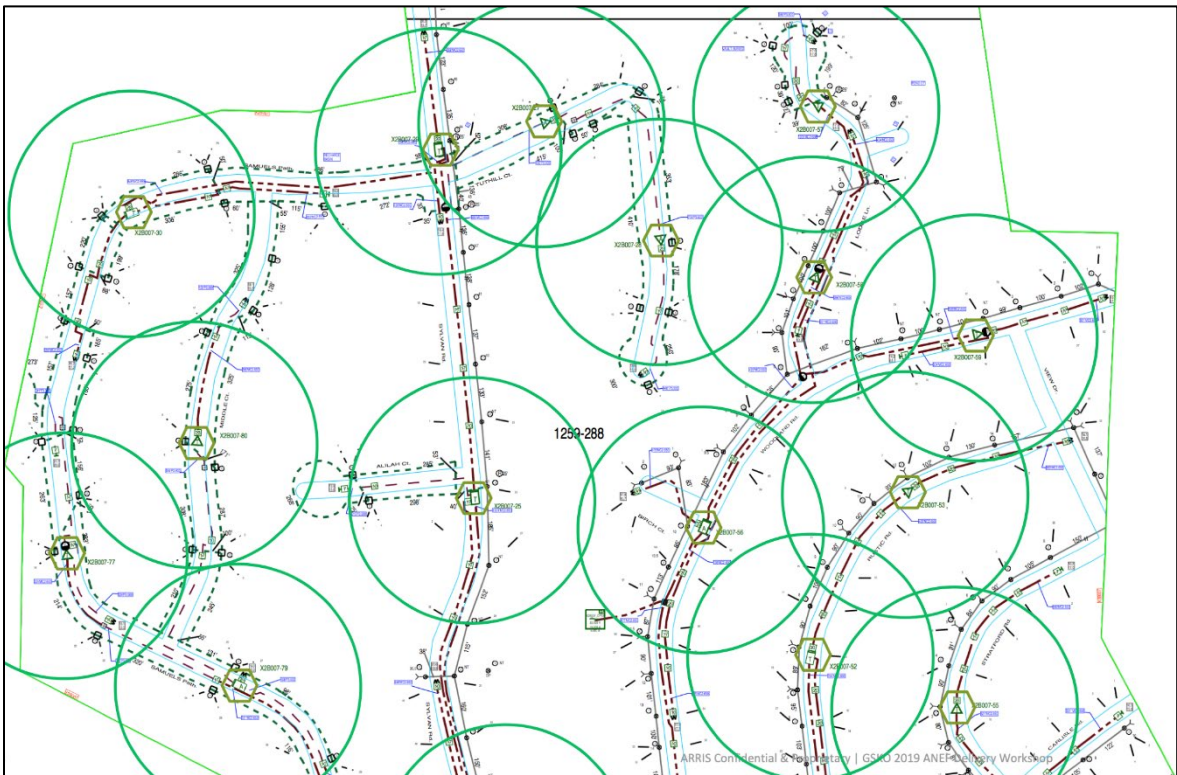


**Figure 24 - Lower density (~60 HP/mile) HFC overlaid with ~200m/150m cells**





**Figure 25 – Lower density (~60 HP/mile) HFC overlayed with ~150m/100m cells**



**Figure 26 – Lower density (~60 HP/mile) HFC overlayed with 100m cells**

Both the high density and low density use cases show that existing HFC amplifier locations can be quite effective for providing 5G millimeter wave coverage, even in the worst case of low home density with a small cell radius (e.g. 100m or less). In these scenarios, it is still possible to add additional cells that are in other locations. It is feasible to add other strand mounted cells that are between RF amplifier locations that get their power from the coax and use DOCSIS over the coax for its backhaul. Another option is the Inside-Outside strategy discussed earlier in this paper. Looking at figure 26, there might only need to be a couple homes selected to install a wireless cell (i.e. Home as a Tower, HaaT).

The previous examples have also shown uniform wireless cell radius. There may be additional optimizations that an operator can make by having wireless cells with different transmit powers. This will allow the operator to maximize coverage for a minimum amount of power consumption.

As time progresses and bandwidth needs continue to rise, an operator might want to migrate from a DOCSIS based backhaul to a fiber backhaul. One solution that supports this is a full fiber deep FTTLA upgrade with its accompanying power reductions. But as our previous case studies have shown, there might not be a cell at every amplifier location. This is especially true in moderate to high density areas. In this scenario, an operator could choose to do a partial FTTLA upgrade where only the RF amplifiers with a small cell adjacent are converted to fiber nodes. This approach is termed FTT5G, or Fiber to the 5G cell.

### **3.3. HFC Adaptability in a changing World**

One observation is that the HFC plant already is where the customers are; more densely populated square miles have more linear cable miles. Per [NNT\_2012], some “things”, both in the nature and in human-made world, benefit from high variability and disorder because of the “thing’s” ability to adapt and gain a competitive advantage in such highly varied and disorderly environment. Could the HFC plant be one of these “things” and if so how?

HFC adaptability is second to none, able to feed consumers from super rural to super urban, whether it is a few homes per square mile up to thousands of homes per square mile. HFC node size and the resulting service group size are highly flexible as well. One node could be 1 or 2 or 4 service groups via node segmentation; then many nodes could be aggregated into a single service group, as in the FTTLA example of Figure 18. A fiber-deep, N+0 last active node could feed just tens of customers directly.

The HFC is extensible. If an adjacent area needs coverage, the operator can build the HFC plant out in a Lego® block like extension. They can add another node, or coax-fed bridger amplifier to connect many dozens of homes or use a line-extender amplifier to do the job if only a few dozen homes. The operator has an unlimited number of combinations of how to form a service group and how to cover an area that is typical or atypical.

The fiber portion of HFC is extensible too. A fiber to the building and FTTP further increases those possibilities. An “inverse node” could take a signal from an RF port, convert it to fiber, then feed a remote housing development that simply was not in any of the plants when the original network was built. The long story short, the more variability encountered in the field, the more the HFC way of serving the customer needs shines.

5G deployment requirements may be yet another “make it shine” aspect of HFC. When comparing various HFC network densities and the 5G access point densities required, it turns out that as many 5G access points are needed per area as there are RF amplifiers in an existing network. In order to deploy efficiently, those access points need backhaul and power supplies nearby. In the near term, DOCSIS 3.1 backhaul will work just fine. Longer term, the FTTLA fiber-deep upgrade of Figure 18, gets fiber within 300 feet of the last tap, and the existing powered hardline coaxial plant now has more kilowatts to spare.

Both backhaul and powering are the two key factors required for the 5G. Densification may be the right term when going from 4G to 5G; right-sizing may be the better term to use when matching the existing HFC plant to the 5G deployment needs.

## Conclusion

### 4. Cable 10G vs. Wireless 5G – Friend or Foe?

So, is Wireless 5G a foe or friend to Cable 10G? The bandwidth capacity of 5G high band millimeter waves does enable multi-gigabit per second services to the home. At first glance, using 5G for Fixed Wireless Access (FWA) might appear as potential stiff competition to cable. But in reality, Cable 10G offers two to three times the bandwidth capacity of 5G with a roadmap to even higher capacities in the future [see CLO\_2019].

A gigabit per second wireless service is still formidable. However, it may be very tough to make the economics work outside of densely populated areas or select locations like MDUs. With a typical density of 100 HP/mile, the FWA operator might need to deploy a cell for every 10-20 homes. At low 50 HP/mile densities, the cell may only cover 2-5 homes.

In addition to these challenges, the FWA operator will need to come up with both a backhaul and a powering infrastructure. FTTP operators may have the backhaul portion covered, but how do they power their small cells? Developing the backhaul and powering infrastructure will be a daunting task for any potential FWA operator. Anyone that is except a cable operator.

Perhaps a better question to answer is: “Who is strategically in the best position to mesh the capacity advantage of wired and the untethered access advantage of wireless?” and the answer may be cable operators! Because, behind a successful wireless network, there is a wired network, with power and bandwidth capacity to boot.

As shown in our case studies in this paper, the HFC plant aligns quite nicely with the needs for 5G backhaul and powering across a wide variety of housing densities. Some HFC plants may have adequate power today to support a small cell infrastructure. Other HFC plants can address the power issue with fiber deep FTTLA upgrades.

DOCSIS 3.1 provides adequate capacity for 5G backhaul in the near term. A longer-term strategy of fiber deep can eventually convert this to fiber backhaul. One such strategy discussed is FTT5G where fiber is pulled to those actives with a wireless cell adjacent to it.

The wireless and wired technologies need each other to make a better system. The future high bandwidth, high frequency wireless systems need small cells with many access points requiring a low latency wired backhaul; and APs positioned inside the home/MDU and outside in every neighborhood for optimum coverage. Cable is ideally suited to support this backhaul infrastructure. Meanwhile, Cable 10G can provide multi-gigabit capacity to the home’s entry point but needs a robust high capacity wireless connection for that final 100 meters inside and around the home to every mobile device.

In the end, Cable 10G and Wireless 5G/CBRS are much stronger together and are at the core of a next generation network evolution. We think that they will be best friends.

# Bibliography & References

[ALB\_2019] A. Al-Banna et. al., “Operational Considerations and Configurations for FDX & Soft-FDD,” SCTE Cable-Tec 2019, SCTE

[CHE\_2018], C. Cheevers, “How will Outdoor and Indoor use of 5G Wireless Services really work?”, IBC 2018

[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

[FDX\_PHY] “DOCSIS 4.0 Physical Layer Specification”, CM-SP-PHYv4.0-D01-190628, Cablelabs 2019

[FLE\_2017], J.R. Flesch, et. al., “Can a Fixed Wireless Last 100m Connection Really Compete with a Wired Connection and Will 5G Really Enable this Opportunity?”, SCTE Cable-Tec 2017, SCTE

[FLE\_2018], J.R. Flesch, C. Cheevers, “The New Home as a Hotspot: Wi-Fi Meet CBRS LTE and Meet Your Long Range Brother LoRa”, SCTE Cable-Tec 2018, SCTE

[NAHB\_2018] <http://eyeonhousing.org/2018/08/lot-size-remains-record-low/>

[NNT\_2012] Nassim Nicholas Taleb, “Antifragile: Things That Gain from Disorder”, Random House, November 27, 2012

[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\_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

5G	5 <sup>th</sup> generation (wireless)
10G	10 gigabit platform (cable)
AP	access point
bps	bits per second
AR	Augmented reality
BW	bandwidth
CAPEX	Capital Expense
CBRS	Citizens Broadband Radio Service
CPE	Consumer Premise Equipment
DAA	Distributed Access Architecture
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
FDX	Full Duplex (i.e. DOCSIS)
FTTLA	Fiber to the Last Active
FTTP	Fiber to the Premise
FTT5G	Fiber to the 5G small cell
FWA	Fixed Wireless Access
Gbps	Gigabits Per Second
GHz	Gigahertz
HaaS	Home as a Tower
HD	high definition
HFC	hybrid fiber-coax
HP	Homes Passed
Hz	Hertz
IoT	Internet of things
ISBE	International Society of Broadband Experts
ITU	International Telecommunication Union
IEEE	Institute of Electrical and Electronics Engineers
LLD	Low latency DOCSIS
LoRa	Long Range
LOS	Line of sight
LTE	Long term evolution
MAC	Media Access Control interface
MDU	Multiple Dwelling Unit
MHz	Megahertz
MIMO	multiple-input and multiple-output
MoCA	Multimedia over Coax Alliance
MSO	Multiple System Operator
MVNO	Mobile Virtual Network Operator
N+0	Node+0 actives
NCTA	National Cable and Telecommunications Association
nLOS	Near line of sight
NSI	Network Side Interface
Nsub	Number of subscribers

OFDMA	Orthogonal Frequency Division Multiplexing Access (Upstream)
OFDM	Orthogonal Frequency Division Multiplexing
OPEX	Operating Expense
PGS	Predictive Grant Service (DOCSIS)
PHY	Physical interface
PNM	Proactive Network Maintenance
PON	Passive Optical Network
QoE	Quality of Experience
RF	Radio frequency
R-PHY	Remote PHY
Rx	Receive
SCTE	Society of Cable Telecommunications Engineers
SFU	Single family unit
SG	Service Group
SLA	Service level agreement
Tavg	Average throughput per subscriber
Tx	Transmit
UHD	Ultra-high definition (4K, 8K)
US	Upstream
VR	Virtual reality

CommScope and ARRIS are trademarks of CommScope, Inc. and/or its affiliates. DOCSIS is a trademark of Cable Television Laboratories, Inc. All other trademarks are the property of their respective owners.

# **Upgrading the Plant to Satisfy Traffic Demands**

## **The One Touch Approach**

A Technical Paper prepared for SCTE•ISBE by

**Nader Foroughi**  
Sr. Network Architect  
Shaw Communications  
2728 Hopewell Place NE, T1Y 7J7  
+1 403 648 5937  
nader.foroughi@sjrb.ca

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Content .....	5
1. Scope .....	5
2. Analysis .....	6
2.1. 1.8GHz Extended Spectrum DOCSIS.....	6
2.2. Test Plan .....	9
2.3. Test Results.....	10
2.3.1. TCP for DS and US:.....	10
2.3.2. DS MER Estimations.....	14
2.3.1. Cascaded Deployment.....	23
2.3.2. US MER Estimations.....	30
2.3.3. Capacity Analysis.....	32
2.4. Cost Analysis .....	34
2.5. CM's .....	36
Conclusion .....	37
Abbreviations.....	38

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Nielsen’s Law of Internet Bandwidth – Tom Cloonan (Arris).....	5
Figure 2 – Coaxial Cable Loss (50MHz – 3GHz).....	6
Figure 3 – 412 P3 Test Plant .....	7
Figure 4 – 625 P3 Test Plant .....	8
Figure 5 – 500 P3 Test Plant .....	8
Figure 6 – S Parameters.....	9
Figure 7 – Test Methodology .....	10
Figure 8 – Current Amplifier Output Power Levels/6MHz (50MHz – 1GHz).....	10
Figure 9 – Projected 1.8GHz Amplifier Output Power Levels/6MHz (108MHz – 1.8GHz).....	11
Figure 10 – Projected Drop-Down 1.8GHz Amplifier Output Power Levels/6MHz .....	12
Figure 11 - Projected Flat 1.8GHz Amplifier Output Power Levels/6MHz.....	12
Figure 12 - Projected MDM Transmit Levels/6.4MHz (108MHz – 684MHz).....	13
Figure 13 – S21 – 412 P3 Cable Test Plant .....	14
Figure 14 – 412 Test Plant Pedestal .....	15
Figure 15 – Projected MDM Rx Levels – Continuing with the Tilt.....	15
Figure 16 – Projected MDM Rx Level – Drop Down at 1GHz.....	16
Figure 17 – Projected MDM Rx Levels – Flat after 1GHz .....	17
Figure 18 – S21 – 625P3 Test Plant .....	18
Figure 19 – Projected MDM Rx Levels – Continuing with the Tilt.....	19
Figure 20 – Projected MDM Rx Levels – Drop Down at 1GHz.....	19

Figure 21 – Projected MDM Rx Levels – Flat after 1GHz .....	20
Figure 22 – S21 – 500P3 Test Plant .....	21
Figure 23 – Projected MDM Rx Levels – Continuing with the Tilt.....	22
Figure 24 – Projected MDM Rx Levels – Drop Down at 1GHz.....	22
Figure 25 – Projected MDM Rx Levels – Flat after 1GHz .....	23
Figure 26 – Projected Amplifier Input Port Rx Levels – Continuing with the Tilt.....	24
Figure 27 – Projected Amplifier Input Port Rx Levels – Drop Down at 1GHz.....	25
Figure 28 – Projected Amplifier Input Port Rx Levels – Flat after 1GHz .....	25
Figure 29 – Projected Amplifier Input Port Rx Levels – Continuing with the Tilt.....	26
Figure 30 – Projected Amplifier Input Port Rx Levels – Drop Down at 1GHz.....	27
Figure 31 – Projected Amplifier Input Port Rx Levels – Flat after 1GHz .....	27
Figure 32 – Projected Amplifier Input Port Rx Levels – Continuing with the Tilt.....	28
Figure 33 – Projected Amplifier Input Port Rx Levels – Drop Down at 1GHz.....	29
Figure 34 – Projected Amplifier Input Port Rx Levels – Flat after 1GHz .....	29
Figure 35 – Projected MDM Rx Levels at the Amplifier Port in the US .....	30
Figure 36 – Projected MDM Rx Levels at the Amplifier Port in the US .....	31
Figure 37 – Projected MDM Rx Levels at the Amplifier Port in the US .....	31
Figure 38 – 1.8GHz ESD Spectrum Plan .....	32
Figure 39 – FDX Spectrum Plan .....	32
Figure 40 – Capacity Comparisons – Bar Graph .....	33
Figure 41 – 1.8GHz ESD Upgrade Cost – Materials vs Labour.....	34
Figure 42 – N+2 FDX Upgrade Cost – Materials vs Labour .....	35
Figure 43 – N+0 FDX Upgrade Cost – Materials vs Labour .....	35
Figure 44 – Upgrade Costs – Materials vs Labour.....	36

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Modulation Orders and Effective Throughputs .....	14
Table 2 – Capacity Comparisons Table.....	33

# Introduction

The majority of MSOs outside plant architecture (OSP) consists of N+X. As time goes on the demand for capacity and speed in both upstream and downstream grows. Along with that, the era of symmetrical services is approaching as competitive pressure arises.

Fibre to the premises (FTTP) can help to meet this capacity demand but it is extremely costly. Over the past few years, many innovative alternatives and technologies have been proposed to alleviate this challenge. Full-duplex-DOCSIS (FDX) was one of the developed technologies in response to these demands. Although this concept is revolutionary, it requires the MSO to upgrade the OSP to a passive (N+0) state. This is more cost effective than FTTP, but depending on the operator, area of construction and plant type (aerial or underground), it can be quite costly. Moreover, FDX can be challenging from a technology implementation perspective, as it requires overlapping the downstream and upstream spectrum from 108MHz-684MHz.

Knowing that coaxial cable has 6GHz of useable bandwidth (BW) on average, last year the idea of extending the spectrum from 1.2GHz to 1.8GHz and eventually 3GHz was proposed, which gained a lot of traction in the industry. This can present different approaches to upgrading the OSP, to satisfy the future capacity demands.

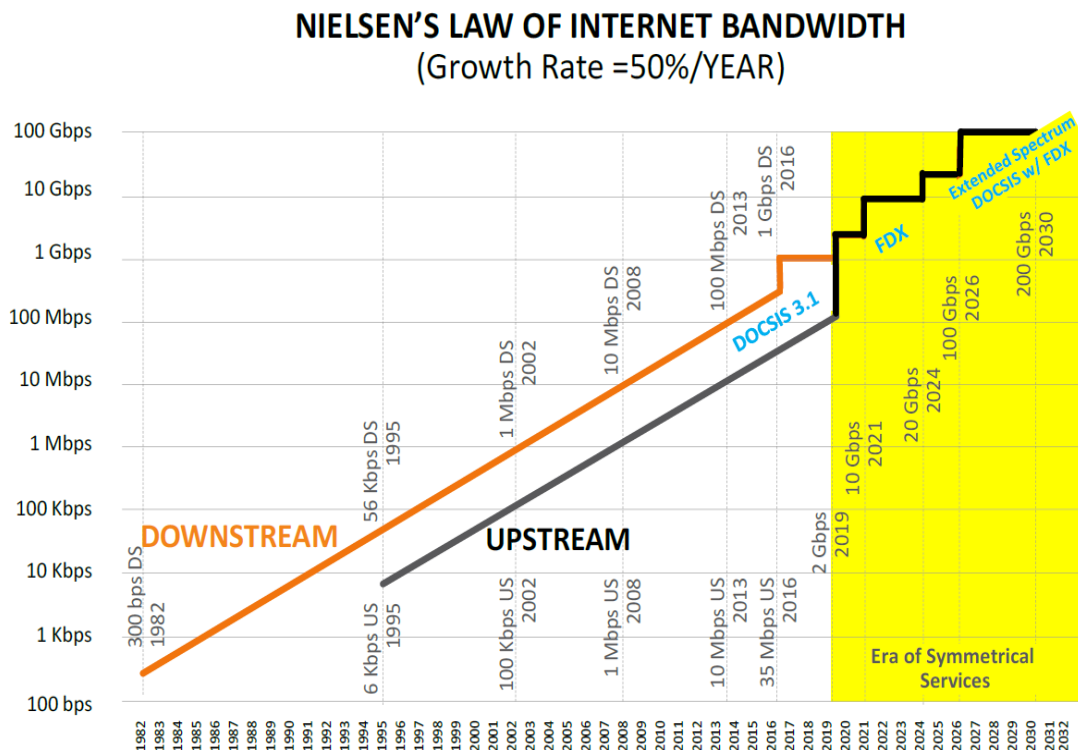
In this paper an analysis has been carried out to evaluate the achievable capacity in an extended spectrum network. Capacity estimates have been based on the field measurements taken from the acquired 2.7GHz taps. Furthermore, a cost analysis has been carried out based on capacity estimations in a cascade of N+4.

Based on the analysis demonstrated in this paper, 1.8GHz extended-spectrum-DOCSIS (ESD) can be a great alternative to both N+2 and N+0 FDX. It can provide matching throughputs as N+0 and N+2 FDX, at a lower cost. This allows an MSO to rapidly deploy this technology throughout their existing cascaded plant, in a cost-effective manner.

# Content

## 1. Scope

Historically both upstream (US) and downstream (DS) capacity demands have been growing substantially year-over-year, as we approach the era demonstrated in Figure 1. The year-over-year growth shown in Figure 1, also referred to as Compound Annual Growth Rate (CAGR) is 50%, which is a common and historical rule-of-thumb for DS CAGR. Upstream CAGR has been on average lower than in the DS but is also more volatile



**Figure 1 – Nielsen’s Law of Internet Bandwidth – Tom Cloonan (Arris)**

This can present many challenges for a multi-system-operator (MSO) as the majority of the outside plant (OSP) architecture consists of N+X. On average, Shaw’s plant consists of N+4. Since business-as-usual (BAU) node splits don’t increase the overall available BW, other strategies must be considered.

In order to quantify the differences between various deployment strategies, the capacities offered by each technology must be evaluated. The categories below have been considered for evaluation:

1. DS and US analysis in a 1.8GHz N+4 ESD plant
2. DS and US analysis in an N+0 FDX plant
3. DS and US analysis in an N+2 FDX plant

In order to estimate the achievable capacity in Scenario 1, a series of acquired 2.7GHz taps were installed in the last span of selected amplifiers. The characterization performance for this test is demonstrated in the analysis section.

The calculated capacities for each scenario can then be compared against the projected demands and the cost of the plant upgrade as it is introduced to satisfy the traffic demands. With all cost variables accounted for, this can then be used to create an NPV analysis for each strategy.


In the subsequent sections of this paper, the RF and capacity analysis for each scenario is demonstrated.

## 2. Analysis

### 2.1. 1.8GHz Extended Spectrum DOCSIS

Historically MSOs have stretched the spectrum to higher frequencies, from 500MHz to 750MHz, 860MHz, 1GHz and currently 1.2GHz. This can be a challenging task depending on how the OSP was designed. Amplifier spacing and tap span lengths can be a concern when this is put into practice.

Moreover, as the spectrum is stretched higher, coax cable becomes subject to more attenuation, which can be challenging when considering upgrading the plant to 1.8GHz. This is demonstrated Figure 2 below:




**Figure 2 – Coaxial Cable Loss (50MHz – 3GHz)**

The following locations were selected to evaluate the feasibility of upgrading the plant to 1.8GHz and eventually 3GHz:




- 412 P3 cable with many splices as a result of plant maintenance (see Figure 3)
  - This cable presents us with the most amount of loss in comparison to other cable types at Shaw, as per Figure 2.



**Figure 3 – 412 P3 Test Plant**

- 625 P3 cable with the last span tap being ~160m away (see Figure 4)
  - This cable is representative of an average cable type at Shaw with average loss



**Figure 4 – 625 P3 Test Plant**

- 500 QR cable (see Figure 5)
  - Just like the 412 P1 cable, this cable has high loss characteristics

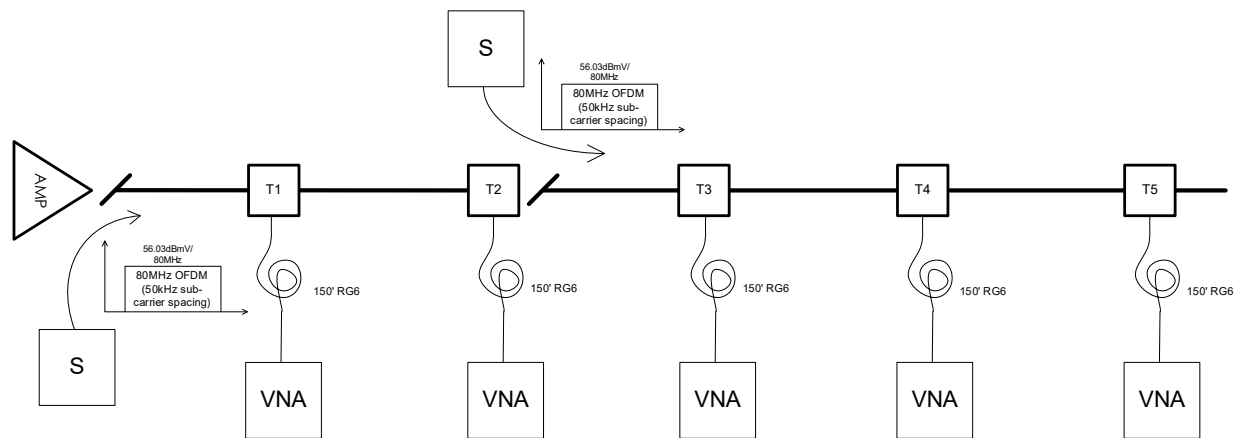


**Figure 5 – 500 P3 Test Plant**

In each of the selected plant locations, the old taps were replaced with a similar value of 2.7GHz BW in order to have a like-for-like comparison.

The test methodology is described below:





**Figure 7 – Test Methodology**


## 2.3. Test Results

### 2.3.1. TCP for DS and US:

Prior to evaluating the S21 and end-of-line capacity analysis, the output power of the amplifier and the total-composite-power (TCP) shall be discussed. The following assumptions have been considered:

- The TCP of the current amplifier gain chips are 73.8dBmV
- The high output gain chip used in N+0 implementations has TCP = 76.8dBmV
- Amplifier performance (gain, noise and distortion) characteristics of the 1.8GHz actives will be similar and/or comparable to the current 1.2GHz ones


Given that ESD will be deployed in cascaded plant, the focus of this paper will be on cascaded levels and tilt. The current amplifier output levels used are shown in Figure 8 below.



**Figure 8 – Current Amplifier Output Power Levels/6MHz (50MHz – 1GHz)**


The levels shown above are in analog. To convert them to digital, 6dB has to be deducted from the power level at any given frequency. Note that TCP will be calculated using digital levels. Since these levels are quite conservative and the TCP of the current amplifier gain-chips are 73.8dBmV, the following three power loading and tilt scenarios have been assumed for the deployment of 1.8GHz capable amplifiers:

1. Continuing with the tilt:
  - This results in a 71dBmV TCP (see Figure 9)




**Figure 9 – Projected 1.8GHz Amplifier Output Power Levels/6MHz (108MHz – 1.8GHz)**

2. Drop-down at 1GHz (see Figure 10)
  - Tilting the spectrum in the ‘legacy’ band up to 1GHz, dropping the level by 3dBmV and continuing with the same tilt up to 1.8GHz
  - This results in a 68.5dBmV of TCP



**Figure 10 – Projected Drop-Down 1.8GHz Amplifier Output Power Levels/6MHz**

3. Flat after 1GHz (see Figure 11)
  - Tilting the spectrum in the ‘legacy’ band up to 1GHz and continuing with same level from 1GHz to 1.8GHz
  - This results in a 58.5dBmV of TCP




**Figure 11 - Projected Flat 1.8GHz Amplifier Output Power Levels/6MHz**

For upstream modem transmit levels (MDM Tx), the following has been assumed:

- The modem will have the same transmit capabilities as described in the full-duplex-DOCSIS specifications

- The spectrum is not overlapped

According to the FDX specifications the modem can transmit with a 10dB tilt from 108MHz to 684MHz with a TCP of 64.5dBmV. The following modem transmit levels have been utilized to estimate capacity in the US:



**Figure 12 - Projected MDM Transmit Levels/6.4MHz (108MHz – 684MHz)**

Although modems in the field will not use the entire 108-684MHz spectrum for upstream burst, upstream capacity has been calculated throughout the entire spectrum.

The assumption for the end-of-line capacity analysis is:

- The modem shall be a point-of-entry (PoE) device, to be installed at the ground block without any splitters
- The limiting factor in achievable modulation order and modulation-error-rate (MER) is the noise floor of the modem and amplifier, if the plant is properly aligned and interference free
- A cascade of four has been assumed for the estimates (N+4)
- A 3dB reduction in signal-to-noise-ratio (SNR) has been included in the estimate, for every doubling of amplifiers. Meaning for a cascade of 4, per assumption above, a 6dB SNR reduction is included as delivered by the network to the modem. This reduction to SNR has been applied to the minimum Rx levels needed at the modem to achieve each modulation order, based on DOCSIS specifications of -15dBmV - +15dBmV per 6MHz at the MDM
- The throughput of each modulation order is based on the theoretical values, not including any overhead, as shown in Table 1 below:


**Table 1 – Modulation Orders and Effective Throughputs**

Modulation Rate	Effective Throughput (Bits/s/Hz)
256QAM	8
1024QAM	10
4096QAM	12

### 2.3.2. DS MER Estimations

#### 2.3.2.1. 412 P1 Location

The insertion loss (S21) results for this location are demonstrated in Figure 13 below:



**Figure 13 – S21 – 412 P3 Cable Test Plant**

From Figure 13, it can be observed that there are “suck outs” present in the spectrum. This can be attributed to a variety of factors, but it is most likely due to the number of coax splices installed in this particular plant in the past ~30 years. A picture of the pedestal where the splices are visible is shown in Figure 14






**Figure 14 – 412 Test Plant Pedestal**

This can be alleviated with regular plant maintenance, using a time-domain-reflectometer (TDR). For the purpose of this paper, the calculations have been done on plant, as-is.


Based on the S21 measurement and the 1.8GHz power loading profiles in Figures 9-11, the following modem receive power levels (MDM Rx) can be calculated for each profile.

1. Continuing with the tilt:




**Figure 15 – Projected MDM Rx Levels – Continuing with the Tilt**

2. Drop-down at 1GHz:



**Figure 16 – Projected MDM Rx Level – Drop Down at 1GHz**

3. Flat after 1GHz:




**Figure 17 – Projected MDM Rx Levels – Flat after 1GHz**

From Figure 15-17 it can be observed that 4kQAM is achievable in the majority of the spectrum across all the taps. Moreover, the only case that 256QAM is not achievable from the figures above is in the flat scenario above ~1.6GHz.

This can be overcome with adjusting the tap values along with using a different drop cable, for example RG11, which should increase the levels by ~2.5dB.

### 2.3.2.2. 625 P3 Location


The insertion loss (S21) results for this location are demonstrated in Figure 18 below:



**Figure 18 – S21 – 625P3 Test Plant**


Based on the S21 measurement and the 1.8GHz power loading profiles in Figures 9-11, the following modem receive power levels (MDM Rx) can be calculated, for each power loading profile:

1. Continuing with the tilt:




**Figure 19 – Projected MDM Rx Levels – Continuing with the Tilt**

2. Drop down at 1GHz:



**Figure 20 – Projected MDM Rx Levels – Drop Down at 1GHz**

### 3. Flat after 1GHz:




**Figure 21 – Projected MDM Rx Levels – Flat after 1GHz**

It can be observed that there are certain frequencies where the achievable modulation order is below 256QAM for taps 4 and 5 in the drop-down and flat scenario. This is primarily due to the fact that this plant was not optimized and designed for 1.8GHz. This is visible since tap 5 can achieve a higher order of modulation in comparison to taps 3 and 4. The tap value can play a crucial role in the MDM Rx power which consequently translates to the achievable modulation order. For example, in the case taps 3 and 4, a lower tap value can increase the achievable modulation order at the modem.

#### **2.3.2.3. 500 P3 Location**

The insertion loss (S21) results for this location are demonstrated in Figure 22 below:



**Figure 22 – S21 – 500P3 Test Plant**

Based on the S21 measurement and the 1.8GHz power loading profiles in Figure 9-11, the following modem receive power levels (MDM Rx) can be calculated for each profile.:

1. Continuing with the tilt:




Figure 23 – Projected MDM Rx Levels – Continuing with the Tilt

2. Drop down at 1GHz:





Figure 24 – Projected MDM Rx Levels – Drop Down at 1GHz



### 3. Flat after 1GHz:



**Figure 25 – Projected MDM Rx Levels – Flat after 1GHz**

It's visible that although 500 P3 cable suffers more attenuation in comparison to newer cable types, 4kQAM is achievable in almost all cases, regardless of the tilt scenario.

#### **2.3.1. Cascaded Deployment**

Given that this technology will be deployed in existing cascaded plant, input levels to the next amplifier in cascade can be a concern. In order to substantiate this the following was performed:

- Mathematically remove the insertion loss of the final self-terminating tap along with 150' of RG6 drop cable
- Mathematically insert the insertion loss of a non-self-terminating tap at the end of line
- Mathematically insert the insertion loss if 100' of 412 cable
  - Given that 412 cable has the highest attenuation in the tested scenarios, as per figure 2, this was kept consistent amongst all test scenarios

Two commonly available classes of amplifiers in the industry today are the Mini-Bridger (MB) and the Line Extender (LE). Typical characteristics of each amplifier type are outlined below:

- Mini-bridger (MB):
  - 42dB of gain
  - 9dB of noise figure
- Line-extender (LE):
  - 34dB of gain
  - 10dB of noise figure


It should be noted that that the same TCP and amplifier output levels as section 2.3.1 has been applied. Below the results have been outlined:

#### **2.3.1.1. 412P3 Location**

The results for each power loading profile in Figures 9-11 are demonstrated in Figures 26-28 below:


Note that all levels are shown in analog/6MHz.

1. Continuing with the tilt:




**Figure 26 – Projected Amplifier Input Port Rx Levels – Continuing with the Tilt**

2. Drop down at 1GHz:



**Figure 27 – Projected Amplifier Input Port Rx Levels – Drop Down at 1GHz**

3. Flat after 1GHz:




**Figure 28 – Projected Amplifier Input Port Rx Levels – Flat after 1GHz**

The results above indicate that the ‘flat tilt’ scenario will not have sufficient level at the amplifier port. This means that the operator will either have to sacrifice the performance of that spectrum (1600MHz – 1800MHz) or they will have to deploy a different tilt scenario (continue with the tilt or drop-down), assuming that the TCP of the currently deployed amplifiers have not be exhausted.

### 2.3.1.2. 625P3 Location

The results for each power loading profile in Figures 9-11 are demonstrated in Figures 29-31 below:

1. Continuing with the tilt:



**Figure 29 – Projected Amplifier Input Port Rx Levels – Continuing with the Tilt**

2. Drop down at 1GHz:




Figure 30 – Projected Amplifier Input Port Rx Levels – Drop Down at 1GHz

3. Flat after 1GHz:




Figure 31 – Projected Amplifier Input Port Rx Levels – Flat after 1GHz

It's visible that in all the output power scenarios, assuming that the current amplifiers' gain and noise characteristics remain the same, the signal can be amplified with enough level to reach the next amplifier in cascade.

### 2.3.1.3. 500P3 Location

The results for each power loading profile in Figures 9-11 are demonstrated in Figures 32-34 below

1. Continuing with the tilt:




Figure 32 – Projected Amplifier Input Port Rx Levels – Continuing with the Tilt

2. Drop down at 1GHz:




Figure 33 – Projected Amplifier Input Port Rx Levels – Drop Down at 1GHz

3. Flat after 1GHz:




Figure 34 – Projected Amplifier Input Port Rx Levels – Flat after 1GHz

It's visible that in all the output power scenarios, assuming that the current amplifiers' gain and noise characteristics remain the same, the signal can be amplified with enough level to reach the next amplifier in cascade.


### 2.3.2. US MER Estimations

There are various challenges with upstream MER analysis. For downstream the main limitations of MER are receive power levels and the number of amplifiers in cascade. It's generally accepted to assume 3dB of SNR loss per doubling of amplifiers. That's not the case for upstream. The main limitation of MER in the upstream is noise funneling, which is hard to model, as this depends on the number of modems that are bursting simultaneously in the US, as well as unpredictable external interference sources that funnel upstream. Moreover, it's hard to simulate what channels are going to be occupied in the US, depending on where the modem sits in the cascade.

With that in mind, since it's generally understood that the modems will have the same transmit capabilities in the US as defined in FDX spec, as per Figure 12, the capacity comparison in the US for 1.8GHz ESD and FDX becomes easier.

The receive levels per 6.4MHz at the amplifier port in the return for each cabling scenario in Figures 3-5 are shown in Figure 35 below, assuming that the modem is transmitting throughout the FDX return-band (108MHz – 684MHz), with the tilt shown in Figure 12:


- 412P3 Location



**Figure 35 – Projected MDM Rx Levels at the Amplifier Port in the US**




- 625P3 Location



**Figure 36 – Projected MDM Rx Levels at the Amplifier Port in the US**

- 500P3 Location



**Figure 37 – Projected MDM Rx Levels at the Amplifier Port in the US**

### 2.3.3. Capacity Analysis

In order to estimate the achievable capacity in the network, an US/DS split needs to be selected. The following assumptions have been considered for this analysis:

- The spectrum is assumed to be 100% IP / DOCSIS
- A guard band of 20% has been assumed for ESD
- The spectrum evolution has been assumed to be completed in this analysis, such that in both FDX and ESD cases the US has been stretched to 684MHz. The spectrum plans in Figures 38-39 below demonstrate this:




Figure 38 – 1.8GHz ESD Spectrum Plan





Figure 39 – FDX Spectrum Plan

- Given the data in section 2.3.2 and 2.3.3 an average of 1kQAM for the plant has been assumed for US up to 684MHz and DS up to 1.5GHz. A conservative modulation order of 256QAM has been assumed between 1.4GHz and 1.8GHz.
- For FDX, without available performance of technology under development, the following has conservatively been assumed
  - For N+0, 1kQAM is used for both US and DS
  - For N+2, 256QAM is used in the US and 512QAM is used in the DS

Based on the assumptions above and figures 38 and 39, the following table and bar chart can be produced:

**Table 2 – Capacity Comparisons Table**

Architecture	DS Throughput (Gbps)	US Throughput (Gbps)
<b>1.8GHz Dynamic FDD</b>	8	6.3
<b>N+0 FDX</b>	8.8	6.3
<b>N+2 FDX</b>	6.5	3.8



**Figure 40 – Capacity Comparisons – Bar Graph**


## 2.4. Cost Analysis

Assuming that N+X FDX amplifiers are developed, the following figures have been produced to visualize the comparisons between 1.8GHz ESD, N+2 FDX and N+0 FDX:

Note: The figures below are produced based on Shaw's experience performing fibre-extensions, as a part of node splits and drop-in upgrades, as a part of mid-split upgrades in the plant. Depending on each MSO's location and cost for construction and/or fibre-extension costs, these results may vary. The following have also been included in this analysis:


- ESD cost analysis:
  - No fibre extensions
  - Drop-in upgrades at existing node, amplifiers and tap locations
- N+2 FDX:
  - Fibre-extensions to existing amplifier locations such that no cascade length is larger than 2 after the installed nodes
  - Drop-in upgrades at existing node, amplifiers and tap locations
- N+0 FDX:
  - Fibre-extensions to existing amplifier locations such that no amplification is required after the installed nodes
  - Drop-in upgrades at existing tap locations

### 1.8GHz ESD Cost Breakdown




**Figure 41 – 1.8GHz ESD Upgrade Cost – Materials vs Labour**

### N+2 FDX Cost Breakdown




**Figure 42 – N+2 FDX Upgrade Cost – Materials vs Labour**

### N+0 FDX Cost Breakdown



**Figure 43 – N+0 FDX Upgrade Cost – Materials vs Labour**



**Figure 44 – Upgrade Costs – Materials vs Labour**

## 2.5. CM's

Although this paper has mainly focused on the analysis of various access architectures and costs associated with them, it is worth noting how the CMs fit in the full story. FDX modems will have the most flexibility and they can be deployed in any architecture, including 1.8GHz ESD, given that the echo-canceller (EC) can be utilized as a 'dynamic' di-plex filter. This will have an impact on the cost of the modem.

Di-plexed modems don't have the same dynamic capability as FDX modems but they do potentially provide faster time to market along with lesser cost.

Given that MSOs may have different deployment strategies for their access networks, both modems mentioned above can potentially be deployed. This means that if the same silicon is developed, accommodating FDX with 1.8GHz of spectrum (DOCSIS 4.0), each operator can determine how the silicon can be implemented in their modems, which can potentially reduce price of the silicon and the modems.

# Conclusion

As per the data outlined in this paper, 1.8GHz ESD is a viable option for the access network. As demonstrated in the plant measurements taken and the analytics carried out, it is reasonable to expect high modulation orders such as 4kQAM and/or 1kQAM from the existing plant. The results of the analysis identified distance as the most significant factor in loss and therefore achievable MER.

Based on the capacity analysis demonstrated in this paper, the following can be concluded:

- **OSP**

1.8GHz ESD equipment (amplifiers, taps and passives) provides the lowest projected initial implementation cost for roughly equivalent total capacity as N+0 FDX. Although N+0 FDX can match and potentially surpass this capacity, it doesn't utilize the full RF capacity potential of the current coaxial cable in the plant as much. Upgrading the current HFC actives to extract that capacity is a viable possibility with DOCSIS 4.0 ESD. For some operators, this also provides for a more evolutionary approach, and more similar to MSO upgrades of the past, for increasing the available bandwidth in the network.

Echo-cancellation technology has the potential to reduce the guard band between US and DS, if they were to be implemented in 1.8GHz amplifiers and nodes, maximizing spectral efficiency and flexibility. Industry alignment on this item can potentially drive down cost and decrease the development timelines.

- **CPE and CMTS (RPHY)**

FDX based silicon with echo-cancellation technology provide the most flexibility since the equipment can be deployed in any architecture. Given that customer premise equipment can be a challenge to change or remove in the future, this silicon can remove that burden.

Industry alignment on this topic would drive down costs, making this a more viable option.

# Abbreviations

bps	access point
BW	bandwidth
dB	decibel
dBmV	decibels relative to one millivolt
DOCSIS	data over cable service specification
DS	downstream
ESD	extended-spectrum-DOCSIS
HFC	hybrid fiber-coax
FDX	full-duplex-DOCSIS
FTTP	fibre to the premises
GHz	giga hertz
Hz	hertz
ISBE	International Society of Broadband Experts
LE	line extender
MB	mini bridger
MDM	modem
MER	modulation error rate
MSO	multiple system operator
PoE	point of entry device
QAM	quadrature amplitude modulation
Rx	receive power
SCTE	Society of Cable Telecommunications Engineers
SNR	signal to noise ratio
Tx	transmit power



# **Smart Home & Smart Notifications**

## **Multimodal event detection and notification**

A Technical Paper prepared for SCTE•ISBE by

**Henk Heijnen**

Project leader

Technicolor France

975 Av des Champs Blancs CS17616 – 35576 Cesson-Sévigné - FRANCE

+33 2 99 27 3000

henk.heijnen@technicolor.com

**Philippe Gilberton**

Senior scientist

Technicolor France

philippe.gilberton@technicolor.com

**Jean-Ronan Vigouroux**

Senior scientist

Technicolor France

jean-ronan.vigouroux@technicolor.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
1. Introduction .....	4
1.1. Exceptional sounds recognition .....	4
1.2. Multimodal sensors, abnormal situations detection.....	4
2. Exceptional Sound Detection .....	5
2.1. Introduction .....	5
2.2. Some definitions.....	5
2.2.1. Rare/exceptional sound events .....	5
2.2.2. Sound detector .....	5
2.3. ML based sound detection demystified .....	7
2.3.1. Sound detector service architecture .....	7
2.3.2. Sound detector functional architecture.....	7
2.4. Conclusion .....	16
3. Anomaly Detection.....	17
3.1. Introduction to Anomaly Detection .....	17
3.1.1. Introduction .....	17
3.1.2. An example of Anomaly Detection algorithm.....	17
3.1.3. Different kinds of anomalies .....	18
3.2. Anomaly Detection Algorithms.....	18
3.2.1. Algorithms performance evaluation: the ROC curve.....	19
3.2.2. Algorithms not using the temporality.....	20
3.2.3. Algorithms using the temporality.....	21
3.3. Evaluation .....	22
3.4. Conclusion .....	25
3.4.1. Choice of the algorithm .....	25
3.4.2. Kind of anomalies detected .....	25
4. Usage examples .....	26
4.1. Problem to solve.....	26
4.2. Existing solutions.....	26
4.2.1. Pre-configuration requirement .....	26
4.2.2. Cloud vs. Local processing.....	26
4.3. Proposed approach .....	26
4.4. Use cases .....	27
4.4.1. Sensors .....	27
4.4.2. Audio recognition .....	27
4.4.3. Actors, entities .....	27
4.5. Possible usages .....	28
4.6. More difficult usages, Geophone-related & others.....	30
5. Possible Implementations .....	31
6. Conclusion.....	32
Abbreviations.....	33
Bibliography & References .....	33

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 acoustic onset, offset time.....	6
Figure 2 Sound Detector Service Architecture.....	7
Figure 3 Functional Architecture .....	8
Figure 4 Sound Detector Functional Architecture .....	8
Figure 5 MFCC Computation Pipeline.....	9
Figure 6 SVM Vector View In The Case Of A Binary Classification.....	11
Figure 7 Typical NN Architecture With 2 Hidden Layers .....	12
Figure 8 Exemplary Of A Two CNN Layers Architecture.....	13
Figure 9 SoundNet NN architecture .....	14
Figure 10 k-NN algorithm.....	18
Figure 11 Receiver Operational Characteristic .....	20
Figure 12 GMM Model .....	20
Figure 13 Autencoder .....	21
Figure 14 Training and Detection using Temporality.....	21
Figure 15 simplified system-level architecture .....	28

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 : Classification results on the training, validation and test set using different audio feature extractors.....	15
Table 2 : SVM versus NN classifier performance comparaison.....	15
Table 3 : AUC of the scoring algorithms .....	24
Table 4 : Possible usages.....	28
Table 5 : Possible usages (2 <sup>nd</sup> level).....	30
Table 6 : Possible implementations.....	31

# **1. Introduction**

In this paper we will explain the outcomes of our research project on the detection of anomalies in the home.

The goal of our experimentation is to be able to detect such anomalies as unexpected sounds (dog barking, gun shot, baby cry ...) and reports from simple sensors measuring temperature or humidity for example.

## **1.1. Exceptional sounds recognition**

Because listening to in-home sounds 24/365 is very critical regarding privacy, all the processing is done locally, and only recognized events are sent to the cloud / backend.

The technology we are using (neural networks) requires very large sound databases during the model build (10000+ sounds per class), therefore, the models are pre-calculated offline and are downloaded to the recognition brick. We can attain extremely good recognition rates although we can't learn new sounds nor improve the model locally.

## **1.2. Multimodal sensors, abnormal situations detection**

Another strong constraint is that we do not want to force users to create scenarios for controlling the device behavior regarding sensors management. We have developed a self-learning, non-supervised technology that will allow our product to learn from "know good" situations and to report events that are outside of this zone without the user having to configure the product.

## 2. Exceptional Sound Detection

### 2.1. Introduction

This section aims to provide some insight into the tremendous gap that Neural Network (NN) architecture has delivered, over the last several years, in the sound classification space. This paper will also include a more detailed focus on current NN architectures, and their performances compared to one of the most popular classifiers relying on the Support Vector Machine (SVM).

### 2.2. Some definitions


First, let's provide the reader with some definitions that should be considered as well as the device requirements of a sound detector.

#### 2.2.1. *Rare/exceptional sound events*

A rare/exceptional sound event is defined as a sound event that occurs once within a fixed period of time which is significantly longer than the event duration itself. A typical example is a glass breaking that would last less than a few seconds and would only occur at most once every couple of months. In the acoustic Machine Learning (ML) domain it signifies that there is a significant unbalance ratio between positive (occurrence of sound) and negative (no occurrence of sound) samples. Some years ago, the ML algorithms were trained on almost continuous and well-balanced audio samples which give good results with controlled test configurations but not so good in a real situation. This imbalance situation constitutes a challenging but more realistic problem to solve for a Machine Learning which must learn on few positive samples. Furthermore, it involves the usage of a supervised learning approach to build a satisfactory acoustical model. For comparison purposes an unsupervised learning approach would need regular and recurrent positive samples which are not compliant with exceptional sound detection requirements.

#### 2.2.2. *Sound detector*

A sound detector is a device whose purpose is to detect a sound event within a fixed temporal but sliding window. As shown in the upper rare/exceptional definition, the sound detection has a coarse sound detection granularity which means that the exact time of the sound event occurrence presented figure 1 (onset and offset time) is not provided. Additionally, the sound detector will operate in a domestic environment in which background noise may exist. As will be detailed further in this paper, this will have some consequences on both the training performance of the ML and the audio dataset composition.

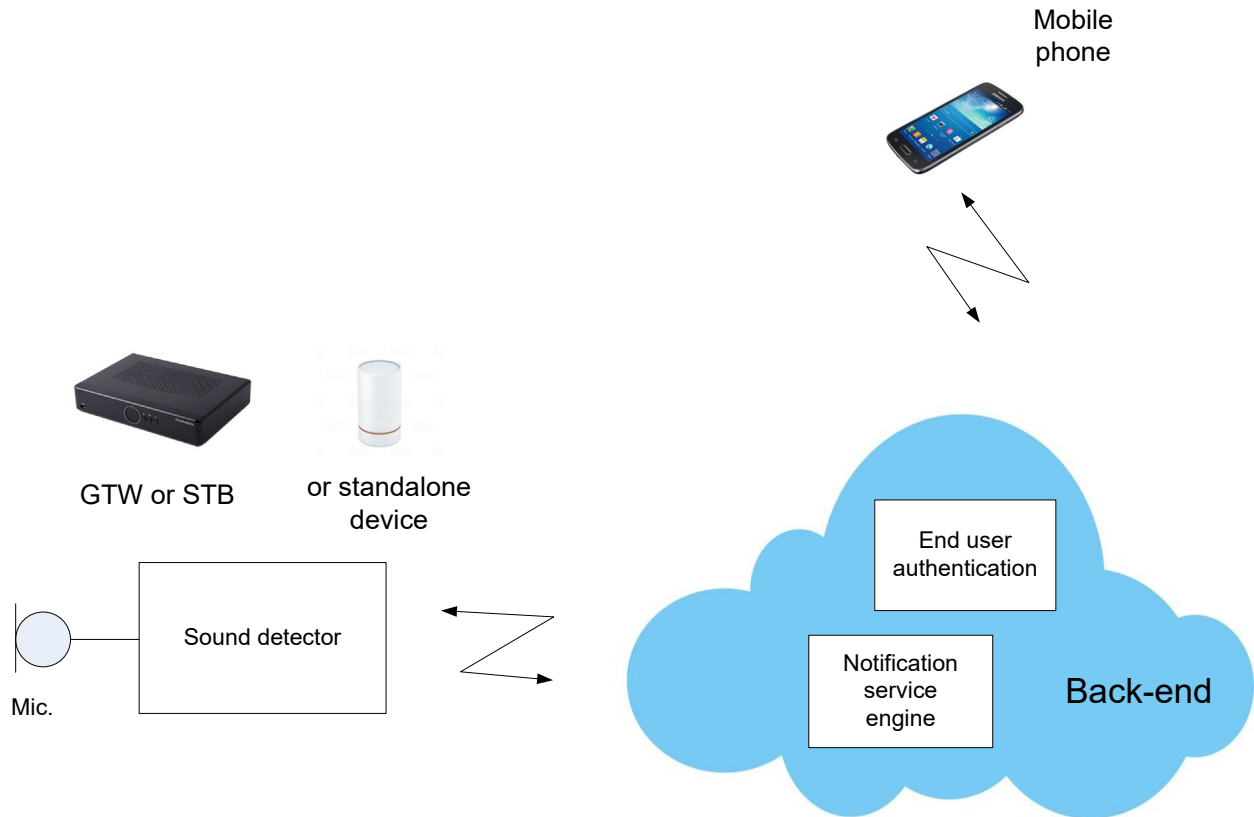


**Figure 1 acoustic onset, offset time**

This means that the ML must learn within an ambient noise environment and with a weakly annotated dataset. More specifically, a weak annotation means that, for each sample of sound used for training, the ML will only know that the sound of each class to recognize is present but without any indication of when it occurs within this audio sample. That is an important parameter to pay attention to as it makes easier to build the audio dataset, while another more complex approach, involving a detailed temporal annotation, would be too time consuming and not cost effective at the end.

## 2.3. ML based sound detection demystified

### 2.3.1. Sound detector service architecture

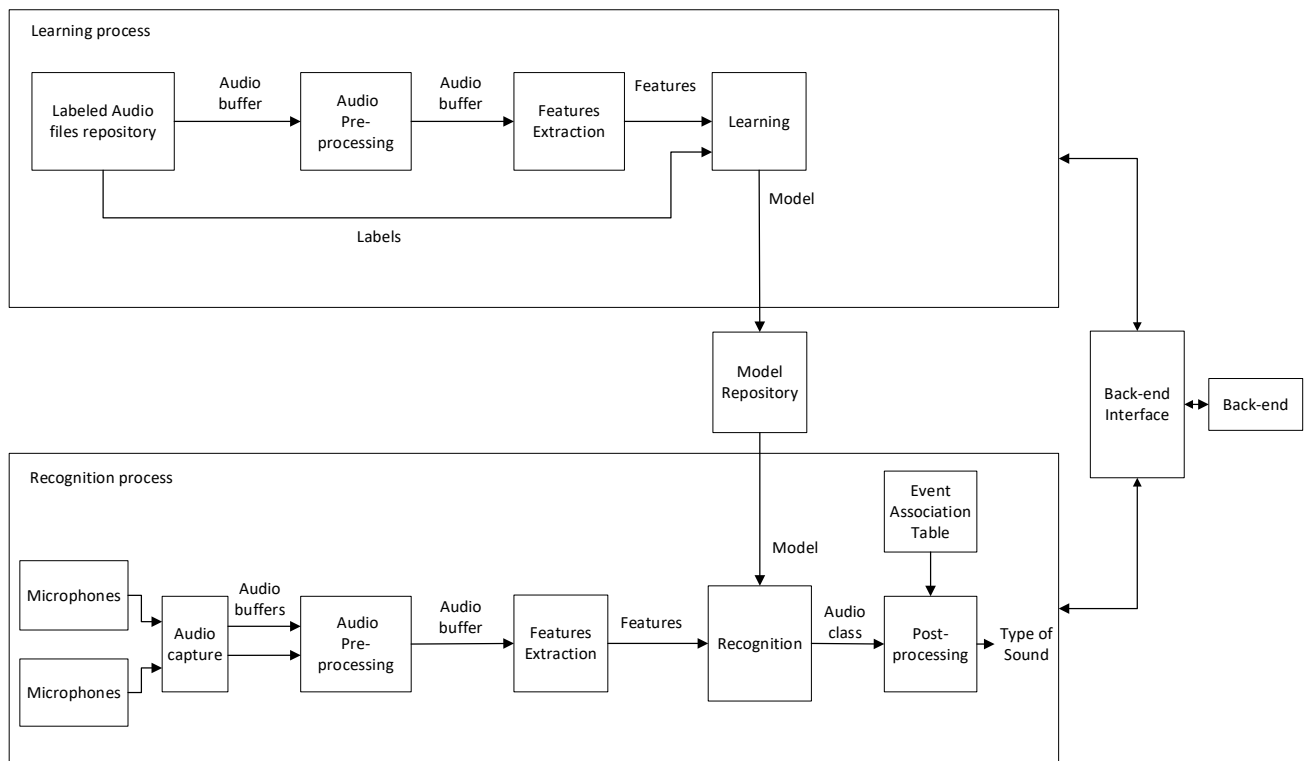


**Figure 2 Sound Detector Service Architecture**

This figure presents an example of a sound detection service architecture that would rely on a mono modal service approach (i.e. only the sound modality) but that could be easily integrated into a multi-modal approach, as exposed in the anomaly detection section (see section 3). In this simplified figure 2, the sound detector would provide the sound class label for each processed temporal window (i.e.: every 6 seconds) without sending and saving any audio samples captured in the household. The sound detector has been optimized to be embedded in small footprint hardware to preserve data privacy of the user (refer to §5). Such a local processing approach has a significant advantage versus cloud-based audio recognition systems like Alexa Guard because it protects user's privacy: the backend, with no access to the user data, will manage only the transmission of the notification to the user and/or the service provider and will not have access to audio samples. A smart notification system could be deployed according to the nature of the detected sound.

### 2.3.2. Sound detector functional architecture

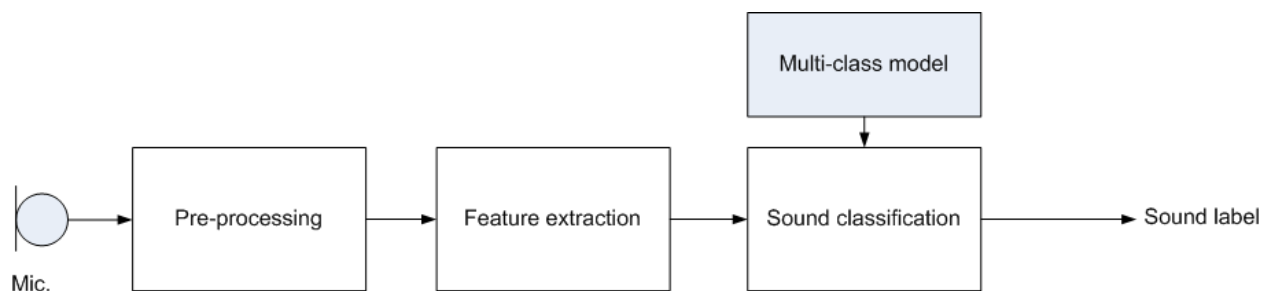
As an introduction to a more detailed description of the sound detection system based on ML, an overview of a functional architecture of the system is presented in the figure hereafter:



**Figure 3 Functional Architecture**

What can be noticed in figure3 is an architecture based on two blocks that operate two separate and sequential processing sessions. As in any ML system, those two required processes are first the training part, usually performed off-line and once to build the acoustic multiclass model, and then the acoustical detection part. As mentioned in paragraph 2.2 the ML approach is based on supervised learning which means that the acoustic model is trained on annotated dataset audio samples. This approach will justify building a huge dataset that is as diversified as possible (see paragraph 2.3.2.1).

Let's focus now on the sound detector functional architecture presented in Figure 4 Sound Detector Functional Architecture and a more detailed description of the design:



**Figure 4 Sound Detector Functional Architecture**

### 2.3.2.1. Dataset

Prior to training the supervised ML, getting the multiclass model and then evaluating its performance, a sound dataset shall be carefully built by following the typical folder organization:



A folder for the training set and a folder for the test set. Each of those 2 folders is subdivided into a number of folders according to the number of sounds to detect. Usually the file number ratio between training and test sets are 80% and 20% respectively as the learning effort must be put on the training session to get the best acoustical model as possible.

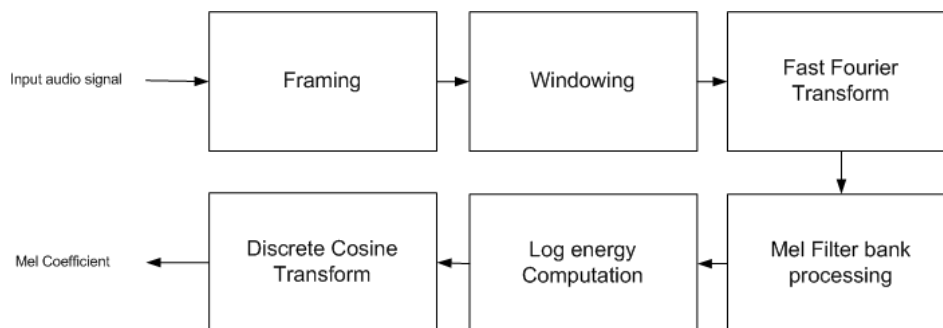
Each sound of the dataset is a mix of the sound event occurrences to detect and different background noises that fit with a realistic acoustical household environment. Many data sources are freely available like <https://freesound.org/> or <https://research.google.com/audioset/> for which a very convenient search engine tool is available. But even with such huge sound libraries, the acoustical diversity and sometime quality criteria are not reached, especially with sound classes that are not as popular (i.e.: hand clapping, snapping or coin rolling on a table). More recently, a new technique has emerged that uses data augmentation to dramatically diversify the dataset by adding in random manner small signal perturbations on the initial training set. Simple algorithms perform such synthetic data augmentations by adding noise, stretching time, changing pitch or speed or even adding acoustical room reverberation. This enriched training set is now much more suited to reflect the acoustic sound diversity present in different household environments.

### 2.3.2.2. Audio pre-processing

This function aims to capture audio samples from the microphone or the array of microphones with fixed parameters such as sampling rate, number of channels, compression algorithm if required and bit depth (i.e: 44.1 kHz, mono, uncompressed, 16bits). Optionally, frequency filtering is applied to the audio signal to minimize the effect of sensitivity dispersion and non-linearity of the microphone. A rms input power level estimation may also be performed to disable the classification process when the input power audio signal is under a fixed threshold value.

### 2.3.2.3. Features extraction

This function is essential in ML usage as it will extract the relevant acoustic characteristics that will be used to train the ML. There are different techniques available, the most popular of which is the computation of MFCC [1] (Mel Frequency Cepstrum Coefficient) presented hereafter:



**Figure 5 MFCC Computation Pipeline**

Prior to computing the features, framing and windowing functions are required. This is based on applying a sliding temporal window that divides the incoming audio samples into small frames (ranging usually from 10 to 50 ms). The frame duration determines the temporal resolution targeted by the system. To avoid acoustical artefact on the edge of the frames, a Hann or Hamming algorithm is applied with an overlapping ratio of 50%. The signal is then ready for FFT processing. As the signal is now transposed in the frequency domain, a filter bank is applied with a Mel frequency spacing scale. The number of digital filters constituting the filter bank fixes the number of Mel coefficients that are computed. The Mel coefficients

are followed by a normalization function that usually uses a logarithm function, but more recently, the normalization of Mel coefficients has been improved by using PCEN [2] (Per channel Energy Normalization) that have a better signal saliency property in the presence of background noise. The number of coefficients is a parameter to fix to a value usually ranging from 10 to 100 depending on the targeted spectral frequency resolution.


#### **2.3.2.4. Audio classification**

This function is critical as it performs the probabilistic classification of sounds per captured audio samples. In other words, it will provide a probability value per audio class for which the sum must equal to one. The detected sound will be the one which gets the highest probability value. A threshold can be applied at the output based on the maximum of the computed probabilities. From a satisfactory user experience perspective, the threshold must be determined carefully to get a fair balance between an acceptable low level of false positive (wrong classification of a sound event occurrence) and an acceptable low level of false negative (missing of a sound event occurrence).

To provide more focused insight on the state-of-the-art classification algorithms, only SVM (Support Vector Machine), the most popular one, and NN (Neural Network) approaches will be presented.

##### **2.3.2.4.1. SVM based algorithm classification**

The SVM (Support Vector Machine) is part of a family of algorithms well suited to classification, regression or anomaly detection problems. They were initially developed in the 90s and rely on finding the simplest way to separate the classes of data by maximizing their distance. The edge of separation of those data is also called the margin. For audio purpose classification problems, the classes of data correspond to the different type of sound to recognize. The acoustical data are represented by vectors which are the extracted features computed in §2.3.2.3. In a simple binary classification problem, the vectors located on the edge of the separation of the 2 classes are named the support vectors as illustrated in figure 6:




**Figure 6 SVM Vector View In The Case Of A Binary Classification**

In this simple example, the data are represented in a 2-dimensional space, where the edge of separation of data is the black line and the margin is the range between the black line and 2 blue and yellow lines. The support vectors that are closest to the edge are represented by circled blue and yellow dots. In this particular case, a linear separation is possible but in most actual cases this is not true, and other techniques are required. They are named kernels and allowed to separate data by projecting them in a higher dimensional space. Those kernels are based on polynomial or gaussian functions. SVM is convenient to setup as it requires few hyperparameters that are usually the regularization function, the kernel function and C which is a coefficient that penalizes more or less the cost of misclassification.

#### **2.3.2.4.2. NN based algorithm classification**

As an introduction, we will present the NN principle. Shown hereafter is a typical NN architecture:



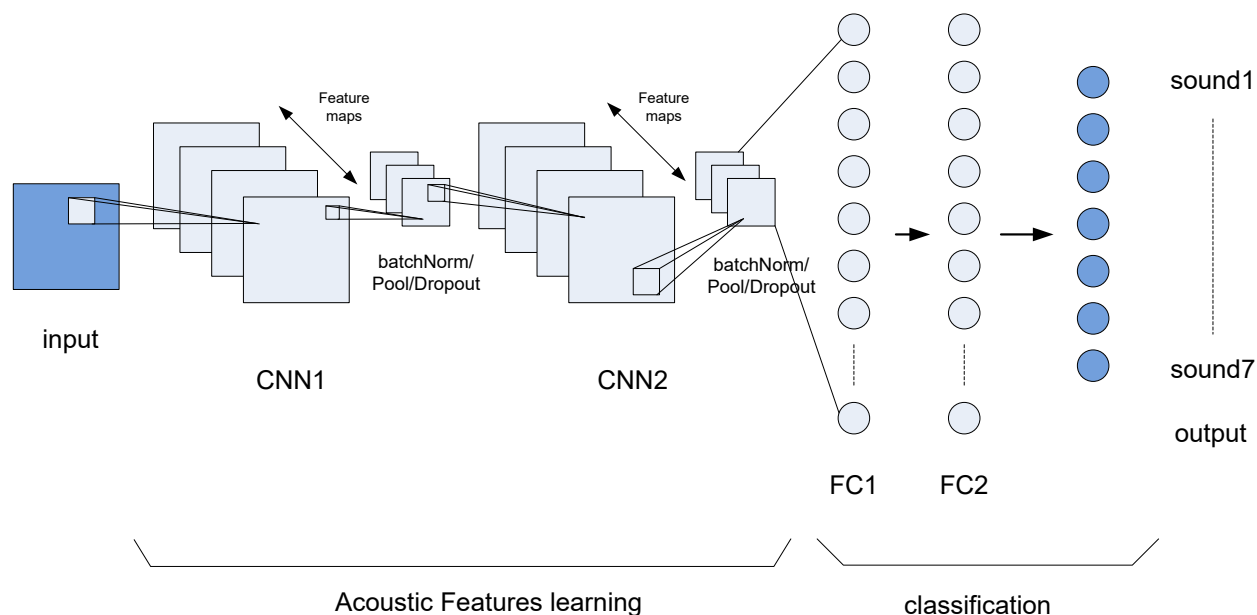
**Figure 7 Typical NN Architecture With 2 Hidden Layers**

The circle represents the neuron for which a weight and a bias are applied, and the arrow represents the link between one neuron to the next one, also going from the previous layer to the next one. In ML, a neuron makes a linear combination of input data (i.e.: extracted feature vectors) on which a bias value is added. The resulting output data value goes through a non-linear activation function (i.e.: hyperbolic tangent or sigmoid). This is performed for all neurons of the first NN layer (input layer). Then those output values are transmitted to the next NN layers (hidden layers) which will perform the same data computation and, finally, will end with an output layer that will provide an example of a classification problem, allowing the probability of each label to correspond with the audio classes that are necessary to be recognized.

What is remarkable is the simplicity of the elementary computation functions that are a combination of addition and multiplication. However, their number could be large as they increase from one layer to the next, except for the last layers. Compared to SVM, NN requires the selection of many hyperparameters to finally get to the most accurate multiclass model. The most common of those are typically the number of hidden layers, the type of layers (convolutional, recurrent, fully connected, ...), the number of neurons per layer, the activation function, the number of epochs, the optimization function. In addition to that some other intermediate functions are added like batch normalization, dropout or pooling for which specific parameters must be selected as well.

#### **2.3.2.4.2.1. Full trained NN model**

Full trained NN model means that the model is trained and tested with the entire available dataset (see §2.3.2.1). The NN architecture used for training relies on CNN for which a simplified architecture is presented below followed by several acronym definitions:



**Figure 8 Exemplary Of A Two CNN Layers Architecture**

The input is a 2-dimensional acoustic features matrix with columns and rows respectively corresponding to the number of acoustic features coefficients and frames. Then the NN architecture is split into 2 parts, the feature learning one including the CNN layers and the classification one including the fully connected layers that ultimately provide the probability of each sound class required to be recognized.

Feature maps are the result of convolutional operations using a small filter matrix performed at each location of the CNN layer. The number of filters determines the level of granularity you would like the NN to be trained on.

BatchNormalization is a normalization function performed on the incoming features to adjust and scale the activations. It helps learning in stabilizing the NN weight computation especially when the dataset is diversified and for which the feature values vary a lot from 1 class to the other.

Pooling (also named MaxPooling) is a downsampling function to reduce feature dimension and minimize overfitting. Usually the maximum value of a subpart of the filter matrix is computed which is the reason the name usually has the prefix of Max.

Dropout is a regularization function that randomly deactivates some neurons. The effect is that the NN becomes less sensitive to the specific weight of neurons. As a consequence, the NN model is capable of better generalization and is less likely to overfit on training data.

FC (Fully Connected) layer also named Dense layer executes a linear operation on the input vector and is located at the end of the architecture to transition from CNN layers. In a multiclass classification problem, for which a probability per class is required, a softmax function is applied in the last Fully Connected layer.


Another type of NN layer, less popular for sound classification because of the occurrence of model convergence problems and of a much longer training time, is the RNN (Recurrent Neural Network). The main difference from the CNN approach is that the next vector value depends not on the previous one but also on all the previous history. RNN has a memory of what happened in the past compared to CNN. The

last generation of RNN uses LSTM (Long Short-Term Memory) that improves during the training phase model convergence problem like gradient vanishing effect. Some results will be exposed in § 2.3.2.4.3.

### 2.3.2.4.2.2. Pre-trained NN model

The pre-trained NN model has become more popular meaning the model is already trained on a large dataset and is by consequence very well generalized for most of the classification/segmentation/object detection problems. Many pre-trained models exist for image processing (Xception, VGG19, ResNet50, InceptionV3, MobileNet, etc...) and fortunately some are available for sound classification as well like:

- VGGish(<https://github.com/tensorflow/models/tree/master/research/audioset>). It uses CNN architecture and it takes as a spectrogram an input dimension of 96 x 64 followed by 4 CNN and maxPooling layers. It ends with a 128-wide fully connected layer ready to be linked for example to a last FC layer for sound classification purpose.
- soundNet(<https://github.com/eborboihuc/SoundNet-tensorflow>). Its architecture is presented below:



**Figure 9 SoundNet NN architecture**

SoundNet [3] is composed of up to 8 CNN layers and 5 maxPooling layers. As shown on figure 9, SoundNet addressed the acoustic scene/object classification and significantly improved the state-of-the-art results.

The main advantage of using this technique is to benefit from the quality of the model that is well generalized because it was trained on a large dataset with many NN layers. Typically, VGGish used the audioset (<https://research.google.com/audioset/>) database which contains more than 2 million annotated videos representing around 6000 hours of sounds. What constitutes the main constraint is you must strictly follow the same input data format that was used for training the model.

Those pre-trained models could be used in 2 ways: first, as a feature extractor or second, to fine tune the pre-trained model to better fit it to your own dataset. In the first case it means you don't care anymore about how to extract features because the pretrained model, based on NN, will do it for you. The other benefit is the significant reduction of training time as only the last fully connected layers are required (with optionally the dropout function) because the NN layers were already used in the pre-trained model. In the second case,

there is not that much training time saved as the pre-trained weights of the model not only have to be updated on all or a part of the NN layers but also on your own dataset.

An interesting comparison on performance of classifiers, based on the usage or not of a pre-trained model, is summarized in this paper [4]. In this case, when using it as a feature extractor in a binary acoustic classification problem, it clearly demonstrates the accuracy improvement:

**Table 1 : Classification results on the training, validation and test set using different audio feature extractors**

Audio features	#dim.	Accuracy %		
		Training	Validation	Testing
MFCC	1212	83.10	83.84	84.05
MFCC	128	89.84	84.93	84.79
CQCC	1404	87.58	84.64	84.95
CQCC	128	86.59	85.53	85.29
SoundNet	512	89.45	86.00	86.87
VGGish	128	88.58	86.05	88.04
<b>VGGish+SoundNet</b>	<b>640</b>	<b>90.40</b>	<b>88.54</b>	<b>90.43</b>
VGGish+SoundNet + CQCC+MFCC	512	89.77	86.36	88.68
Baseline model of ASVspoof 2017	2223	76.31	76.30	72.63

#### 2.3.2.4.3. NN versus SVM performance comparison results

Another comparison was performed to justify the significant improvement of using NN versus SVM approaches. That was a performance evaluation conducted on 2018 for each classifier in using the same rare sound database extracted from the acoustic challenge DCASE2017 (<http://www.cs.tut.fi/sgn/arg/dcase2017/challenge/task-rare-sound-event-detection>).

The generated dataset consists of isolated sounds (Baby cry, Glass breaking and Gunshot) mixed with different types of background noise (Beach, Bus, Cafe/Restaurant, Car, City center, Forest path, Grocery store, Home, Library, Metro-station, Office, Park, Residential area, Train, Tram). The sound event occurred only once over the 30 second duration of incoming sound and located randomly within the file. The signal to background noise ratio was also adjustable and randomly chosen among 3 values, -6, 0 and +6dB. 500 files per class were generated both for training and test phases. The challenge relied on the low level of the event occurrence and the level of background noise. It meant the model had to be trained on a very unbalance training set with a lot of negative data (no event) or very little positive data (event). The results are summarized in the table 2 hereafter:

**Table 2 : SVM versus NN classifier performance comparaiison**

classifier	Event F-score				Event Error Rate			
	average	babycry	glassbreak	gunshot	average	babycry	glassbreak	gunshot
RNN	<b>93.10%</b>	92.20%	97.60%	89.60%	<b>13.00%</b>	15.00%	5.00%	19.00%
CRNN	<b>85.12%</b>	81.99%	97.03%	76.33%	<b>28.00%</b>	38.00%	6.00%	39.00%
SVM(kernel rbf)	<b>75.93%</b>	81.51%	86.67%	59.63%	<b>44.00%</b>	43.00%	27.00%	61.00%

CRNN is an architecture that combines CNN and RNN layers. The metrics used during the test phase are listed as follows:

- the event F-score ratio computes the ratio of successful event detections and positioning per audio test file
- the Event error rate computes the ratio of misclassification and bad positioning per audio test file

Table1 clearly shows that, in this particularly challenging but realistic audio dataset, the NN architecture outperforms the SVM one by at least 10% for the event F-score.

## **2.4. Conclusion**

Exceptional sound detection, that a decade ago was a difficult ML challenge to overcome, has now demonstrated promising results thanks to the availability of NN architectures. The emerging usage of the pre-trained models inspired from the last tremendous advancements done on imaging computation (face recognition and tracking, object detection and tagging, etc...) is another track to explore by refining the detection of more complex sequences of sounds composed of multi labelled occurrences.



## 3. Anomaly Detection

### 3.1. Introduction to Anomaly Detection

#### 3.1.1. Introduction

Machine Learning is an essential tool to implement Artificially Intelligent systems, that is systems analyzing and making decisions about their environment. The most well-known problems in Machine Learning are:

- Supervised Learning: learning to recognize classes of objects from examples,
- Unsupervised Learning: learning to find similarities between objects, learning to group or to represent objects without a class objective,
- Reinforcement Learning: learning to select the most profitable actions.


Aside from these points, Anomaly Detection is a lesser known subject that however has a huge interest for IOT devices. The idea is to learn to detect unusual or anomalous situations without examples of these situations. Put more clearly, the idea is to learn the normal behavior of a system and to trigger alarms when an unusual situation occurs.

This is particularly interesting for IOT as it is impossible for the end user to produce examples of anomalous situations (think of glass breaks or gunshots at home), or to label them. In addition, the variety of anomalous situations may be so huge (think of the different ways an old person may fall at home), that it looks much more reasonable to carefully model the normal situations and to trigger an alarm when the situation measured by the sensors deviates from normality.

It should be noted that the objective of these methods is to detect anomalous situations, and not to identify them. When an alarm is risen it will be necessary for the end user, or a third party (relative, neighbor, operator...) to check what happens, and see if there is really an urgency or if it is a false alarm, and to select the appropriate action.

#### 3.1.2. An example of Anomaly Detection algorithm

Being able to detect anomalous situations without having seen any of them may seem puzzling. Let us however present an algorithm that will convince the reader that this is feasible. The objective of an anomaly detection algorithm is to compute a score that is expected to be low (or even negative) in a normal situation, and high (and generally positive) in an abnormal situation. The k-nearest neighbors can be used for anomaly detection as follows. Let us suppose that all the inputs to the algorithm are vectors in a n-dimensional space, for instance a set of n simultaneous measures. Let us suppose that we have a collection of N normal points at our disposal. For a new point one measures the distance to the k-nearest neighbor, that is the distance of the k closest point. This distance is clearly an interesting candidate for being an anomaly score, as a point close to many others is likely to be a normal point, whereas a point remote to all the other points is likely to be an anomalous point.



**Figure 10 k-NN algorithm**

In general, one will take  $k$  equal to a few units or tens, to avoid producing a low score for repeated anomalous points.

### **3.1.3. Different kinds of anomalies**

The anomalies are unexpected events or situations occurring in the input stream. Following Chandola [5] we distinguish three kinds of anomalies.

- Point anomalies occur when an input vector is clearly different from all normal input vectors. This occurs for instance when a sensor is out of order and returns erroneous values, or an excessive temperature is measured, indicating a possible fire.
- Contextual anomalies correspond to situations which are abnormal in the context in which they occur. For instance, somebody preparing their breakfast in the middle of the night if this is not their habit.
- Collective anomalies occur when each measure in a segment of a series is normal, but the whole segment is anomalous. For instance, a fridge consuming no electricity for many hours.

It will appear that detecting point anomalies can be readily implemented, whereas the two other kinds of anomalies are more difficult to tackle.

## **3.2. Anomaly Detection Algorithms**

Anomaly detection algorithms aim at computing a score that will characterize the likelihood that a situation is abnormal. It is expected that a low (or negative) score will be produced in normal situations, and a high (positive) score will be produced in anomalous situations.

Algorithms can be divided in two kinds. Many algorithms consider samples independently. However, in IOT situations, it is reasonable to hypothesize that there is a continuity between the successive measures of a sensor (for instance the temperature in a house is not likely to change very fast in normal situations). Therefore, it is interesting to use this continuity principle to characterize normal situations.

Below, we present examples of these two kinds of algorithms, after a reminder on algorithm performance evaluation.

### **3.2.1. Algorithms performance evaluation: the ROC curve**

ROC curves are the classical way to evaluate detection algorithms, since the early days of radar technology.

A detection algorithm produces a score  $f(x)$  signaling the likelihood of a detection. If we fix a threshold  $S$  on the score this defines a classification function  $D(x)$ :

$$D(x) = \begin{cases} \text{True} & \text{if } f(x) > S \\ \text{False} & \text{if } f(x) \leq S \end{cases}$$


Knowing a ground truth on the objects, it is possible to characterize the performance of the classification by the numbers of:

- True Positives (TP): objects which should be detected and are actually detected,
- True Negative (TN): objects which should not be detected and are actually not detected,
- False Positive (FP): objects which should not be detected and are actually detected,
- False Negative (FN): objects which should be detected and are actually not detected.

The performance of the classification function for a value  $S$  of the threshold is defined by the two ratios:

- True positive rate:  $\text{tpr} = \frac{\text{TP}}{\text{TP} + \text{FN}}$
- False positive rate:  $\text{fpr} = \frac{\text{FP}}{\text{FP} + \text{TN}}$

Letting the threshold  $S$  vary defines a curve named the Receiver Operational Curve, located in the square  $[0, 1] \times [0, 1]$ . The area under this curve AUC is the measure of the performance of the algorithm. It shows if the algorithm can have a high true detection rate with a low false positive rate.




**Figure 11 Receiver Operational Characteristic**

The AUC is expected to be as close as possible to 1.0. If AUC is equal to 0.5 then the algorithm behaves as a random decision.

### **3.2.2. Algorithms not using the temporality**

Many algorithms can be used to define an anomaly score, considering the successive measures as independent realizations.


Gaussian mixture models (GMM) model the series of samples as realizations of a mixture of Gaussians. Each sample is supposed to be produced by a two-step process: first select a Gaussian in a set of  $n$  Gaussians, each having a probability, and then draw a sample using this probability. The parameters of the model can be learned using the EM (Expected Maximization) algorithm [7].



**Figure 12 GMM Model**

One-Class SVM is another example of classical algorithms aiming at detecting anomalies. The idea here is to separate the examples from the origin in a high dimension space, using a non-linear transformation.

Neural algorithms (see 2.3.2.4.2) have been used for detecting anomalies. A well-known example is the use of autoencoders. In this algorithm the parameters of the neural network are optimized to reconstruct the input vector on the output, after a projection in a lower dimension space in the middle of the network. In this case the anomaly score will naturally be the reconstruction error.



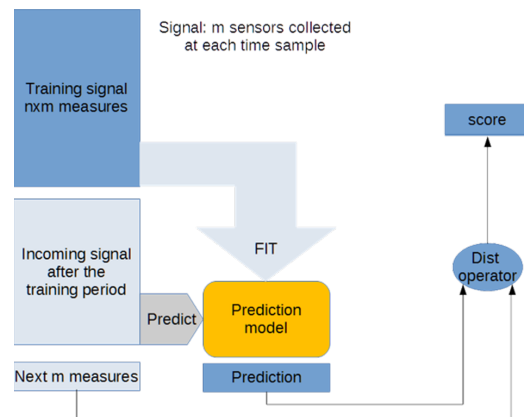
**Figure 13 Autencoder**

In an autoencoder the parameters of the network are optimized to make the input  $X$  of the network as close as possible to output  $Y$ , despite a projection on a lower dimensional space.

In general, these algorithms exploit the fact the input vector in normal situations does not take all the possible values of the sensors, but rather confines to specific regions of the input space, such as regions around specific points. In other words, this signal is compressible.

### **3.2.3. Algorithms using the temporality**

In IOT situations the sensors tend to measure continuous data, such as temperature or pressure. The next temperature measure is in general very close to the previous one, and a deviation to this is generally a hint that something abnormal is happening. It is therefore pertinent to set-up a prediction algorithm that will try to estimate the next sensor measures, and to use to prediction error as the anomaly score.



**Figure 14 Training and Detection using Temporality**

In this case a segment of the input signals is used to train a predictor. This predictor is used for a certain period, until the input signal has changed or after a fixed amount of time, when the predictor is re-trained. At each time step the already known examples are used to predict the next measures. The actual measure is compared to the prediction, and the deviation is the anomaly score. Many algorithms can be considered to implement the predictor.

Classic non-neural algorithms have been used for a long time to implement predictions.

The moving average is an example of such a simple algorithm. In this case the next sample is predicted by:

$$\hat{x}_i = \frac{1}{N} \sum_{k=1}^N x_{i-k}$$

This simple algorithm does not imply any training phase. Similarly, the exponential smoothing uses the previous measures and the previous prediction to produce the estimation of the next measure.

$$\hat{x}_i = \alpha x_{i-1} + (1 - \alpha)\hat{x}_{i-1}$$

In this case all the previous measures (potentially an infinite number) have an influence on the measure, with an exponentially decreasing weighting.

Holt-Winters algorithms are an example of multi-level regression, enabling it to consider seasonality effects in the data.

We have evaluated many neural architectures for providing the estimation of the next measure. The simplest architecture is a fully connected neural network, which takes the already known measures as input and produces the estimation of the next measure. See paragraph 2.3.2.4.2 for a description of dense neural networks.

Such a network can be implemented using a well-known library such as Tensorflow or Pytorch, trained using a retro-propagation algorithm.

Convolutional networks are an evolution of these networks, where weights are shared on the first layers to implement convolutional filters.

Finally, LSTM (Long Short-Term Memory) is a recurrent architecture, where the output of the network is reused to make the next prediction.

### 3.3. Evaluation

We have evaluated the algorithms on a dataset provided by EDF [6], the French electricity provider. This dataset is composed of measures of voltage, intensity and power in different part of a house, one measure per minute over a duration of four years. Each measure point is composed of seven individual measures.

We have introduced known anomalies on this dataset to measure the performance. The anomalies are:

- Out of order sensor: random values in the same min/max interval that the initial distribution,
- Activities at unexpected times: permute some segments of the data,
- The fridge is stopped: subtract the fridge consumption from the meter in the related room.

We recognize here instances of point anomaly, contextual anomaly and collective anomalies.

Different algorithms were tested on this dataset and the anomalies specified above. The algorithms are not optimized to try to obtain the best results for each of them, but are rather parametrized to operate in comparable conditions:

- The temporal algorithms use a buffer of one-hour length (60 measures) to make the prediction,
- The algorithms are periodically retrained after  $60 \times 24 \times 61 = 87840$  measures, that is every two months,
- The dataset for training the algorithms (for the ones that need a training) is equal to at most eight months.
- The algorithms are trained on the normal data and the performance is measured on anomalous data. This is to avoid the problem of deciding what to do in the training with the data that is declared anomalous: skip it? Replace it by most probable values? This renders the evaluation of the algorithms far easier.

The conditions seem to be realistic for the data that we have used: we expect to operate the different algorithms on a small device, with limited memory and computation resources, as described in section 5.

This enables us to estimate what is the relative performance of the different algorithms in comparable conditions. It does not give the optimal performance that could be reached if the parameters of an algorithm were optimized using for instance a grid search method.

The algorithms are implemented using Python 3.7, Scikit-learn, Keras and Tensorflow. The detailed parameters of each algorithm are as follows:

- Moving average:
  - The prediction is the moving average of the series on the prediction window, that is one hour (60 elements).
- Exponential smoothing:
  - The alpha factor (see paragraph 0) is equal to 0.01
- GMM likelihood:
  - The number of kernels is equal to 16, the covariance for each kernel is a full matrix.
- Dense Neural Network:
  - The size of the input layer is always 420 (60 samples of size 7). The output layer is always of size 7. Different number of intermediate layers with 60 or 420 neurons have been tested. Different activation functions: *relu*, *tanh*, *sigmoid* have been tested.
  - The network noted *60 relu sigmoid* below corresponds to a network with 420 neurons on the input layer, then *relu* as activation function on an intermediate layer of 60 neurons, then *sigmoid* as activation function, then the output layer which has size 7.
  - The network noted *420-60 tanh, tanh, tanh* is a network with 420 neurons on the input layer, a first hidden layer of 420 neurons, a second hidden layer of 60 neurons and an output layer of 7 neurons. The activation function between the input and the first hidden layer, the first hidden layer and the second hidden layer, the second hidden layer and the output layer are all *tanh*.
- Convolutional Neural Network:
  - The network is composed of an input layer with 420 neurons, a convolution layer Conv1D with 20 or 40 filters of length 8, a max pooling layer (pool size = 2), a second Conv1D layer with the same configuration as the first one, a max pooling layer (pool size = 2), and a dense layer with 7 neurons. The activation function is the *relu* function.
- LSTM:
  - The network is composed of two layers of LSTM functions with 60 units.

The AUCs measured for each algorithm are reported hereafter. We have retained only the best results for each algorithm to limit the data.

**Table 3 : AUC of the scoring algorithms**

<b>Algorithm</b>	<b>Anomaly 1: Point anomaly</b>	<b>Anomaly 2: Contextual anomaly</b>	<b>Anomaly 3: Collective anomaly</b>
Moving average	0.832	0.663	0.507
Exponential smoothing	0.817	0.629	0.505
GMM likelihood	<b>0.961</b>	0.496	0.503
Dense Neural Network 60 relu, sigmoid	0.936	0.603	0.518
Dense Neural Network 60 sigmoid, tanh	0.923	0.622	0.511
Dense Neural Network 420-60 relu, relu, relu	0.818	0.521	0.492
Dense Neural Network 420-60 tanh, tanh, tanh	<b>0.929</b>	0.595	0.508
Dense Neural Network 420-420-60 relu, relu, relu, relu	0.855	0.536	0.448
Dense Neural Network 420-420-60 tanh, tanh, tanh, tanh	0.866	0.568	0.507
Dense Neural Network 420-420-420-60 relu, relu, relu, relu, relu	0.850	0.529	0.469



Dense Neural Network 420-420-420-60 tanh, tanh, tanh, tanh, tanh	0.925	0.579	0.509
Convolutional Neural Network 20 filters, filter length = 8	0.900	0.647	0.513
Convolutional Neural Network 40 filters, filter length = 8	0.907	0.648	0.516
LSTM nb_units = 60	<b>0.967</b>	0.569	0.529

### 3.4. Conclusion

#### 3.4.1. Choice of the algorithm

The LSTM is the best algorithm (AUC = 0.967). The dense neural networks performance is lower (AUC = 0.929).

The GMM is nearly as good (AUC = 0.961), but its training time is far lower than the Neural Networks.

The numbers are not the optimal performance attainable with these algorithms, but it is reasonable to think, that for this problem of anomaly detection, considering that the training that must be done on a small device, the GMM is the solution to choose.

#### 3.4.2. Kind of anomalies detected

We see that the algorithms are good at detecting point anomalies but have virtually no performance for the other kinds of anomalies, contextual and collective. We have not tried to make any specific adaptation of the algorithms to these kinds of anomalies. But it is clear that:

- Contextual anomalies, here where the context is the time, could be detected if we would use specific models for specific times of the day (the day is divided in three or four periods for instance). However, the learning time would be much longer.
- Collective anomalies could be handled here using the integration of the features on different time periods. This would help to detect a null consumption if the normal consumption is generally not null. However, this would entail a large delay in detection.

## 4. Usage examples

### 4.1. Problem to solve

Users are interested in home systems that will report events, mostly abnormal ones. For example, an unexpected sound inside their home could be reported (glass breaking ...).

### 4.2. Existing solutions

#### 4.2.1. *Pre-configuration requirement*

Existing “security” or “elderly care” systems already properly report events but are limited to detecting pre-configured situations that have been programmed by using, for example, a scripting language.

By default, all “non-pre-programmed” situations will be either ignored or reported, by default, as abnormal.

#### 4.2.2. *Cloud vs. Local processing*

Existing solutions generally use cloud-based processing, local HW is limited to sensors + sensor interface to the cloud (backend). This is for example the case for the “Hive” solution. Note that this solution is extremely intrusive in term of privacy as, for example, 24/365 sound capture could be sent to the cloud for recognition.

We provide LOCAL processing thus protecting privacy and limiting the volume of data flowing to the backend. This solution might appear as more costly because it requires a larger CPU. In fact, modern GW or STB or IoT-dedicated SoCs are now powerful enough and can run neural networks algorithms locally.

### 4.3. Proposed approach

The system will use AI to learn from a known-safe situation during a long period of time. A model will be created by polling all the sensors that are available. Notice that only the anomaly detection model is trained at home. The acoustic classification model is trained in lab by the manufacturer.

When the model being created is considered to be good enough, all the sensor data will be monitored and any data that drifts from the “normal behavior” will be reported as “abnormal”.

This two-phase-model (learning / operational use) will be continuously updated with new data.

The system is able to report events that are outside of the automatically learned “normal situation” and will do that without any user-level configuration or programming (no scripting language).

The system also proposes a more complex sensor dedicated to the recognition of exceptional sounds. This brick uses neural networks technology and requires 10’s to 100’s of thousand sounds to create a good enough model. It is not possible to learn from user inputs because of this huge data requirement, it is also not possible because, for some sounds, it is not possible to ask the user to produce the targeted sound (glass breaking ...).

## **4.4. Use cases**

### **4.4.1. Sensors**

The system is sensor-agnostic, any continuous variation data can fit. In the following use cases, we will consider these sensors:

- Temperature
- Light (color & level)
- Atmospheric pressure
- Sound level (sound pressure)
- Humidity
- Organic volatile gas
- Presence detector (passive infrared detection, PID)
- Geophone (very low frequency microphone (< 100Hz), vibration detector). OPTIONAL (expensive)
- (time / date): not really a sensor but a crucial element.

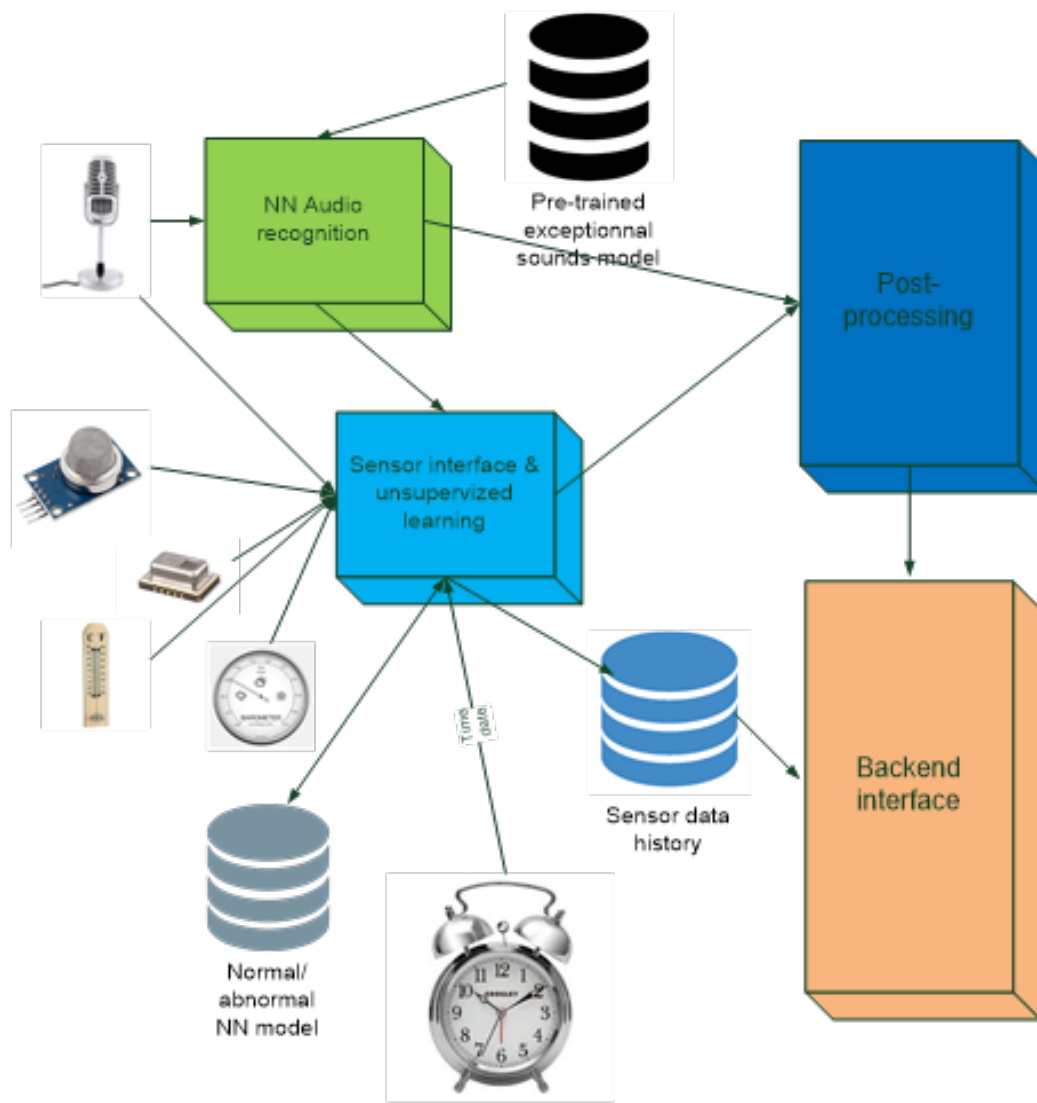
### **4.4.2. Audio recognition**

We will also use the output of the audio sound recognition brick that can detect pre-trained abnormal sounds:

- Dog barking
- Glass breaking
- Gun shots
- Baby cries
- Alarm (CO, fire ...)
- Human voice (no speech recognition, just “someone is speaking” information)

### **4.4.3. Actors, entities**

- NN Audio recognition + pre-trained sounds models recognize rare sounds (dog barking, glass break ...). The NN model is prepared offline (require huge computing capabilities)
- Sensor interface: drivers
- Unsupervised learning + NN model: local learning model aimed at detecting abnormal situations
- Post processing: prepares the data, does filtering if needed
- Backend interface: interfaces to the cloud
- (sensors)
- (clock): time & date



**Figure 15 simplified system-level architecture**

#### 4.5. Possible usages

**Table 4 : Possible usages**

What?	Detected as ...	Comments / limitations
Fire detection	Very fast temperature rise  High temperature is reached, for instance above 50°C  And/or  Fire alarm sound detection	Direct sunlight on the sensor might be detected as fire and can happen only a few days per year (high seasonality)

Freeze warning	In case of heating failure, the temperature falls under 5°C. There is a danger of frozen pipes in the house.	
Flood	Very fast humidity rise	
Device reliability	Magnetometer / gyroscope: device is moved	If the device is moved, its previous learning is not relevant anymore
Open window Heater / AirCo issue	Temperature rise / fall above limits – maybe with a humidity modification at the same time	
Dog barking for too long	Dog barking detection + time	Dog barking from the sound recognition brick filtered with “normal” duration by the unsupervised learning
Glass breaking, gun shots, fire alarm sound	Rare sound recognition	No filtering
Baby cry for too long	Rare sound recognition + time	Filtered by the unsupervised learning / time
Presence at wrong time	Presence detector + time + unsupervised learning	Presence detection at unusual time/date
Smoke / CO <sup>2</sup> / Gas	Smoke detector (if any) and/or alarm sound detection	No filtering?
Abnormal noise level (too loud, too long), for example, faucet kept open	Sound energy detector	(not managed by the sound recognition)  Detected by measuring sound energy (average), any loud and lasting sound will be recognized the same way
Door slamming	Pressure and/or loud sound	Abnormal pressure rise / fall
Shutters not opening or opened at the wrong time	Light / light color detector	

Light kept switched on at the wrong time	Light detector + time/date	
--	----------------------------	--

#### 4.6. More difficult usages, Geophone-related & others

**Table 5 : Possible usages (2<sup>nd</sup> level)**

<b>What?</b>	<b>Detected as ...</b>	<b>Comments / limitations</b>
Earthquake / explosion	Geophone activity	The geophone is optional
Steps	Geophone activity	The geophone is optional
Steps at the wrong time	Geophone activity & time	Steps at the wrong time (night?). Geophone is optional
Everything related to power (mains) usage	Currently, no sensor	Could be easy to use, needs adding a dedicated sensor (or set of sensors) on mains
Abnormal steps	Geophone activity	The geophone is optional. “abnormal” steps might be “too loud”, “too slow”, “too fast” ...

## 5. Possible Implementations

The recognition bricks have been designed in order to run on relatively low-performance hardware while keeping the important goal of running all the AI software locally.

Simplified requirements are:

- CPU: between 50 and 100% of one 32 bits ARM core (Cortex A53 class, ARMV7 instruction set). Note that most A53 implementations have 2 or 4 cores.
- Linux or Android. 256~512MB extra RAM required
- Between 2 and 4GB of storage / code space

Using 64 bits implementation (ARMV8) can improve the porting, having access to a TPU (or a GPU) can also help but it is not a strong requirement.

Several implementations are possible:

- Run the SW bricks inside a CE equipment, for example a set top box
- Run the SW bricks inside a gateway
- Create a dedicated standalone accessory
- Run the SW at the Edge / inside the cloud

**Table 6 : Possible implementations**

Solution	Pluses	Minuses	Potential improvement / remarks
Set Top Box	Cheap. STB CPUs are fast enough. Ample RAM/Flash.	Not necessarily well located. Might be switched off at random time. Sensors might be missing	Add remote mics
Gateway	Very cheap (but remote mics/sensors will add cost)	CPU & RAM might not be enough. Remote mics are required (can be costly). Always ON device.	Dedicated USB key (including CPU, RAM and sensors/mics)
Standalone device	Very convenient, multiple devices might cooperate. Sensors can be part of the accessory.	Somewhat expensive (but not so much compared to remote sensors + mics devices).	One intelligent accessory connected to simpler sensors boxes.
Cloud / Edge	No CPU/RAM/Storage limit while being cheap	Still need local sensors/mics (expensive). Does not meet the privacy constraint	Less privacy protection is a major issue inside the EU (GDPR). Privacy protection is one of the main differentiators of the technology.

## 6. Conclusion

The technology we developed allows seamless anomaly detection and reporting for many home-based situations.

Privacy is preserved because no data is sent outside of the home, with only event notifications uploading to the cloud.

Exceptional sounds detection is done using world class AI with models that are created using thousands of sounds for each class, thus achieving extremely good recognition levels.

Installation is simplified because the system can learn by itself from a “normal” situation and then can report any event that is not sufficiently in line with this situation.

Implementations can be either as a standalone accessory or as a SW brick running in existing CE-grade equipment.



## Abbreviations

ML	Machine Learning
RMS	Root Mean Square
PCEN	Per Channel Energy Normalization
FFT	Fast Fourier Transform
SVM	Support Vector Machine
NN	Neural Network
CNN	Convolutional Neural Network
FC	Fully Connected
RNN	Recurrent Neural Network
LSTM	Long Short-Term Memory
RBF	Radial Basis Function
EM	Expectation Maximization

## Bibliography & References

- [1] *Distance measures for speech recognition, psychological and instrumental*, Paul Mermelstein, 1976
- [2] *Trainable frontend for robust and far-field keyword spotting*, Wang, Y., Getreuer, P., Hughes, T., Lyon, R. F., & Saurous, R. A. (2017, March). In *Acoustics, Speech and Signal Processing (ICASSP)*, 2017 IEEE International Conference on (pp. 5670-5674). IEEE
- [3] *SoundNet: Learning Sound Representations from Unlabeled Video*, Yusuf Aytar, Carl Vondrick, Antonio Torralba, 30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.
- [4] *DISCRIMINATE NATURAL VERSUS LOUDSPEAKER EMITTED SPEECH*, Thanh-Ha Le<sup>1,2</sup>, Philippe Gilberton<sup>1</sup> and Ngoc Q. K. Duong, *Technicolor and Eurecom*, ICASSP2019, Brighton, UK
- [5] *Anomaly Detection: A Survey*, Varun Chandola, Arindam Banerjee and Vipin Kumar, *ACM Computing Surveys*, vol. 41, no. 3, July 2009.
- [6] *Individual Household electric power consumption data set*, Georges Hebrail, and Alice Berard, in *UCI Machine Learning Repository*,  
<https://archive.ics.uci.edu/ml/datasets/individual+household+electric+power+consumption>
- [7] A.P. Dempster, N.M. Laird et Donald Rubin, « Maximum Likelihood from Incomplete Data via the EM Algorithm », *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 39, no 1, 1977, p. 1–38 (JSTOR 2984875)

# **A Machine Learning Pipeline for D3.1 Profile Management**

A Technical Paper prepared for SCTE•ISBE by

**Mahe Harb**

Senior Principal Data Scientist

Comcast

1800 Arch Street, Philadelphia, PA 19103

267.260.1846

mahe\_harb@comcast.com

**Jude Ferreira**, Comcast

**Dan Rice**, Comcast

**Bryan Santangelo**, Comcast

**Rick Spanbauer**, Comcast

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
DOCSIS 3.1 Hardening .....	5
1. Standardized CMTS configurations and software .....	5
2. HFC Roll-off Reduction .....	5
3. PHY-Link Channel (PLC) Location .....	6
4. Ingress.....	6
5. Partial Channel/Service impairment handling.....	7
Problem Statement.....	9
Overview of Solution.....	15
MER & Time.....	18
Core Algorithm .....	26
6. Clustering of Modems .....	26
7. Modulation Efficiency Assignment .....	26
8. Segmentation.....	26
9. Profile Consolidation (Pruning).....	29
Lab Environment .....	31
Pattern Detection.....	37
Future Work .....	40
10. FEC   40	
11. Up-Stream Signal Path Implications .....	41
12. Near-Real-Time Operation .....	41
Conclusion .....	42
Abbreviations.....	42
Bibliography & References .....	43

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1. Distribution of MER across billions of D3.1 subcarriers (plus shaped data points).....	4
Figure 2 - MER distribution for select OFDM Interfaces.....	6
Figure 3 – MER spectra for two cable modems showing LTE Interference.....	7
Figure 4 – Schematic illustrating the CM-STATUS message operation.....	8
Figure 5 - Example of Profile Flapping. Profile failure is followed ~25 sec later by recovery.....	9
Figure 6 - MER measurements for a group of 20 CMs shown on a dual y-axis plot.....	12
Figure 7 - Each of the same 20 CMs is now assigned a modulation profile from a pool of a total 5 profiles.....	13
Figure 8 - Architecture view of the PMA systems. It includes 3 components: Data Engine, Analytics Engine, and Configuration Manager.....	15


Figure 9 - Block diagram showing the Analytics Engine key data sources, core functions, and post generation activities. ....	17
Figure 10 - Distributions of variation in MER over time by OFDM interface for one CMTS.....	19
Figure 11 - Analysis showing the impact of MER variation on the stability of modulation efficiency assignments.....	20
Figure 12 - MER time samples for CM84. ....	21
Figure 13 - MER time samples for CM95. ....	22
Figure 14 - MER time samples for the same 20 devices collected over a period of 10 days. ....	24
Figure 15 - The same 20 devices assigned profiles that are based on the aggregated MER values rather than a single point-in-time MER.....	25
Figure 16 - Illustration of the segmentation process. ....	27
Figure 17 - Example illustrating segmenting profiles to satisfy vendor constraints.....	28
Figure 18 - Impact of segmentation on reassignment of bitloads as measured across Comcast's full D3.1 network.....	29
Figure 19 - Exploring capacity gain as the local (interface-level) and global (CMTS-level) constraints on the number of profiles are varied. ....	30
Figure 20 - Full footprint sensitivity analysis exploring change in capacity gain as the MER consideration period and MER selection percentile are varied. ....	31
Figure 21 - Development lab RF connectivity showing our current development lab signal flow. ....	33
Figure 22 - MER data received from a D3.1 modem operating on our HFC plant.....	35
Figure 23 - Spectrum analyzer capture of the SDR generated spectrum based on MER data shown in Figure 22. ....	35
Figure 24 - Example of an impairment generated via SC-QAM carriers from an otherwise unused port on the CMTS.....	36
Figure 25 - Example MER charts for sweep generator insertion points in OFDM channel. ....	39
Figure 26 - Example of detected LTE Interference patterns, enriched with mobile carrier license data. ...	39
Figure 27. Illustration of subcarrier configuration in OFDMA. In OFDMA, contiguous subsets of subcarriers are assigned to different users. ....	41

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Minimum MER values that support the corresponding modulation. ....	11
Table 2 - List of Algorithm tuning parameters .....	30

# Introduction

We have entered an era where leveraging Machine Learning to optimize the performance of cable access networks is possible and, perhaps, even a must. The fast arising opportunities in this realm are due to advances in cable technology and increasing investment in data science functions across organizations. Specifically, DOCSIS 3.1 (D3.1) includes Orthogonal Frequency Division Multiplexing (OFDM), enabling the possibility of tailoring the modulation of OFDM channels to realize much improved spectral efficiency and impairment resiliency. Additionally, due to the nature of the wider OFDM channels, Comcast identified several opportunities and key deployment challenges affecting network stability & performance. As part of operationally hardening D3.1, it became clear that an effective modulation Profile Management Application (PMA) is essential for operating D3.1 to its full potential. The initial perspective -- that PMA was an optimization technique to maximize capacity in the future -- changed to a conclusion that PMA is really a table-stakes feature required to ensure network stability, manage operational metrics, and ensure a great customer experience. This document describes the profile management solution developed to address these challenges. As a primer to the ensuing discussion, consider the distribution (shown in Figure 1) of Modulation Error Ratio (MER) collected from Comcast's entire population of D3.1 devices. The distribution, while encapsulating information aggregated across billions of subcarriers, hints at the core idea of PMA: since the quality of the spectrum varies across the network, customizing modulation across subcarriers and devices holds the opportunity to enhance network performance in terms of increasing both capacity and resiliency. The goal of PMA is to pursue this ideal.



**Figure 1. Distribution of MER across billions of D3.1 subcarriers (plus shaped data points). The vertical gray lines indicate recommended thresholds for the respective modulation efficiency. The vertical blue and red lines represent distribution mean and median respectively.**

The paper is organized as follows: **DOCSIS 3.1 Hardening** describes our efforts in addressing D3.1 deployment challenges; these efforts are an important precursor to PMA. **Problem Statement** introduces formulation of the PMA problem. **Overview of Solution** presents the high-level PMA solution architecture. **MER & Time** describes our solution for addressing MER variation in time. **Core Algorithm** introduces the algorithm developed for constructing D3.1 Profiles. **Lab Environment** describes establishing a lab for testing the PMA solution. **Pattern Detection** describes a host of algorithms aimed at detecting impairments; these are complimentary to the PMA effort. **Future Work** comments on the future evolution of the PMA algorithm.

# DOCSIS 3.1 Hardening

After initially deploying D3.1 across the Comcast network, the operational service performance was evaluated and compared to DOCSIS 3.0 (D3.0) customer metrics. The initial analysis indicated that D3.1 was under-performing D3.0, based on a variety of operational metrics including truck rolls and call-in rates. At first reaction, this seemed very unlikely. All theoretical models and lab testing had shown that D3.1 is fundamentally more robust, with higher capacity based on the principles of OFDM technology, including the advanced Low Density Parity Check (LDPC) algorithm used in Forward Error Correction (FEC). As the root cause of the operational challenges were evaluated, a consistent theme emerged: While D3.1 is fundamentally much more capable than D3.0, it requires operational consideration when deploying, in order to fully realize its potential. Therefore, operational hardening of the D3.1 solution is required.

We began the D3.1 hardening effort with the goal of achieving parity of operational key performance measures relative to our mature D3.0 services. Along the way, we identified several opportunities to improve processes, configurations, and in some cases, entirely new techniques that ultimately would improve the service and the customer experience. While many of these were related to software maturity of the gateway, the operational metrics were equally challenged for customer-owned cable modems (CM). Upon our initial allocation of root cause, we believed that about 40% of the issues were related to D3.1 technology. We quickly moved to develop a collection infrastructure, to gather data and improve our visibility of the network, and diagnosed many specific customer challenges to isolate the key D3.1 operational challenges. In parallel, we began working on a solution for PMA.

Several key areas were identified for operationalizing D3.1 and are briefly described here, supporting the essential nature of profile management for D3.1.


## 1. Standardized CMTS configurations and software

Inconsistencies were discovered in detailed D3.1 CMTS (Cable Modem Termination System) configurations and software (SW) across the network and across CMTS platforms. In addition to SW updates, the primary configuration adjustments included setting the Cyclic Prefix to 512 samples, the roll-off period to 256 samples, and the time-interleaver to 16 OFDM frames.

## 2. HFC Roll-off Reduction

The initial spectrum deployment targets were 96 MHz to be placed in the highest spectrum of the Hybrid Fiber Copper (HFC) network above the existing video and D3.0 channels. In many service groups, a 96 MHz OFDM channel required that some of the OFDM channel be located in the HFC design roll-off area; where roll-off refers to spectrum above formal plant design. For example, in a 96 MHz channel, in some instances, the highest frequency was at 774 MHz for a 750 MHz HFC network, resulting in 24 MHz of channel bandwidth in the roll-off. In many cases, this was just fine, as this spectrum had been vetted and used previously for additional SC-QAM channels. However, in other cases, such as the example in Figure 2, the roll-off spectrum had reduced MER with the higher attenuation and response of the network, depending on N+x cascade length, where x represents the number of amplifiers in the signal path between the node (N) and the customer. In Figure 2, the red line represents the level for a flat QAM-256 modulation. The solid colored lines on each chart represent select percentiles in the distribution of MER per subcarrier. As can be seen in many of the channels, 75% of the MER samples for the modems will support modulation higher than QAM-256, but some will not support QAM-1024 or 2048. However, a small number of modems have degraded MER level in the roll-off region. To mitigate this, in the short term, we reduced the size of the channel back to the HFC roll-off edge, thereby reducing capacity. For the longer term we are able to

run into the roll-off spectrum effectively by using the PMA solution to modulate the subcarriers consistent with the network roll-off characteristics.



**Figure 2 - MER distribution for select OFDM Interfaces. The dashed red line represented the QAM-256 threshold. While the spectrum supports higher modulation overall, the roll-off region (upper edge) has MER levels falling below QAM-256.**

### 3. PHY-Link Channel (PLC) Location


On some of the CMTS population we noticed a high number of PHY-Link Channel (PLC) CM-STATUS 21 failure messages, along with PLC error rates. Some of these were resolved by fixing some CMTS and CM software. In several other cases, we noticed that the PLC channel was placed within spectrum used by the mobile carriers and was failing due to ingress. To resolve these issues, we performed analysis described later in this paper to join together OFDM channel locations with spectrum licenses obtained from a spectrum allocation database. As a result, we were able to find “gaps” in the spectrum allocation for each county associated with each HFC node and, through a rule set, identify the most common spectrum locations to reduce the probability of overlap between our PLC channel and offending mobile carriers. In most regions we were able to find a single spectrum location based on the HFC bandwidth (625, 750, 860, and 1000 MHz) and OFDM channel location to consistently place the PLC in a lower risk location. In several network regions, different localities required different PLC options where the mobile spectrum is widely licensed and deployed. We also developed an algorithm based on MER that could be used to find the best possible spot for the PLC channel, when there was not a great opportunity for lower risk, due to mobile spectrum licensing. We moved all of the PLC locations across the network to mobile spectrum gaps and reduced the rate of PLC failure events. The PMA software will be able to recommend optimal PLC locations based on statistical analysis of MER per subcarrier. The PLC current locations are shown in charts as described in Section 7 (Pattern Detection), and alternate locations are identified by finding the best possible 400 kHz or 6 MHz of bandwidth based on MER.

### 4. Ingress

We identified a variety of ingress sources in the network; many were based on wireless ingress such as LTE (Long Term Evolution). Several wireless ingress issues were caused by uncoordinated configuration of

access network systems. For example, a small number of ingress sources traced to the use of our sweep generator systems placing insertion points or QAM carriers into the service group, which were not moved out of the spectrum when the OFDM channel was added. As a result, the sweep generator moving through the spectrum with a transient tone every 6 MHz introduced errors that were difficult for the CM-STATUS messaging-based management to deal with.

For example, in Figure 3, the LTE interference in this case was at a level that periodically caused the OFDM channel to go into partial service or move back and forth between profiles at such a high rate that traffic forwarding performance was severely degraded. Many of the subcarriers would easily support QAM-2048 or QAM-4096, but for a small section of the channel, even QAM-256 was not supported well all of the time. With many of these types of smaller impairments, and from a spectrum perspective, it was clear that a flat modulation profile would not be as effective in achieving the high capacity performance potential of D3.1, and a PMA solution would be needed to achieve the anticipated D3.1 capacity benefits and stabilize the network.



**Figure 3 – MER spectra for two cable modems showing LTE Interference. The dashed red line corresponds to the MER threshold that supports QAM-256.**

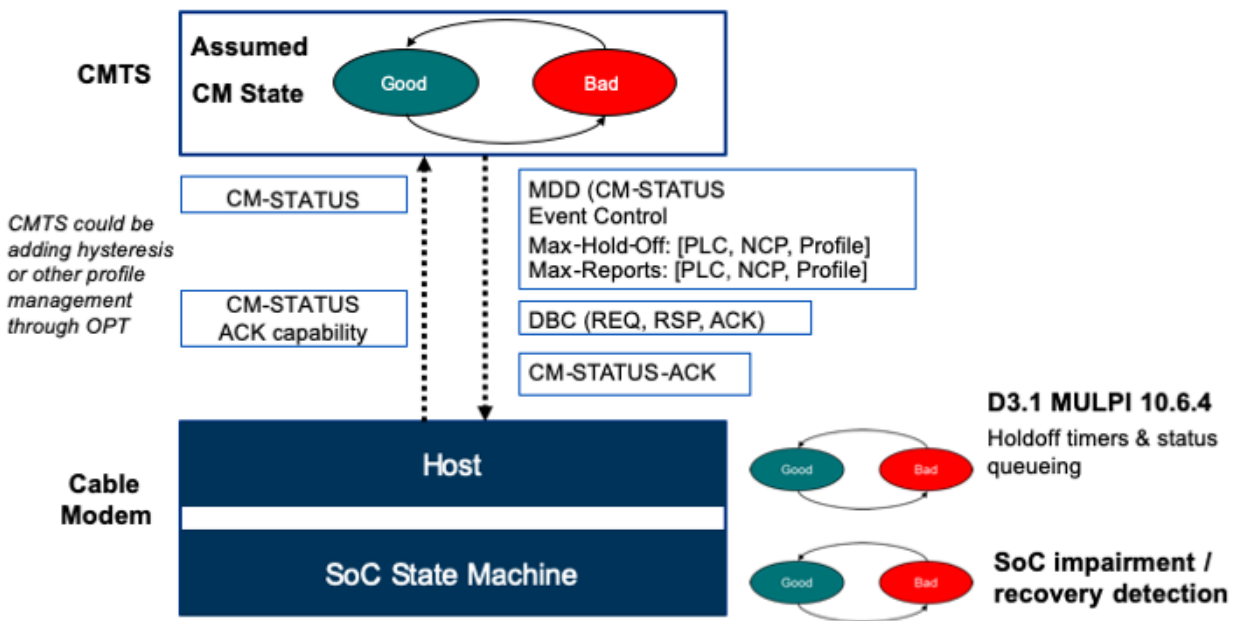
A high percent of OFDM channels had some subset of subcarriers across some subset or all cable modems in the service group that had RF physical layer impairments. For D3.0, when these impairments occur, the partial service and bonding mechanisms, along with a very mature capacity planning process, protects the customer experience very effectively. If one or two channels are impacted, there are more than 30 additional D3.0 channels available to provide a seamless customer experience. Respectively, without a PMA solution on an OFDM channel, a flat modulation profile is very inefficient from both a capacity and robustness perspective. Reducing modulation for all subcarriers, or losing 96 or 192 MHz of capacity, is a much more significant capacity and network stability challenge than a corresponding single D3.0 SC-QAM 6 MHz channel.

## 5. Partial Channel/Service impairment handling

Often, in the presence of transient ingress, or noise levels near the modulation boundary, we discovered stability issues with “profile flapping” (the condition where traffic is moved between two different profiles, back and forth, sometimes leading to instability or slow performance) based on our configuration for managing CM-STATUS messages. The CM-STATUS approach (illustrated in Figure 4) for notification from the CM to the CMTS on the health of the OFDM profile and channel elements is fairly complicated, with 3 primary state machines that must be tuned to work together, and are still immature in their implementation. The state machine, on the silicon System on Chip (SoC), makes decisions based on vendor-proprietary algorithms as to when the OFDM channel is impaired or when the channel is considered recovered from an impairment. Additionally, the time scale for consideration within those proprietary



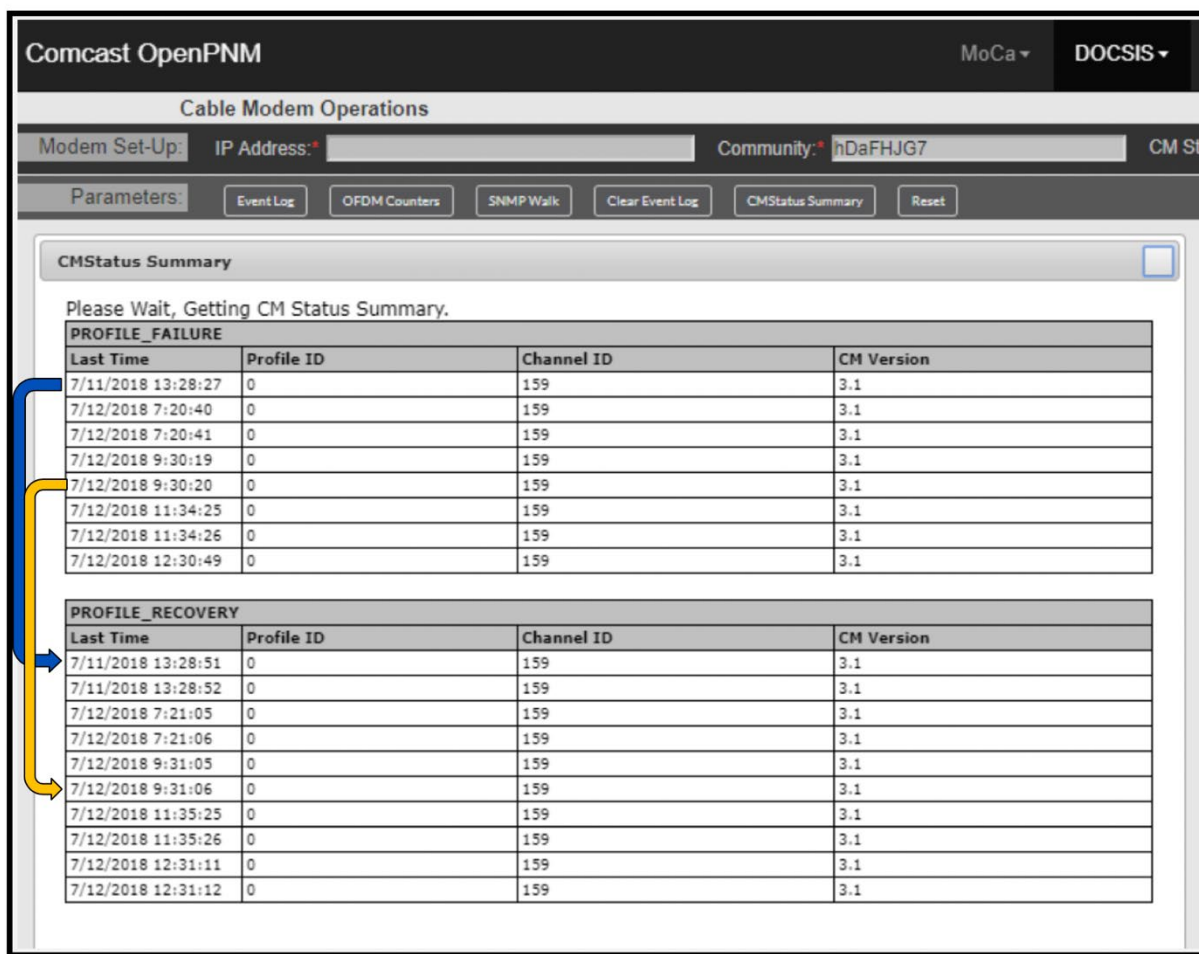
algorithms can vary. In one of our gateways, the event messages were by default sent on time frames of milliseconds or 100s of microseconds, depending on the noise characteristics. In the other gateway, the event messages were sent in timeframes of seconds or 10s of seconds. In either case, after the SoC decides on the condition, it messages the DOCSIS MAC layer, which sends it to the CMTS based on a configurable hold-off timer and state machine. These hold-off timers were set differently across CMTS platforms, creating variable behavior once the event is received. The CMTS decides how to react to the CM-STATUS event, and may apply proprietary controls in response to that message.



**Figure 4 – Schematic illustrating the CM-STATUS message operation.**

We were able to identify and reproduce a variety of scenarios that resulted in profile flapping. One example of profile flapping in the field is shown in Figure 5. Due to the holdoff settings, the CM state machine, and the CMTS configuration, the traffic would be moved from a profile that was marginal to a more conservative profile -- but then immediately moved back to the marginal profile based on the CM-STATUS messaging. With the 30 second holdoff timers configured on the CMTS, the modem would be stuck on the marginally performing profile for most of the time, with very brief transitions back and forth. The end result was traffic not getting forwarded through the CM, further demonstrating why a PMA solution is important. Profile decisions based on statistics from a data lake of channel performance, as opposed to the microsecond and second-by-second decisions that can be made by a CM and CMTS, helps to ensure network stability.

After a lot of testing and experimenting with different noise impairments recorded from the network, we were able to model and test CMTS- and DOCSIS-based settings that added the appropriate controls and hysteresis to stabilize the “profile flapping” challenges. The next section introduces the main idea behind PMA and the challenges that need to be addressed for a successful PMA deployment.



**Figure 5 - Example of Profile Flapping. Profile failure is followed ~25 sec later by recovery. The cycle of failure and recovery continues, and traffic gets blocked because the CM stays on an impaired channel.**

## Problem Statement

To support the increased capacity and stability of D3.1 OFDM channels, opportunities for optimization were identified. D3.1 supports the use of OFDM, in which the usable spectrum is divided into multiple narrow band (25 or 50 kHz wide) subcarriers. With OFDM, it is possible to tailor the modulation of signal to the specific spectral conditions of each of those subcarriers. There are two obvious benefits to gaining such flexibility in customizing modulation to the device (cable modem) and subcarrier levels: (1) Increasing the total capacity of an OFDM interface since the modulation scheme will no longer be dictated by those devices with relatively poor overall signal-to-noise ratios (S/N). (2) Increasing robustness of all devices by assigning impaired regions of the spectrum (as shown in Figures 2 & 3) a suitable modulation scheme, or even entirely blocking the impaired regions if needed.

To highlight the potential benefits of customizing the modulation profile, consider the plots, shown in Figure 6, of measured MER across the OFDM spectrum for a group of 20 CMs attached to the same OFDM channel. QAM-256 is the recommended modulation scheme for MER falling within the 27 to 30.5 dB range. Without the ability to configure modulation profiles, a QAM-256 modulation plan (8 bits/symbol)

would be adopted across the channel to accommodate the lower performing CMs. Yet, most CMs shown in Figure 6 have superior MER characteristics; thus, they could support the use of modulation higher than the assigned QAM-256. Additionally, one device (CM84) with an impaired region, likely due to LTE interference, has otherwise favorable MER. This device may benefit from using a modulation lower than QAM-256, exclusively in the impaired region, for added robustness. Therefore, it is evident that a flat, one-size-fits-all modulation scheme is far from ideal; this is true whether the device has a healthy or an impaired spectrum.

Ideally, each device would be assigned its own fully custom-matched modulation scheme. In practice, however, customizing modulation for each device is not practical due to protocol efficiencies and current CMTS capabilities. D3.1 allows customizing modulation through the concept of *modulation profiles*, where a profile defines a specific modulation scheme for each subcarrier, over the entire channel spectrum, for which devices within an OFDM channel can share a set of tailored profiles. While the number of allowed profiles per OFDM channel with the current D3.1 specification is 16, current vendor-specific implementations limit this figure -- in some instances, as low as 3 per OFDM channel (one of which is the control profile, known as profile 0 or profile A.) Restricting the number of profiles introduces one of the main optimization problems, in which the objective is to construct suitable profiles such that the total capacity of the interface is maximized within known MER conditions, so as to not negatively affect customer experience.

Let's define the problem mathematically by first introducing some useful notation. Assume that the usable spectrum contains  $L$  frequency subcarriers, and that there are  $M$  devices/cable modems (CM) attached to the same channel and  $N$  profiles ( $P$ ) to work with. The measured MER of a device is a vector of MER values of length  $L$ . A constructed profile on the other hand is defined as a vector of modulation efficiency values, also of length  $L$ . These are represented below for device  $j$  and profile  $k$ :

$$CM^{(j)} = [x_1^{(j)}, x_2^{(j)}, \dots, x_i^{(j)}, \dots, x_L^{(j)}], \quad i \in \{1, \dots, L\} \text{ and } j \in \{1, \dots, M\}$$

$$P^{(k)} = [y_1^{(k)}, y_2^{(k)}, \dots, y_i^{(k)}, \dots, y_L^{(k)}], \quad i \in \{1, \dots, L\} \text{ and } k \in \{1, \dots, N\}$$

A useful property of a profile is its total symbol size  $B$  defined as:

$$B^{(k)} = y_1^{(k)} + y_2^{(k)} + \dots + y_L^{(k)}$$

The problem at hand is to maximize the OFDM interface capacity. Following the convention in CableLabs seminal paper on PMA [1], the capacity is defined as:

$$C = S \times \left( \sum_{k=1}^N \frac{R^{(k)}}{B^{(k)}} \right)^{-1},$$

where  $S$  is the symbol rate (fixed to 25 or 50 Ksym/s),  $B^{(k)}$  is the total symbol size for profile  $k$ , and  $R^{(k)}$  is the ratio of number of devices assigned to profile  $k$ . The optimization problem has several constraints that will be discussed in this paper. The most basic constraint is that for any device  $j$  assigned to a profile  $k$ , the MER of the device across the spectrum must support the chosen modulation efficiency values. Let's assume that the MER thresholds for assigning modulation efficiency values in the downstream are represented by the mappings listed in Table 1 (obtained from the D3.1 specification [2]).


**Table 1 - Minimum MER values that support the corresponding modulation.**

<b>MER Threshold (dB)</b> <b>(<math>x_{threshold}</math>)</b>	<b>Modulation efficiency</b> <b>(<math>y</math>)</b>
0	0
9	2
15	4
21	6
24	7
27	8
30.5	9
34	10
37	11
41	12
46	13
52	14

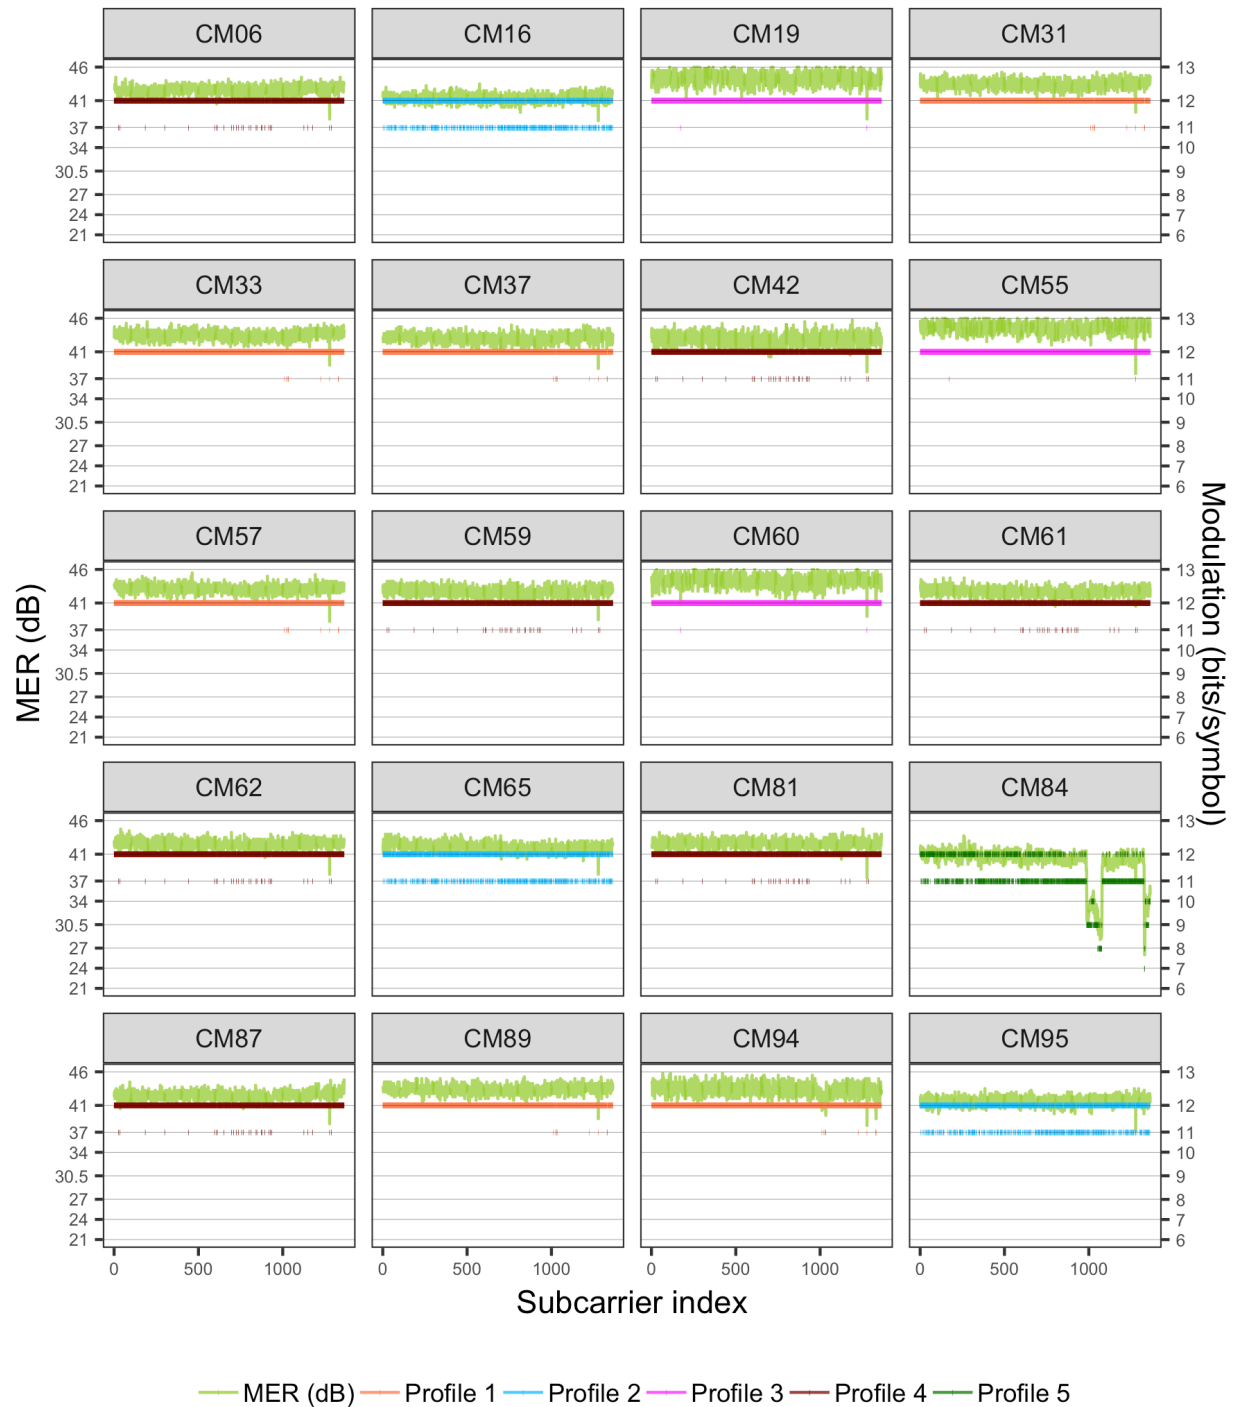
Then the constraint above is mathematically expressed as:

$$\forall x_i^{(j)}, x_i^{(j)} \geq x_{threshold} \text{ for the assigned } y_i^{(k)}$$

The statement of the problem so far is naïve and is used for illustrative purposes. In reality, the thresholds listed in the table are not hard thresholds but rather recommendations that have to do with keeping the symbol error rate below a certain level. Revisiting the thresholds and allowing some subcarriers to operate above their recommended modulation efficiency values should be permissible as long as it does not result in generation of uncorrectable errors and packet loss above some desired value (refer to Section 8: Future Work for a discussion of how we plan to integrate error rates into PMA).



**Figure 6 - MER measurements for a group of 20 CMs shown on a dual y-axis plot. The left y-axis indicates the MER value in dB and the right y-axis the corresponding modulation level (bits/symbol). In the current configuration, all devices use a QAM-256 modulation across the entire OFDM spectrum (1,364 subcarriers in total). It is clear from the shape and level of the MER curves that most devices benefit from using higher modulation while one device (CM84) exhibits impairments (possibly in two regions of the spectrum). Note that the device identifiers shown in the headers of panels are anonymized versions of the MAC addresses of devices.**



**Figure 7 - Each of the same 20 CMs is now assigned a modulation profile from a pool of a total 5 profiles. The profiles were algorithmically constructed to maximize the interface capacity. Notice that the impaired device (CM84) gets assigned its own unique profile (Profile # 5), whereas the rest of the devices share the remaining profiles. The control profile (Profile #0), which is common to all devices, is not shown on this plot.**

Following this illustrative representation of the problem, one may come up with a suitable algorithm to construct profiles and assign devices appropriately. For example, Figure 7 shows the same group of 20 modems, assigned to 5 profiles yielding ~42% gain in capacity over a flat QAM-256 (benchmark) configuration. The chosen profiles also highlight the problem with treating the values in Table 1 as hard decision boundaries. Notice that Profile #2 in Figure 7 fluctuates between modulation values of 11 and 12 bits/symbol. The observed fluctuation is due to MER curves for the devices assigned to Profile #2, tracking the 41 dB decision boundary line. This scenario is not uncommon, and highlights that basing the construction of profiles on single point-in-time MER likely leads to unstable profiles that follow the inherent noise in the MER measurement.

The mechanics of the algorithm used to construct the profiles shown in Figure 7 will be discussed in a later section (Core Algorithm). But before diving into the details of the algorithm, we'll argue next that the problem is far more challenging than this illustration due to the following factors:

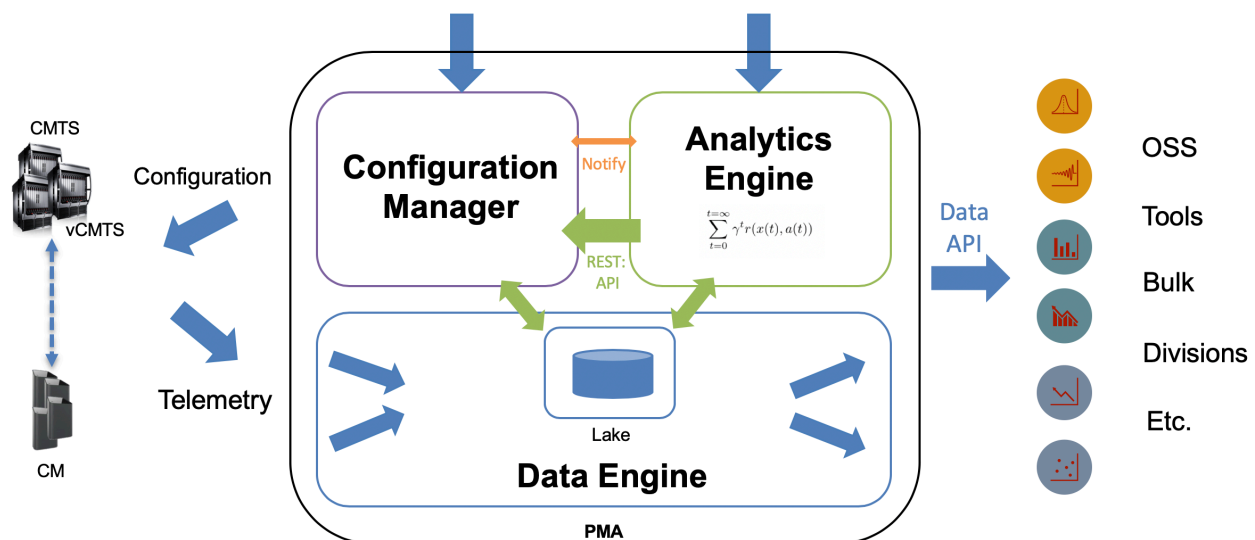
- **The time dimension:** MER measurements have inherent noise; a single-point-in-time MER curve does not capture the “true” but rather the hidden state of the S/N for a given device. The variance in MER should be addressed by the algorithm. In addition to the inherent noise in the measurement, true changes in S/N occur over time: Impairments may be introduced in the spectrum; impairments may be fixed; some may be seasonal, weather, and/or temperature related; some effects may transiently show up and disappear. These factors should be considered to ensure that the recommended profiles function properly and do not become outdated/obsolete shortly following their application.
- **Vendor constraints:** CMTS vendors introduce a host of additional constraints to the optimization problem, depending on their implementation of the D3.1 specification. In addition to the variation in supported profiles per OFDM interface across CMTS make/models, some of the encountered constraints include the following:
  - o Limit on the number of profiles per CMTS. While the optimization problem is defined at the OFDM interface level, we've encountered additional constraints on the total number of unique profiles across the CMTS.
  - o Limit on the number of segments within a profile, where a segment is defined as a contiguous block of subcarriers assigned the same modulation efficiency value.
  - o Requiring that segment width be a multiple of some fixed frequency value. Example, a segment width that is to be a multiple of 1 MHz translates into assigning the same modulation value for each group of 20 subcarriers (assuming 50 KHz subcarrier width).
  - o Requiring that the absolute frequency value at which a segment block starts be divisible by a certain value representing the CMTS “grid spacing”. Example, if the grid spacing is 250 KHz and the start frequency is 700.1 MHz, then the first 3 subcarriers are skipped or the segment starts in the excluded subcarrier guard band; resulting in the segment starting at 700.25 MHz to conform with this constraint.
  - o Allowing a subset of QAM modulation values from the list in Table 1. Example, certain CMTS models support only square QAM constellations (even number modulation efficiency values).

These are just few examples of encountered constraints. They do not apply across all CMTS makes and models and are not an exhaustive list. To deal with this challenge, the end-to-end data science pipeline built for managing profiles must be able to dynamically apply constraints depending on the specific make/model of the CMTS. At the same time, the list of known constraints should be curated and kept up-to-date to follow any CMTS hardware, firmware, and software development that may result in change to the constraints.

- **Policy considerations:** with a broadband network as large as Comcast's, managing the application of profiles goes beyond the pure data science problem. Maintaining a good degree of control over algorithmic recommendations is a must. Dictating parameters that directly influence the algorithm -- such as time and frequency aggregation statistic thresholds, applying overrides on the output, deciding how frequently profiles should be updated, monitoring performance of the program, and managing notifications & alerts -- are just few examples of the control capabilities that are a must for operationalizing PMA.
- **Forward Error Correction (FEC) consideration:** D3.1 uses a different error correction mechanism (LDPC) compared to D3.0 (Reed-Solomon). The recommended mappings between MER values and modulation efficiency values shown in Table 1 are not hard truths and will vary based on the statistics and nature of the noise and network linearity. Ultimately, the response of the system to the application of profiles is manifested in the rate of Corrected and Uncorrectable Codewords (CCER & CER). Therefore, FEC has to be measured and considered within the PMA program in some to-be-defined form.

With all this complexity in mind, we will venture next into describing an implementation of a PMA solution from a high-level “architectural” viewpoint.

## Overview of Solution



**Figure 8 - Architecture view of the PMA systems. It includes 3 components: Data Engine, Analytics Engine, and Configuration Manager. These are described in the text below.**

At a high level, the solution is comprised of the 3 systems shown in Figure 8. The focus of this paper is on the Analytics Engine (AE); but here's a brief description of the 3 systems:

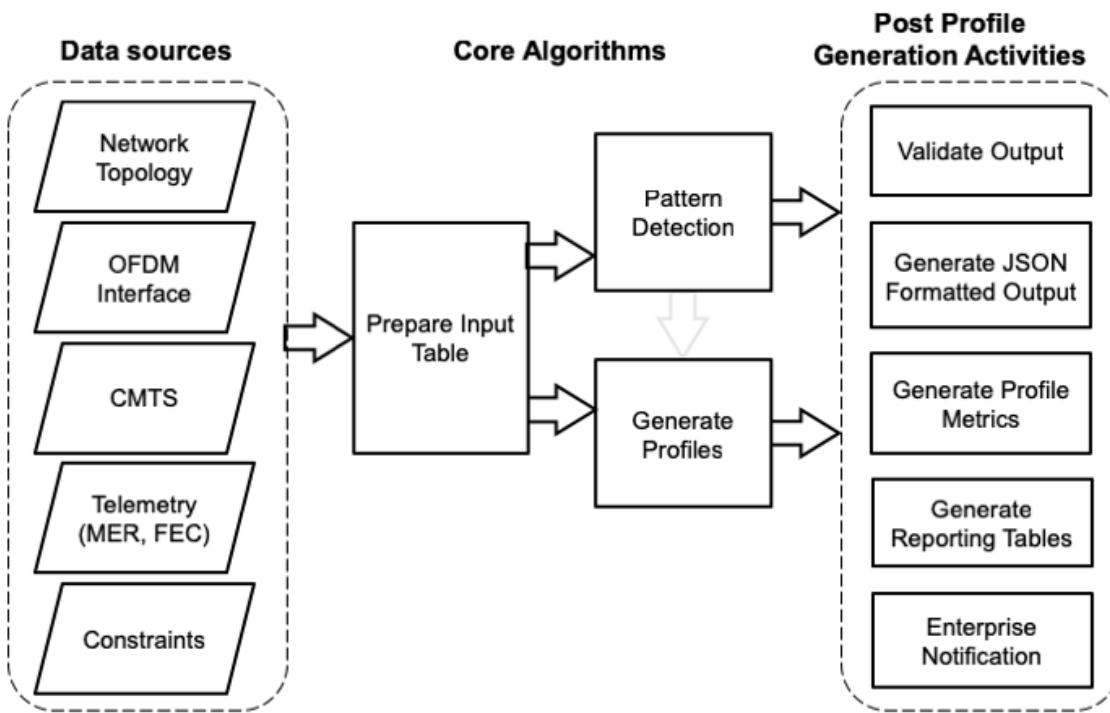
1. **Data Engine:** Collects and maintains all data required for the creation and management of modulation profiles. Data includes:
  - o **Network Topology:** connects a CM to an OFDM channel, MAC Domain, CMTS, and related Comcast Region & Division.



- **OFDM Channel Characteristics:** contains OFDM channel configuration such as channel width, subcarrier width, start frequency, active & excluded regions, position of PLC channel, and other configurations.
- **OFDM Subcarrier Configuration:** The modulation efficiency for each subcarrier and subcarrier type
- **CMTS Characteristics:** contains information on the make, model, hardware & software versions for each CMTS in the network.
- **Telemetry data:** includes MER, FEC, traffic and other network metrics for each CM and CMTS in the network.

Each of the different data sources is retrieved/updated with a certain frequency, depending on the dynamic nature of the data. For example, MER data from each device in the network may be acquired at n-hour intervals, while FEC data per profile is acquired at n-min intervals, with the topology acquired on yet a different frequency. The data architecture places the raw data in our data lake, with structure designed for ease of retrieval. The Data Engine has the responsibility of ensuring scalability of collection, performing validations to ensure integrity of the data, and managing retention policies. Note that the Data Engine is the foundation for D3.1 Proactive Network Maintenance (PNM) activities in addition to the PMA solution.

2. **Analytics Engine:** Consists of an end-to-end data science pipeline and is primarily responsible for the generation and recommendation of modulation profiles, based on the evolving conditions of the network. It retrieves the relevant raw data from the data lake, then cleans, aggregates, shapes, joins, and transforms the data. It further enables complex construction and invocation of the different algorithms for profile construction, validating the output, and generating useful metrics pertaining to the output. The Analytics Engine maintains its own data on policy and vendor related constraints. It also employs REST APIs for standard access to the recommended profiles and profile assignments as well as for managing constraints data and for making the output easily accessible by the end user. The Analytics Engine is described in more detail below.
3. **Configuration Manager:** The format of the output from the Analytics Engine is agnostic to the CMTS vendor make and model, i.e., it simply defines profiles as contiguous blocks of subcarriers, assigned specific modulation efficiencies, as well as which devices should be assigned to each of the defined profiles. The Configuration Manager is responsible for translating the raw output from the Analytics Engine to the CMTS-specific API and configuring the profiles with transactional integrity to the targeted CMTS/OFDM channel. Additionally, the Configuration Manager serves as a safeguard, conducting validation independent of the Analytics Engine, as well as checking the conditions and state of the CMTS. This is to make sure a profile update would not be disruptive to other activities, such as other configuration operations, 911 calls, CM registration status, and automatically opens and closes operational change management events. The Configuration Manager has the authority to reject or defer profiles recommended by the Analytics Engine based on policy. Finally, the Configuration Manager records the response of the CMTS to the application of profiles and can remediate or rollback, should the change cause adverse effects.



**Figure 9 - Block diagram showing the Analytics Engine key data sources, core functions, and post generation activities. Note that Pattern Detection, discussed in a later section, includes a set of algorithms aimed at detecting ingress to ensure appropriate network maintenance.**

Figure 9 shows the process flow within the Analytics Engine. The first step involves preparing an input data table that combines all needed information from the different data sources. The generated input table has observations (rows) corresponding to OFDM channels in the network and variables (columns) corresponding to features needed to generate profiles and metrics. The main features are:

- **Organizational Topology:** String representing location where CMTS is located
- **CMTS:** String representing the CMTS domain name
- **OFDM Channel:** String representing OFDM interface (includes port, card, and slot numbers)
- **Devices:** Array of strings containing list of mac addresses of devices attached to the interface
- **MER:** Array of an array of floating numbers, where each inner array contains the corresponding CM's MER vector
- **Constraints:** Data structure (key-value pairs) containing all vendor-related and global policy constraints

In addition to these features, characteristics of the OFDM channel are also included. The input data is used by two processes running asynchronously: Pattern detection and Profile generation. The latter includes the core algorithms for construction of suitable modulation profiles, given the MER characteristics of the devices and the corresponding CMTS constraints. Following the generation of profiles, several processes run in parallel:

- Profiles are validated against the constraints to ensure that the algorithm indeed generated valid profiles. For example, if a certain CMTS make/model allows a maximum of 5 segments and a generated profile

is found to contain segments exceeding this number, the violation is recorded, and the profile is flagged as invalid.

- The raw output describing profiles and device assignment is formatted in accordance to an agreed upon JSON data structure.
- Metrics for the specific profile generation run are calculated. These include the OFDM interface capacity, segmentation impact, various counts of the number of interfaces, profiles, devices, etc.
- Aggregated reporting tables are constructed for metrics of interest (e.g. the total capacity gain over QAM-256 by organizational topology). The main purpose of the reporting tables is to support operations and deployment of dashboards.

The entire pipeline presented in Figure 9 was implemented and run in the Apache Spark environment. A main orchestration function kicks off the pipeline and manages invocation of the other functions. All form of output produced by the pipeline is saved into the data lake under unique run IDs, which identify the specific pipeline run for system traceability.


## MER & Time

MER data for D3.1 devices in Comcast's network is polled at frequent intervals and used to conduct exploratory analysis to understand and quantify the amount of variability of MER, the effect of the time dimension on the stability of profiles and, most importantly, to turn these insights into actions that inform how MER data is to be processed by the Analytics Engine. Insights from the analysis are shared in this section.


First, we considered the variation in MER on a subcarrier level over a period of ~1 week during which MER is sampled at ~n hour intervals. The analysis covered a single CMTS containing 29 OFDM channels serving DOCSIS devices. The outcome of this analysis is presented in the form of the boxplots shown in Figure 10. Each boxplot corresponds to the distribution of the range (max-min) of MER variation across all combinations of devices and subcarriers in the corresponding channel. The boxplots show that the variation in MER is around 3 dB (median represented as vertical bold line in boxplots). Though there are outliers (data points in the boxplots) extending as far as >20 dB. These outliers indicate highly unstable subcarriers within the considered 1-week period. The impact of the variation of MER on modulation efficiency assignment was further investigated. Figure 11 shows that by the first day, the ratio of subcarriers that retain the modulation efficiency assignment, based on their initial MER point-in-time sample, dropped to ~0.83. After ~10 days, this value further dropped to ~0.75. The analysis suggests that fluctuations in MER values cause somewhat impactful changes early on, after the initial modulation efficiency assignment. Thus, the problem warrants some consideration of the underlying causes, as well as a strategy based on utilizing longitudinal samples rather than attempting to follow the changes in real time. Note that not operating PMA in real-time is an implementation choice we made; a choice that could be reconsidered once CMTS protocols achieve certain level of maturity that allows seamless switching between profiles. The dynamic nature of the channel and profile changes was one of the challenging stability challenges identified as part of the D3.1 hardening work.

The ~3 dB variation in MER represents a combination of inherent noise in the MER measurement (i.e. even without any material changes in quality of the spectrum, repeated measurements will show some variation) as well as slight fluctuations in the MER level. When investigating the causes for the larger variations, these were found to relate to impairments that appear intermittently, periodically, or that have an unstable character (e.g. they shift position in the spectrum or exhibit a change in their level and/or shape). An

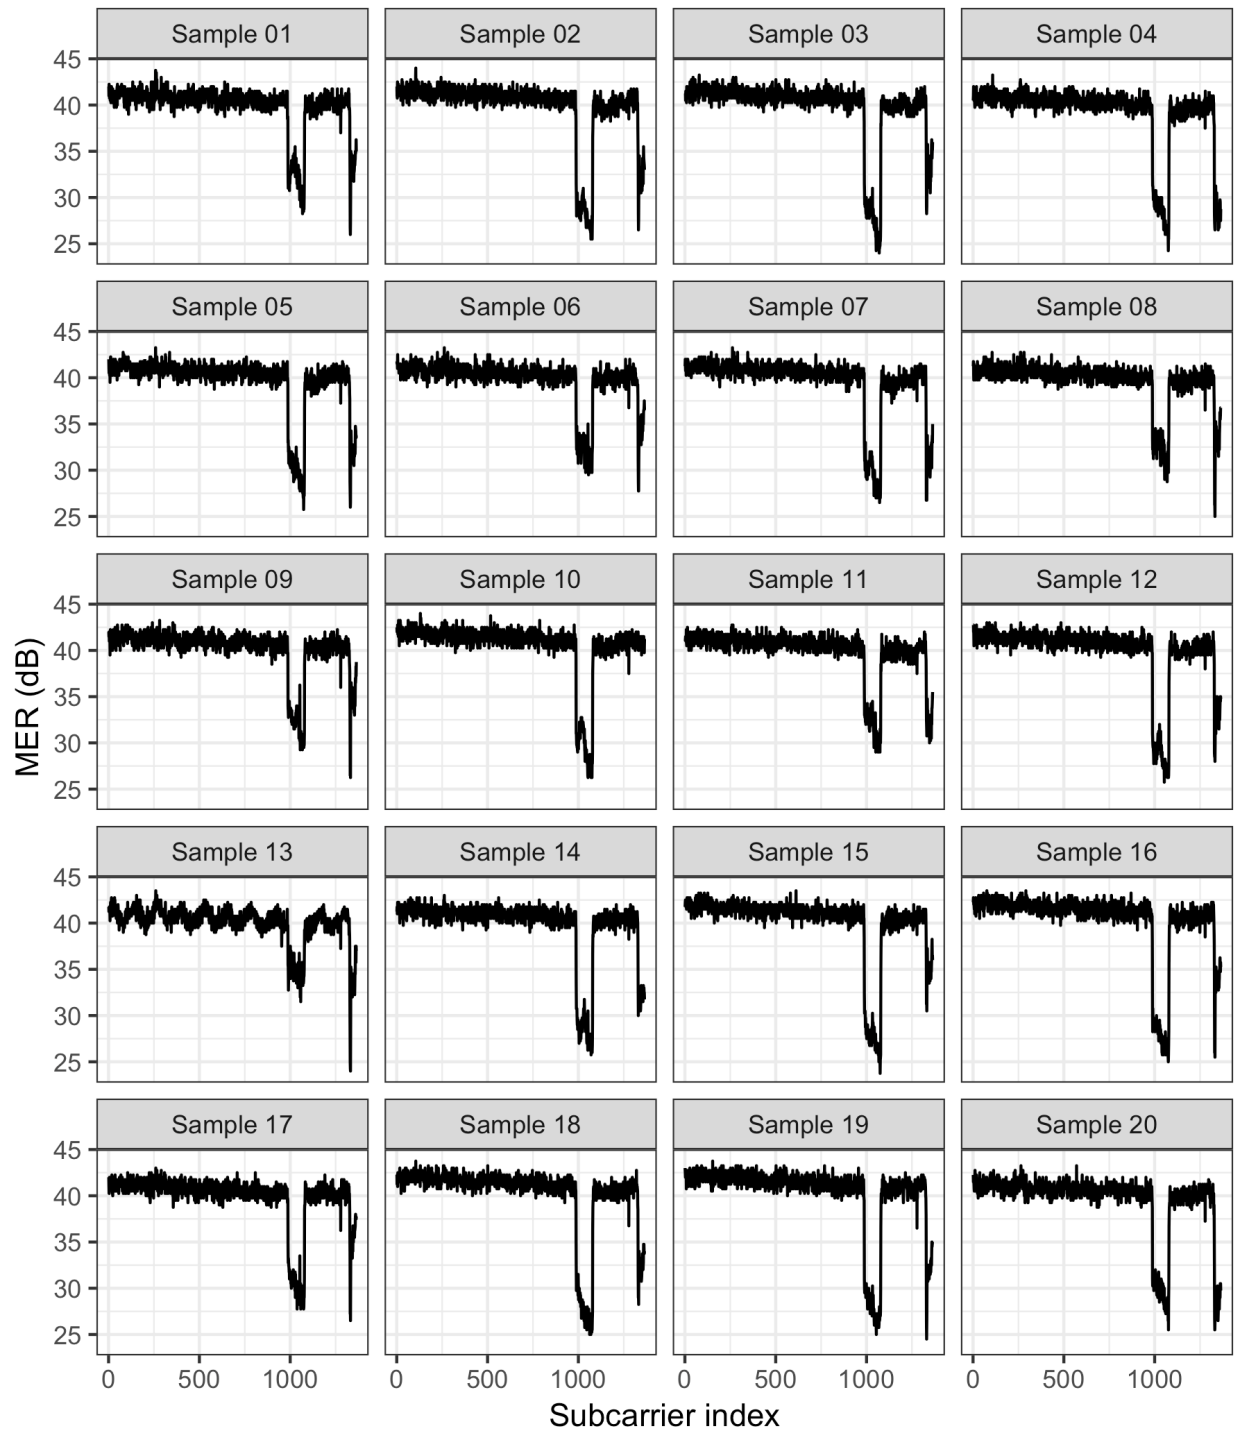
example of an impairment that causes large variation in the character of the MER curve is shown in Figure 12 for CM84, which is one of the 20 devices we've been following in this paper. Such varying MER presents a challenge to the generation and assignment of profiles. This is because each of the different MER samples presented in Figure 12, if considered separately, would suggest a different modulation profile for the device. Ideally, the fluctuation/instability in MER must be factored into the algorithm to maintain network stability, given knowledge of the MER history for devices in the network. Another example of the shortcomings of basing profile generation on a single-point-in-time MER is shown in Figure 13 For CM95. In this example, the device appears to be healthy overall. However, impairments do show up in 2 of the time samples.



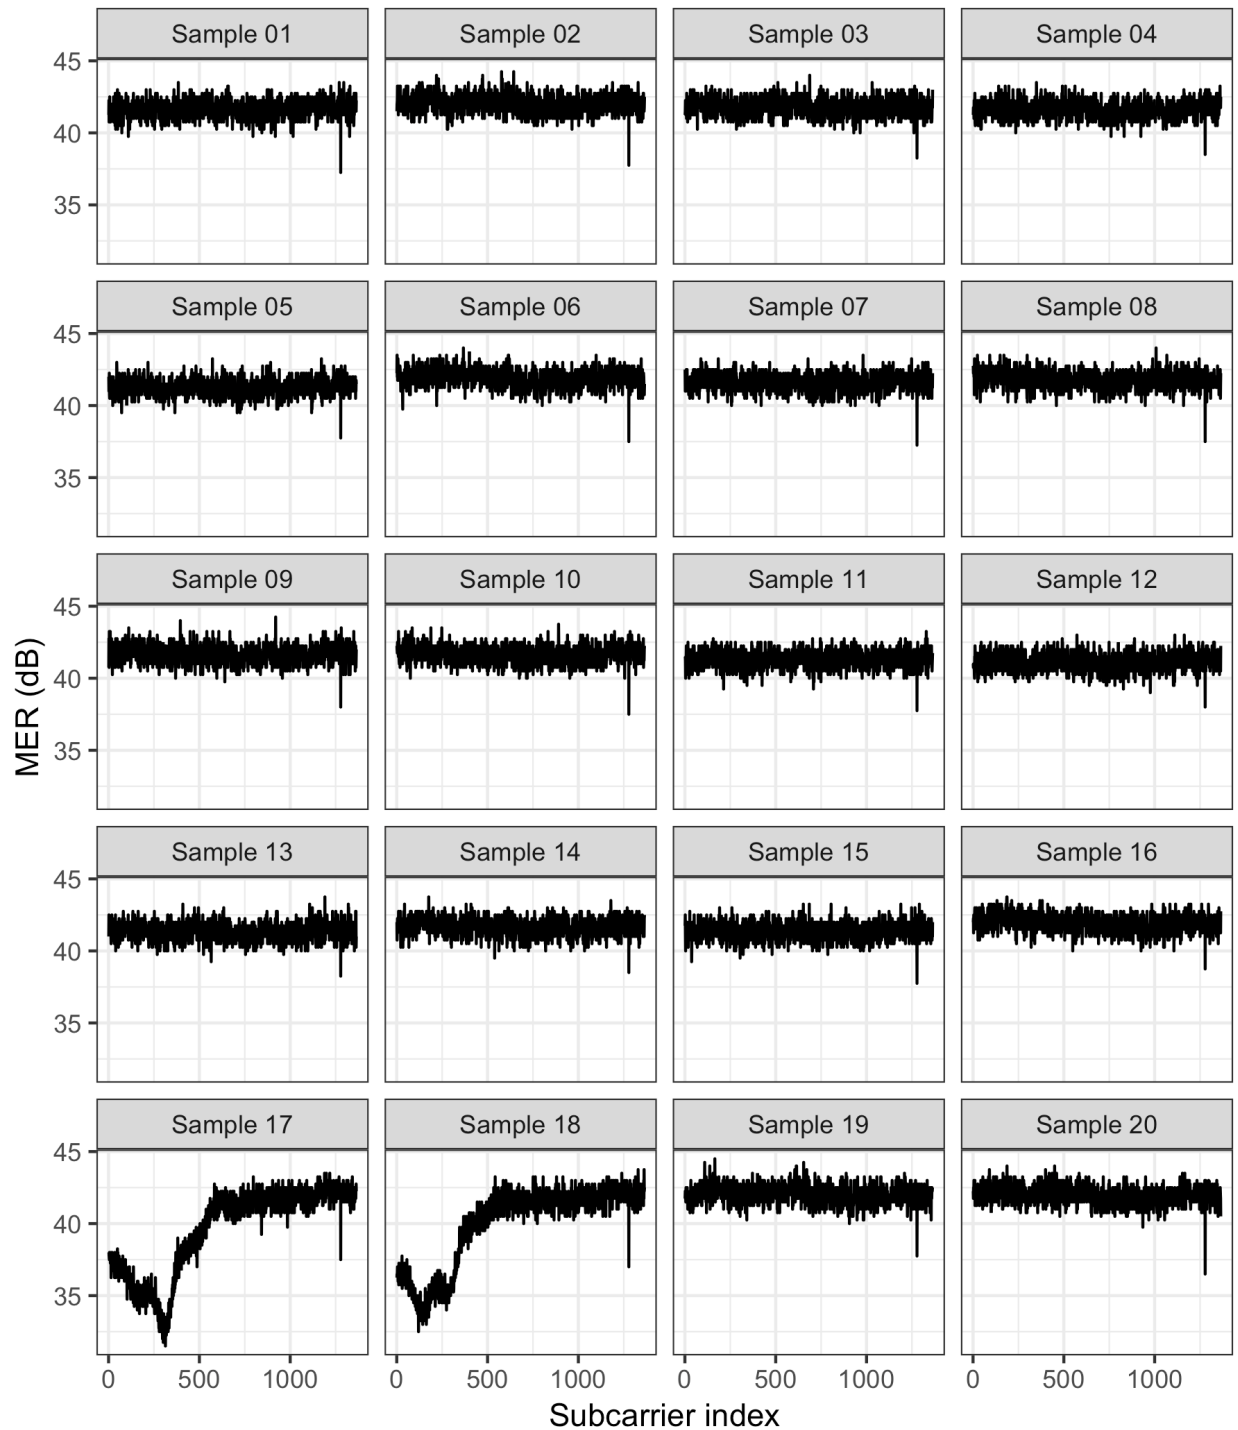
**Figure 10 - Distributions of variation in MER over time by OFDM interface for one CMTS. Each boxplot represents the distribution of MER range (max – min) for the considered ~1-week time period for a population of device-subcarrier combinations. The number of samples represents the total number of devices multiplied by the number of frequency subcarriers in the OFDM spectrum. In each boxplot, the bold vertical line corresponds to the median, the rectangular box to the inter-quartile range (IQR), and the upper/lower whiskers to  $(Q3 + 1.5IQR)/(Q1 - 1.5IQR)$ . Data points represent outliers.**



**Figure 11 - Analysis showing the impact of MER variation on the stability of modulation efficiency assignments. The ratio of subcarriers that retain their original bitload assignment as their MER varies drops in time: to ~0.83 within one day and ~0.75 within 10 days.**



**Figure 12 - MER time samples for CM84.** The sampling period was ~n hours. CM84 exhibits instability in its MER curve due to impairment, with the level of MER in the impaired region changing over time. Certain time samples also show oscillatory shape in the MER curve outside the impaired region (e.g. sample #13).



**Figure 13 - MER time samples for CM95. The sampling period was ~n hours. CM95 shows healthy MER level overall, but an intermittent impairment appears in 2 of the time samples (#17 & #18). If profile generation was based on a single point-in-time, this impairment may be missed.**

The MER instability was addressed by selecting the appropriate MER level for each device's subcarriers based on the collected time samples from history. The basic idea is to consider the measured MER curves for a device over a period of several days. With a  $\sim n$  hour sampling interval, this should result in  $\sim$ tens of MER samples collected per device. Thus, for each combination of device and subcarrier frequency there will be a distribution of MER values, from which the  $n$ th percentile is selected to represent the MER level for the device-subcarrier. The adopted method allows 3 tuning parameters that are specified through the pipeline global policy:


- The historical period considered for collecting MER time samples (e.g. last 10 days)
- The maximum number of samples to be considered per device (e.g. 100 samples)
- The value of  $n$ th percentile used for MER selection from the distribution (e.g. 10th percentile)

Controlling these parameters allows the selection of MER curves to vary on a spectrum ranging from the very conservative to the very aggressive. The conservative end corresponds to using small values for the  $n$ th percentile. At one extreme, picking the 0th percentile corresponds to the minimum value for selecting MER from each subcarrier distribution. With 0th percentile, we're effectively preparing for the worst observed state, given historical data. On the other end, large values for the  $n$ th percentile have the opposite effect, with the 100th percentile corresponding to the maximum value for selecting MER from each subcarrier distribution. Increasing the considered period also has an effect of adding some conservatism to the approach. This is because the likelihood of capturing irregularity with the device increases with increasing number of considered samples. The approach also allows the use of a single-point-in-time MER by setting the maximum number of samples to 1 (in this case, the  $n$ th percentile value becomes irrelevant). The maximum number of samples is currently set to a high enough value so not to influence the selection. Note that the default values are viewed as initial trial policy. For instance, the 10th percentile, while conservative enough, still allows for elimination of outliers (extremely low MER levels encountered). The 10-day period is viewed as an onboarding period, during which new devices and devices with recently fixed impairments are expected to demonstrate stable behavior before their "bad" history is erased/refreshed. The Analytics Engine pipeline allows experimenting with changing these tuning parameters and measuring the effect on the key performance metrics. These results are presented in later section, along with additional sensitivity analyses.

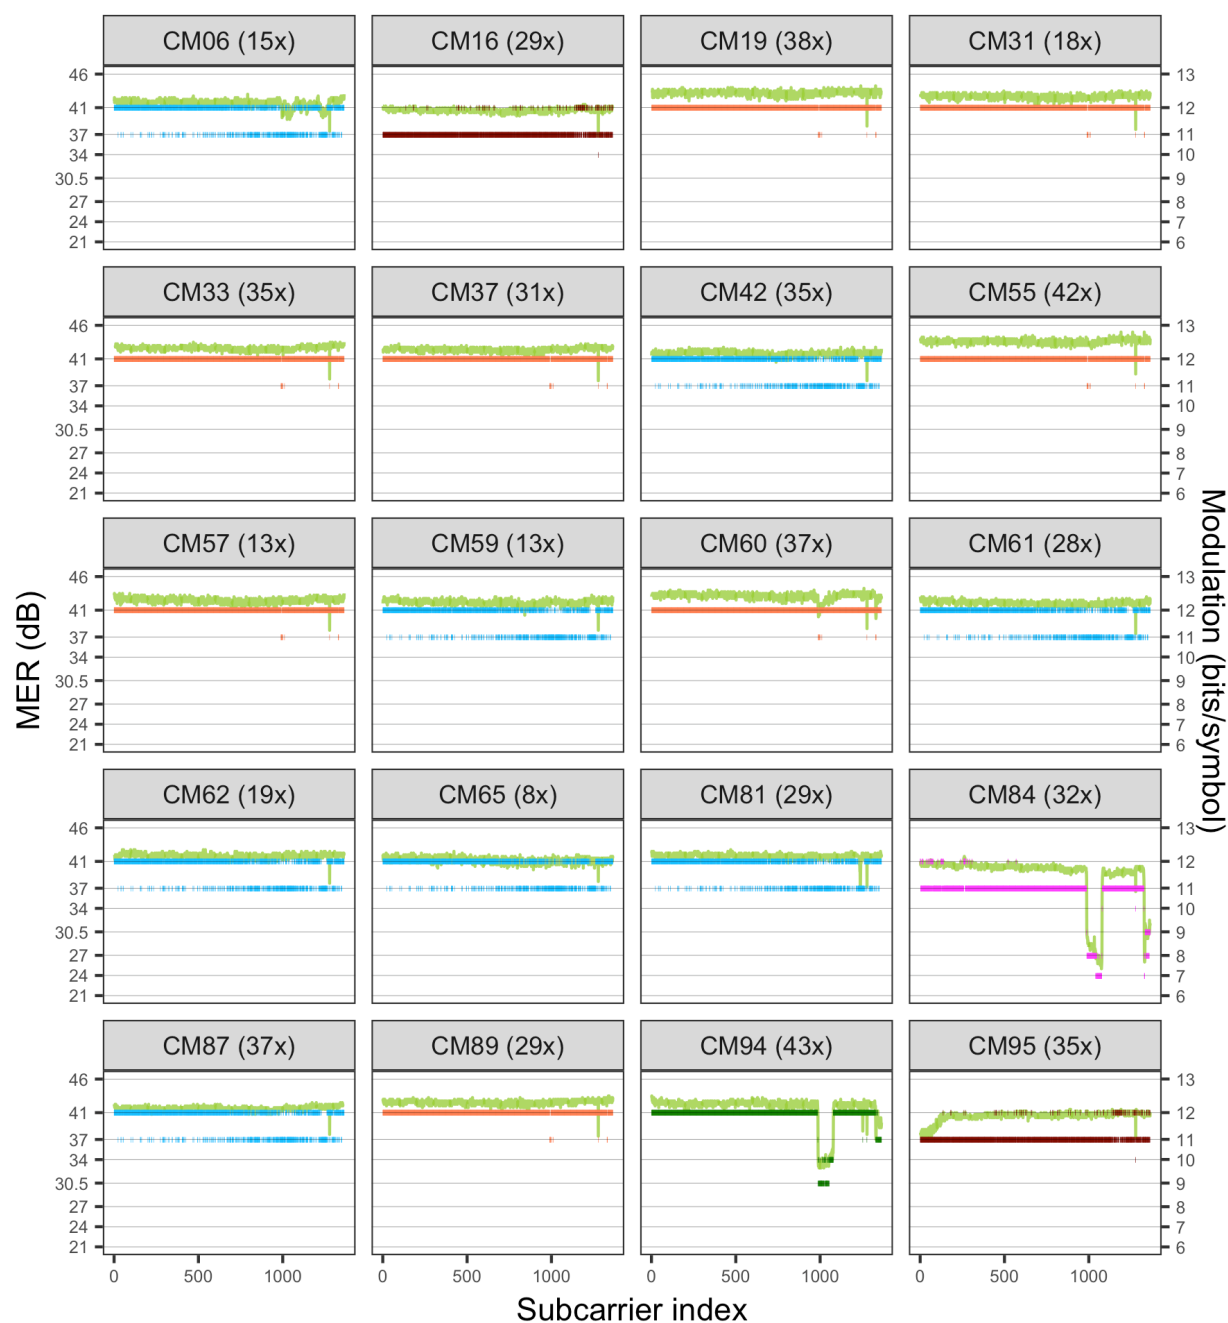
To visualize the processing of historical MER data, let's once again consider the same 20 devices introduced earlier. Figure 14 shows MER data collected for 10 days (thick green band) for the devices, along with the 10th percentile selection curve (solid gray line). Notice that an immediate benefit to using this approach is the reduced variance in the aggregated MER curve. This effect is expected, since the  $n$ th percentile selection has a similar effect to averaging, in terms of increasing the S/N of the measurement. The effect of increased



MER is also manifested in the less fluctuating levels in the profiles generated based on the aggregated MER values (Figure 15).



**Figure 14 - MER time samples for the same 20 devices collected over a period of 10 days.**  
The thick green band represents all MER time samples. The number of time samples varies by device and is indicated in each panel's header. The solid gray line represents an MER curve constructed by selecting the 10th percentile value for the distribution of time samples at each frequency subcarrier. Hence, the aggregate MER curve tracks the lower end of the green band.



— MER (dB) — Profile 1 — Profile 2 — Profile 3 — Profile 4 — Profile 5

**Figure 15 - The same 20 devices assigned profiles that are based on the aggregated MER values rather than a single point-in-time MER. These profiles are expected to be more stable because they consider the variability in MER measured over a period of ~days.**

# Core Algorithm

We adopted a *greedy* approach to the generation of profiles that optimizes OFDM channel capacity and satisfies given vendor and policy constraints. The main steps involved are the following:

1. Clustering of modems within an OFDM interface based on some objective function (e.g. the similarity of their MER curves)
2. Construction of profiles for the clustered modems within the interface
3. Segmentation (reshaping) of the profiles to satisfy the segment-related constraints
4. Pruning of the number of profiles to satisfy the maximum profiles per CMTS constraint

The ordering above represents the natural way to attack the problem: we start with the core profile generation, then reduce the number of segments, then (optionally) reduce the global number of profiles. The approach is greedy because it does not guarantee that the outcome yields the absolute optimum solution. Instead, the approach presented includes a combination of several steps, involving various computational techniques, ranging from machine learning algorithms to rule-based approaches. Choices made are in large part informed by the prior exploratory data analyses, and lab experiments. It is expected that the algorithm will continue to evolve and develop through rollout of future versions. Next is a description of each of those steps.

## 6. Clustering of Modems

The core profile generation algorithm uses hierarchical clustering [3] to group modems within an OFDM interface that share the same MER characteristics into the requested number of clusters. Deciding on which modems to group together proceeds iteratively: in each iteration, the two modems/clusters that have the most similar MER curves are merged together. Similarity is measured using several alternative objective functions. These include the Euclidian distance between the MER curves of two clusters and the Capacity Loss (i.e. the amount by which capacity of the OFDM interface would be reduced following merger of the two clusters). Once a merger occurs, the MER curve representative of the newly formed cluster is calculated on a subcarrier basis by taking the minimum of the MER levels of the two merged clusters along that subcarrier. That is, if cluster  $p$  was merged with cluster  $q$  to form cluster  $r$ , then the MER vector of the newly formed cluster is:

$$CM^{(r)} = \left[ \min(x_1^{(p)}, x_1^{(q)}) \quad \min(x_2^{(p)}, x_2^{(q)}) \quad \cdots \quad \min(x_i^{(p)}, x_i^{(q)}) \quad \cdots \quad \min(x_L^{(p)}, x_L^{(q)}) \right]$$

The requested number of clusters corresponds to the number of profiles per OFDM channel allowed by the CMTS vendor minus one (as one profile is reserved for Profile 0). Once the requested number of clusters is reached, the iterative process terminates.

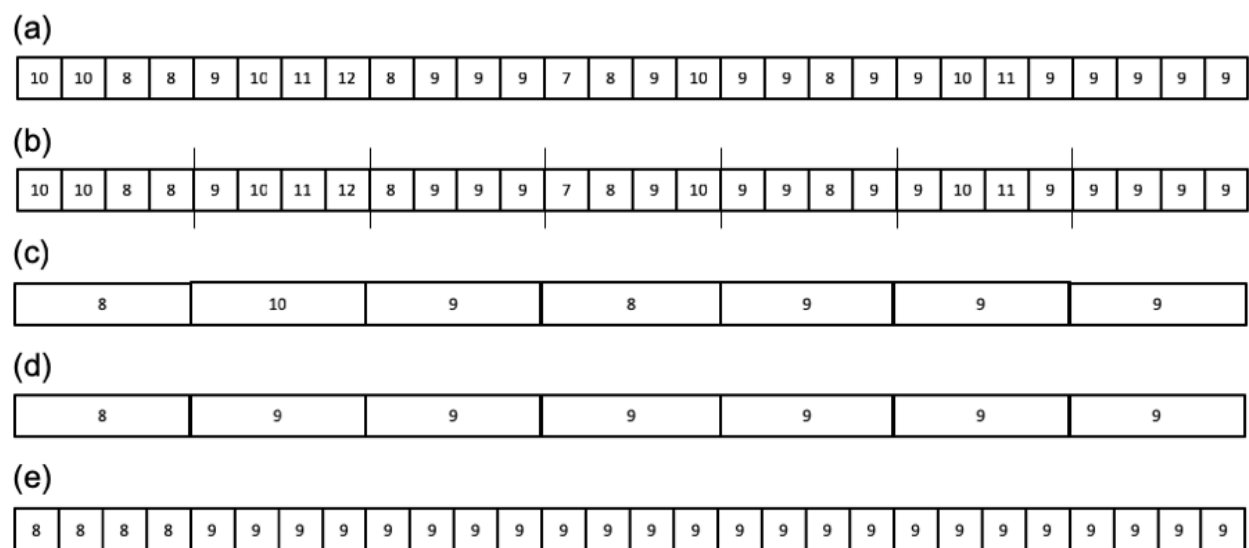
## 7. Modulation Efficiency Assignment

The MER curve for each cluster is mapped into modulation efficiency values based on the mapping in Table 1. No consideration for segment-related constraints is made at this point. Profiles are thus generated, but they are most likely “invalid” from the perspective of the CMTS.

## 8. Segmentation


The two nontrivial segment-related constraints are the segment width and the maximum number of segments. An illustrative example for the segmentation process is shown in Figure 16. The first constraint

is addressed by binning the modulation efficiency vector such that every  $n$  subcarriers get assigned the same modulation efficiency value—the median of the  $n$  modulation efficiencies, with  $n$  being a parameter that is determined by the vendor constraints data (steps a-c). Next, the number of segments is reduced in accordance with the vendor-related constraint on maximum number of segments. This is done by running an aggregation (smoothing) function on the modulation efficiency vector within an initial window size of 3. Each value at the center of the window is replaced by the median value for the window. The window size is increased iteratively until the segments are reduced to the set limit (step d). Once the required number of segments is reached, the modulation efficiency values are expanded by replicating each value  $n$  times (step e).

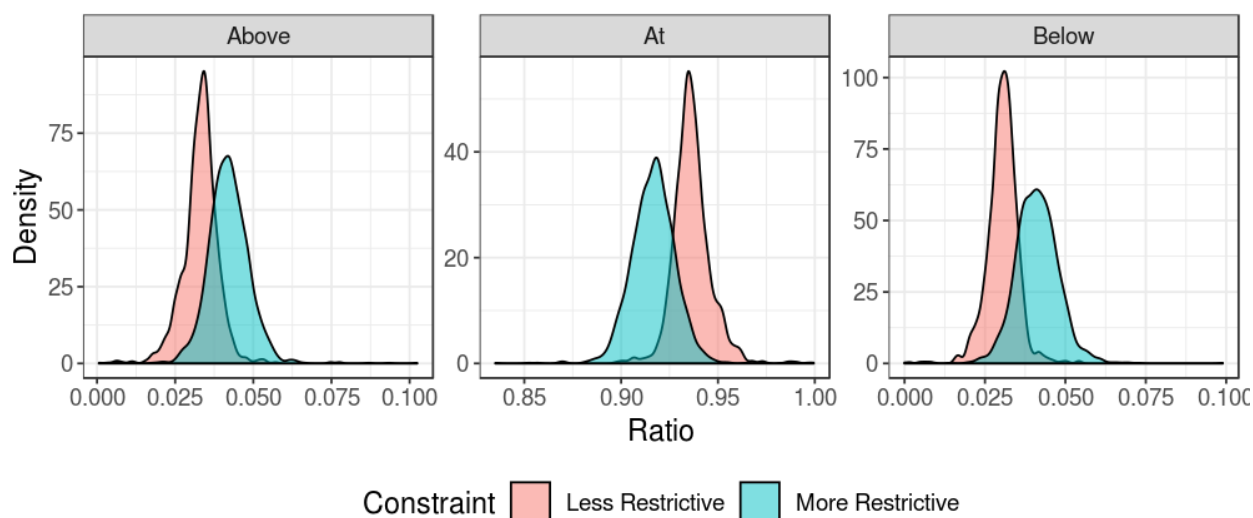


**Figure 16 - Illustration of the segmentation process. (a) The starting profile has 18 segments. The (fictitious) constraints require a maximum of 3 segments and a segment width that is a multiple of 4 subcarriers. (b) Subcarriers are divided into blocks of 4 subcarriers. (c) Each block is assigned the median value. When there's a tie, the lower modulation efficiency value is used. (d) A moving median window of width 3 is applied, resulting in a reduction of the number of segments to 2. (e) The binned modulation efficiencies are exploded to retain the original length of the profile.**

Figure 17 shows the 5 profiles generated for the sample 20 devices, both before and after segmentation. Note that the segmentation procedure outlined here may result in devices operating above/below the recommended modulation efficiency value, according to the mapping in Table 1 for parts of the spectrum. We permitted this violation of the recommended mapping to occur for a couple of reasons. First, the choice of 10th percentile and large MER consideration period already includes conservatism in choice of MER values from time samples. Second, we believe that this effect is minimal because MER levels between neighboring subcarriers are correlated (e.g. impairments tend to affect a chunk of contiguous subcarriers, not randomly positioned subcarriers). We tested this proposition by tracking the ratio of profile subcarriers operating below, at, and above their recommended level. The result in Figure 18 shows that ~90% of subcarriers across the whole network are assigned their appropriate modulation level after segmentation.



**Figure 17 - Example illustrating segmenting profiles to satisfy vendor constraints. Here, the profiles for the 20 example devices were segmented such that the maximum number of segments is 5 and each segment width is a multiple of 20 subcarriers.**




**Figure 18 - Impact of segmentation on reassignment of bitloads as measured across Comcast's full D3.1 network. The analysis shows the impact of segmentation for two hypothetical CMTS vendors: one with more restrictive constraint and one with less restrictive constraint. Overall 90% of the subcarriers keep the modulation efficiency value assigned according to the mapping in Table 1. Less than 5% operate at level higher (lower) than the level recommended for their MER.**

We also confirmed that the segmentation procedure does not have a negative or positive impact on interface capacity. This is because the aggregation algorithm described above is, in effect, a smoothing operation: On balance, the total sum of bitloads across the spectrum does not change when comparing the values before and after segmentation.

## 9. Profile Consolidation (Pruning)

Profiles are consolidated to satisfy the global limit on the number of profiles within the CMTS. This step is invoked only if the CMTS vendor limits the total number of profiles within the CMTS. It uses the exact same algorithm for clustering of modems. The end result is the pruning of profiles to the limited number, such that certain profiles are reused across multiple OFDM channels.

We examined the effect of pruning on the capacity gain for one CMTS. The results are shown in Figure 19. The study shows that pruning slightly reduces capacity gain. More interestingly, this analysis reveals the greedy nature of the algorithm: it does not pay to over-optimize in step 1 (clustering), given that profiles will have to be pruned at a later stage. For the CMTS used in this study, working with 8 to 10 profiles seems to be ideal without other limitations, although most CMTSs are limited to fewer than 6 profiles in the current implementation.



**Figure 19 - Exploring capacity gain as the local (interface-level) and global (CMTS-level) constraints on the number of profiles are varied. The panel on the left shows capacity gain (left y-axis) and number of profiles (right y-axis) for the scenario where the constraint on global profiles is 248. The panel on the right shows the same metrics for the scenario where the global profiles are not constrained.**


The algorithm as presented contains several tuning parameters that allow further optimization of the output and can be set through global policy. These are listed in Table 2.

**Table 2 - List of Algorithm tuning parameters**

Step	Hyperparameter
Processing MER Time Samples	Collection period; default value: n days
Processing MER Time Samples	Percentile selection value; default value: n%
Processing MER Time Samples	Maximum number of samples default value: n samples
Clustering	MER similarity function; default value: Capacity Loss
Clustering	Option to threshold the device MER before clustering; default value: Yes
Modulation efficiency	Global offset applied to threshold in MER mapping table; default value: n dB
Segmentation	Initial size of smoothing window; n

The effect of changing the hyper-parameters on the capacity was explored through various sensitivity analyses. One analysis is shown in Figure 20, where the MER selection percentile and the global decision

boundary offset was changed. The analysis was conducted on the entire D3.1 footprint. Qualitatively, the observed trend in capacity gain is expected. A 3 dB offset (increased/decreased thresholds) can change the capacity gain by about 10% (decreased/increased capacity). However, note that the main tradeoff at play is the decrease/increase in uncorrectable error rates. We are currently working on measuring the actual response of the system as manifested in the FEC data.



**Figure 20 - Full footprint sensitivity analysis exploring change in capacity gain as the MER consideration period and MER selection percentile are varied. A 2-week consideration period (left) is mostly similar to a 1-week consideration period (right). The main difference is the extended y-axis range of the curves for the 2-week period as more extreme samples are encountered. In both panels, the global offset has the effect of shifting the curves up (reduced thresholds) or down (increased thresholds).**

Based on the analysis above, we believe that a conservative estimate for implementing PMA will be >30% increase in capacity across Comcast's network. Future opportunities to increase by an incremental ~10% exist. The estimate corresponds to 2-week MER consideration period, 10<sup>th</sup> percentile selection value, and 0 dB offset (i.e. keeping the mapping in Table 1 as is).

## Lab Environment

We recognize the need for lab environments that can support CMTS feature testing and qualification, pre-deployment testing of the system, algorithm development, and automated testing for our planned Continuous Integration & Continuous Deployment (CI/CD) pipeline. To ensure good test coverage, our labs are configured to represent the most common combinations of CMTS and cable modem populations present in our deployed systems, rather than exhaustive testing of all possible combinations of CMTS and cable modems in use. From the perspective of feature and hardware support required by profile management, CMTS variants are further defined by particular downstream line card features and CMTS firmware versions.

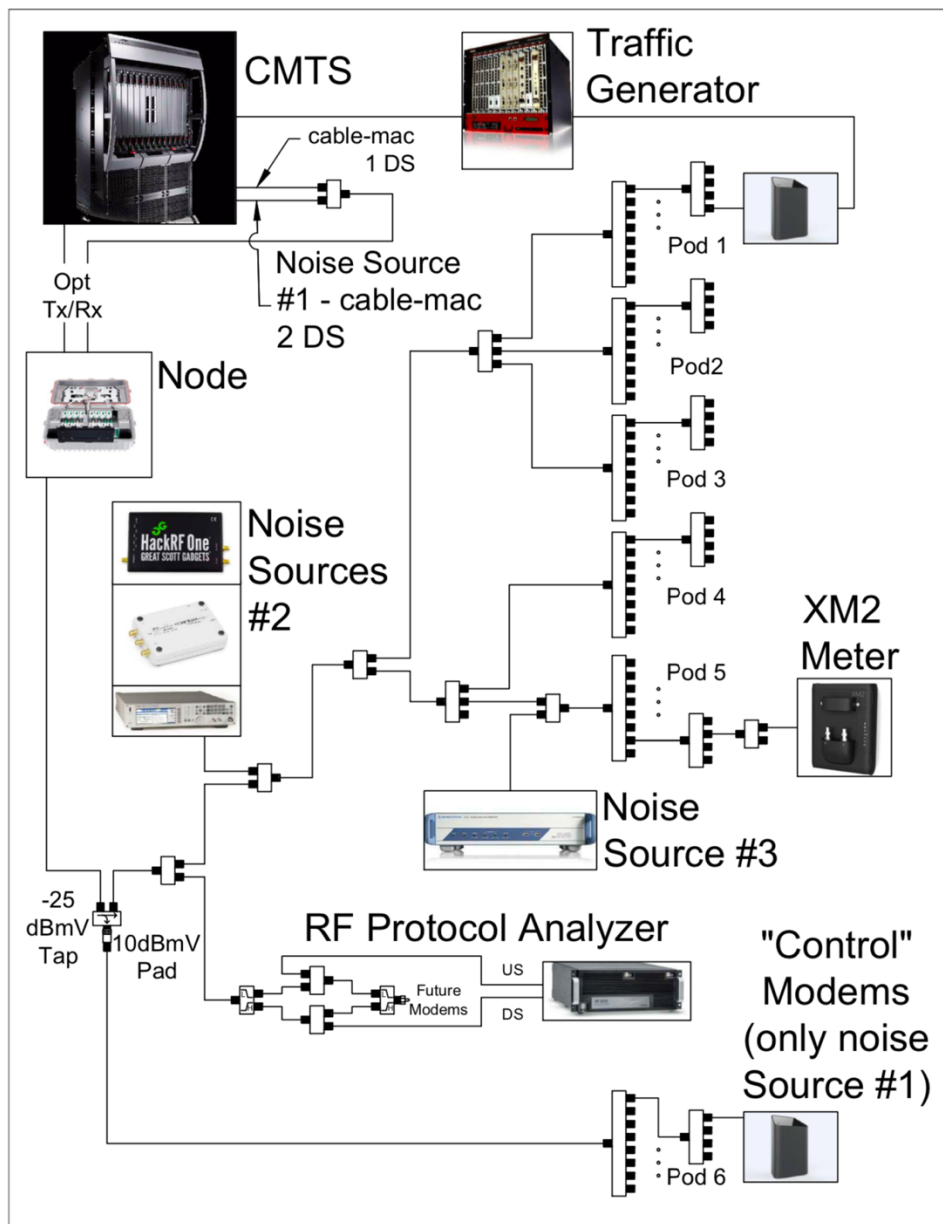
At Comcast, we support two primary integrated CMTS solutions and our vCMTS (virtual-CMTS) solution. Our D3.1 modem population is a mix of Comcast provided gateways and customer-owned retail modems. Two primary cable modem chipsets are in use in the modem population. Here again, the firmware version further defines and constrains support for profile management type applications. Our Labs provide large groups of test modems, selected to achieve good test coverage of the numerous combinations of chipset, firmware, and retail brands.



Early in the development process, it was apparent that support for profile management related features in all of our CMTS platforms was quite different, and highly dependent on the actual hardware and software configuration of the devices. One of the first use cases for lab activity was focused on documenting profile management related hardware and software support for each CMTS variant, testing how the features worked, and verifying that the two cable modem SoCs were properly supporting the profile management features. For example, early in our D3.1 deployments, we experienced modem instability due to profile flapping, described earlier, which was traced, based on careful lab testing and analysis, to hold-off timers implemented differently in each CMTS and the CMs.

Additionally, we established two primary test labs, one of which is used for development related activities, and the other for pre-deployment verification of proper operation of our solution. To fully test the solution, each lab provides groups of modems with different RF channel impairments, in order to verify proper operation of the PMA solution. As of this writing, the development lab is configured with three groups of modems, two of which experience different impaired downstream channels, and the other being a control group with an unimpaired channel. Each of the labs has all CMTSs, test populations of modems, RF signal sources used to produce impairments, traffic generation, and the necessary RF plumbing that would ordinarily be present in the HFC RF network.

To test any DOCSIS profile management solution, some means of generating realistic HFC channel impairments must be provided. From the data we collected, it was clear that the most common impairments our solution would encounter involved some sort of signal interference. As such, our first test beds focused primarily on duplicating interference sources and presenting them to the test network in a repeatable, controllable way. It was also clear that given the numerous types of interference scenarios we needed to test against, some basic automation of CMTS and test equipment would be required to efficiently configure the channel environment, run tests, and produce reports. Our automation solution is an ongoing work-in-progress, but the initial automation solution was based on simple Python scripts that manipulate the CMTS via the CLI interface. Automation of commercial test gear follows a similar design pattern: Python scripts with appropriate interface libraries to access the test gear (for example, using the VXI-11 standard). Software Defined Radio (SDR) solutions, which appear particularly promising, have been developed to be highly configurable through software control.



**Figure 21 - Development lab RF connectivity showing our current development lab signal flow. Traffic flows are generated by an Ixia TG traffic generator; the three test populations of modems are shown on the right side of the diagram; the CMTS and optical node are shown on the left side of the diagram. Noise generation sources are connected as needed during the test processes, and an Avera DP-1000 is included for DOCSIS packet capture. The combining networks are designed in a way that isolates the impairments sources, to ensure the two impairment groups experience unique RF channel environments.**

In the development lab (Figure 21), impairments are generated by commercial RF signal generators, unused RF output ports of the CMTS, and SDR devices. Unused CMTS RF ports were a first and convenient way to introduce easily-controlled impairments into the development system RF networks. One desirable feature

of this method was to be able to use the same CMTS software drivers we developed for profile management to also turn on, turn off, and adjust levels of the simulated impairment source. Our lab signal generators are used for a multitude of projects, and tend to move between project teams, so using the CMTS RF ports for impairments virtually guaranteed access to a signal source when a particular test flow needed to be run, while use of the signal generators required scheduling of their use. The downside of using the CMTS RF ports for impairment signal generation is that the possible impairment waveforms is constrained to what the CMTS can produce -- either SC-QAM carriers of 6 MHz, or an OFDM carrier of varying width. For LTE-like pulse interference or noise, it was still necessary to use a commercial RF signal generator. For automation purposes, commercial RF signal generators vary in their support for remote control applications. Older generators simply have no means of programmatic control; others use older GPIB interfaces; and the more modern versions often rely on vendor-provided libraries, or require sometimes difficult-to-use frameworks or libraries. Usually, the add-on libraries require expensive software licenses to enable remote control of the device. One final concern with commercial RF generators is sometimes limited waveform memory, which in turn limits the time duration and/or bandwidth of the impairment signal that can be produced by the device.


One alternative to RF impairment generation that was developed is based on commercially available SDR technology. As the name implies, SDR generates waveforms using software, and streams the digital representation of the waveforms over USB or Ethernet to the SDR device. The SDR device converts the digital waveform back into an RF signal. The advantage of this technology is the intrinsically software-based approach to signal generation, which, in turn, meshes very well with our objective of full automation of test flow, and support of our CI/CD pipeline as it matures. In the sections below, we will discuss one novel method we are developing that allows us to reproduce, using SDR, actual impairment scenarios extracted from modem MER data.

Early on, as our data collection systems became operational and began accumulating MER data from our deployed D3.1 modems, it was clear that our modems were operating in channels impaired by locally generated (in-home) and plant-side LTE interference, misconfigured sweep generators, HFC suck-out, roll-off, old pilot carriers mistakenly left on, and a variety of other unusual scenarios, as described in the D3.1 Hardening section. Our profile assignment algorithm is designed to reduce modulation around such impairments, but one question arose: with such a diverse set of channel impairments present in our data, how do we exhaustively test and verify the PMA system will react properly in each of these impairment scenarios?


To generate such a diverse set of RF impairments, we have started down the path of using SDR techniques to generate a library of waveforms used to test our PMA. Initially, we have settled on a relatively low-cost unit named “HackRF”, and a slightly more expensive device from Ettus Research, the USRP B210. Both devices are capable of generating RF signals up to 6 GHz, but support different sampling rates in their analog-digital conversion (HackRF up to 20 mega-samples/sec, B210 up to 56 mega-samples/sec). The useful width of the generated impairments is constrained by the sample rate of the SDR and the data rate the associated Linux system is able to generate. In our current system, the B210 is operating at 38.4 mega-samples/sec and provides roughly 34 MHz of spectrum that is used to generate impairments in the downstream channel. The SDR hardware is driven from software around the popular open-source GNURadio software system. GNURadio offers an easy to work with Python/C++ programming environment and also a tool named GNURadioCompanion that allows construction of simple SDR applications built up from a library of common Digital Signal Processing blocks.

Figures 22 & 23 show a promising approach to developing realistic impairments, using MER data collected from cable modems operating in an impaired channel. The basic concept is to select a subset of subcarriers experiencing the impaired channel, performing an inverse Fast Fourier Transform (iFFT) on them, then


streaming an impairment out of the SDR. The RF signal from the SDR is then combined with the pristine downstream signal from the CMTS, thereby producing the impairment observed on the original cable channel.



**Figure 22 - MER data received from a D3.1 modem operating on our HFC plant. This plot shows two significant interference regions, likely from cellular LTE transmission in the 700 MHz band. The shared red region is the segment of spectrum the SDR impairment generator will duplicate and inject into the lab test network.**



**Figure 23 - Spectrum analyzer capture of the SDR generated spectrum based on MER data shown in Figure 22. The tone artifact at about 740 MHz is due to I-Q imbalance in the SDR and isn't considered an issue for this particular test. The SDR-generated signal would be added to the original, clean lab downstream OFDM carrier to produce the original MER response measured in the field.**



**Figure 24 - Example of an impairment generated via SC-QAM carriers from an otherwise unused port on the CMTS. In this plot, the control group (unimpaired) modems are shown against modems operating in an impaired channel. The colored lines depict how the PMA solution is bit-loading around the impaired regions. Actual impairments observed on the operating HFC plant are not as cleanly defined in either time or frequency as this lab generated impairment is, but for functional testing, SC-QAM-generated impairments were an easy, first test case in our lab work. Time-based behavior of the impairment is simulated by turning on/off the SC-QAM carriers from automated scripts.**

Another complication we encountered was attributable to the nature of how the profile management solution was architected. At Comcast, our D3.1 system is composed of over three thousand CMTS units, outputting more than one hundred thousand OFDM channels, received by millions of D3.1 cable modems. From the outset, we know our profile management solution needed to deal with a very large, and quickly growing cable modem delivery system. The system architecture embraced cloud-based techniques from the beginning, with the core concept that cloud-based architecture, with its inherent elasticity, would deliver the necessary performance, as our D3.1 delivery system grew. Our solution is structured around a closed loop concept where cable modem telemetry is continuously collected at scale, stored in cloud-based technology, and analyzed continuously to design appropriate profiles for each OFDM channel in our delivery system. However, for the purposes of testing, the architecture presented some challenges to our test workflow. We couldn't just configure the test bed, and invoke the profile management solution on demand, in a controlled way. We needed some way of scheduling test workflows around the normal operation of the pipeline.

The profile management solution can be thought of as a conveyor belt, with measurements from the cable modems entering on the left, data analyzed as the belt moves left to right, profiles emitted, and finally CMTS channels configured, once the profiles are available and the system determines that the latest profile materially improves the observed channel for the modems. The virtual conveyor belt turns continuously, asynchronous to test work flow. Additionally, the profiles are based (as discussed earlier in this paper) on data collected over a week or more of calendar time. So as we considered how to structure our test workflow, it was apparent that if we wanted to conduct the tests on the unaltered profile management pipeline, we would somehow have to compress a week or more of data collection into a much shorter period

of time, such that multiple test cases could be managed each day. One concept we developed is called “spot-beaming”. In the actual solution, the system collects MER data from the modems four times a day, and SNMP-based measurements are gathered every twenty minutes. With the spot-beaming solution, the collection infrastructure is configured for the test CMTS to sample at an increased rate of 1-2 minutes per data point, essentially compressing a calendar week worth of measurements into several hours. With the spot-beaming telemetry available, it was then possible to run tests using simple mechanisms such as Linux “cron” to manage individual test cases. From the perspective of the profile management system, the test workflow is simply modifying the channel environment, as would be encountered in the actual deployment of the system. Figure 24 shows example MER charts and profile assignment recommended by the Analytics Engine for a population of devices with lab-generated impairments. With the spot-beaming capability, lab efforts can cycle through many impairment scenarios on a daily basis to validate the implementation of PMA.

## Pattern Detection

A PMA solution has proven to be very effective at managing both the network stability and network capacity. Comcast has also developed methods to reduce the impact on customer experience of network impairments very effectively [4]. The result can be temporary or less effective than desired, if the operational mitigation of network issues is not also considered. In turn, the full benefit and savings on operations metrics, like truck rolls and the customer experience, were not fully realized until the tools and metrics were used, to dispatch technician fix agents to implement network changes that resolved the problems. Networks tend to degrade over time, from an impairment perspective, if not maintained. When both network remediation and automated mitigation were applied, the customer experience improved, operations metrics reduced, network performance improved. The network maintenance was targeted and efficiently scheduled, and repeat truck rolls were avoided, providing significant value.

This PMA solution is following with the same dynamics as the aforementioned case study [4]. To learn from past lessons, pattern detection algorithms and other analytics are provided as a core part of the architecture. Events 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 such as those identified within this paper. In addition to calculating the profiles to optimize customer experience, a pattern detection pipeline is required to ensure that care/fix agents are efficiently dispatched to remediate network challenges. As such, adaptive systems as PMA can easily mask the degradation of the network. The types of pattern detection already developed and under development include:

- Mobile wireless interference enriched with mobile carrier information
- PLC-related ingress in the spectrum, enriched with mobile carrier information
- Non-mobile wireless ingress, such as:
  - o Incorrectly-configured sweep insertion points in the OFDM channel
  - o Incorrect placement of sweep QAMs incorrectly within the OFDM spectrum
  - o Other wireless ingress
- Analog Television non-linear distortion interference falling into the OFDM channel
- Linear impairments due to micro-reflections
- Profile recommendations with X% of subcarriers below Y bps/Hz modulation efficiency
- Recommendation for new PLC locations

Many of these impairments can be diagnosed with high resolution MER per subcarrier and spectrum analysis information, available through D3.1 PNM. The RxMER, for example, can be used to identify ingress that is at a lower power than the signal, and as such not visible on typical spectrum analysis tools. For example, the PLC location relative to mobile spectrum is now leveraged, along with MER, to provide


a preferred PLC location, should new spectrum be licensed and deployed, or other ingress challenges emerge. With each of these detection events, the following bullets describe types of data are made available to other OSS tools that do the triangulation, event prioritization, and triage functions; these OSS tools can correlate with Geographic Information System (GIS) and address information to isolate the common network point of the issue:

- Count and list of impacted cable modems
- Total MHz of impairment and spectrum location
- API reference to access the results of the pattern detection data analysis (such as the data informing Figure 25)
- API reference for historical data stored in the data lake related to the event
- API reference for real-time collection, on-demand, to enable fix agents to confirm the issue is still present, and to assist in both isolation and mitigation confirmation
- Other metadata to enrich the event, such as the mobile wireless carriers, or the more optimal PLC location


Two examples from our pattern detection framework are described below. Each of these has a specific pattern detection algorithm associated with the event forwarded through the notification service. To illustrate the pattern detection process, the following rule-based approach is used:

- **Sweep Generator Detection** (Figure 24)
  1. MER for all cable modems for the past 3 days (~n Polls) collected.
  2. Because the sweep insertion points come and go, we aggregate the minimum MER for all subcarriers to maximize the possibility of detecting sweep.
  3. The algorithm checks individual cable modems to see if there are subcarriers exactly 3 MHz or 6 MHz apart that have significantly lower MER than the overall MER.
  4. The algorithm excludes cable modems where many subcarriers around the potential sweep points also have significantly lower MER, as this may be a sign of another type of interference.
  5. If more than 30% of cable modems on a given OFDM channel meet the above criteria, we classify the OFDM channel as being affected by sweep generator activities, and send notifications to Comcast OSS tools.
- **Mobile Wireless Ingress Detection** (Figure 25)
  1. The following datasets are used for detecting Mobile Wireless Ingress:
    - o MER for all cable modems for the past 3 days (~n Polls) collected
    - o Spectrum allocation by state and county
    - o Account/cable Modem – ZIP code mapping
    - o ZIP Code – state mapping
  2. The algorithm identifies portions of the OFDM channel that overlap with other spectrum owners for individual cable modems by joining the above datasets
  3. For any overlapping spectrum and poll, the algorithm checks for the presence of the following 2 conditions; if both these conditions are met, the cable modem is classified as being impacted by Mobile Wireless Ingress:
    - o The average device MER is lower than the overall average MER by more than 3 dB
    - o Drops exist in MER between consecutive subcarriers of more than 3 dB on either side of the overlapping spectrum

In addition, we are currently developing alternative approaches for Pattern Detection by working with expert field techs to label MER curves with the various impairment categories. With that information, we can train classification models by applying supervised machine learning (ML) approaches to the problem.



**Figure 25 – Example MER charts for sweep generator insertion points in OFDM channel.**  
The sweep insertion points are 3 or 6 MHz apart.



**Figure 26 - Example of detected LTE Interference patterns, enriched with mobile carrier license data. The devices shown in this example experience two ingress patterns that perfectly overlap with the spectrums used by T-Mobile (~35-dB level) and Cellco Partnership (~20-dB level).**



## Future Work

The future work efforts are focused on measuring the system response to the application of profiles through the FEC data, supporting Orthogonal Frequency Division Multiple Access (OFDMA) and D3.0 SC-QAM for the upstream, and enhancing our telemetry collection infrastructure.

### 10. FEC

Noise in a communication channel introduces the possibility that a symbol will be received erroneously, thus degrading the effective capacity of the channel. For the same level of S/N, increasing the modulation orders increases the likelihood of a symbol being errored. Given the clear goal of maximizing the channel capacity, we must also make special consideration for the increased likelihood of receiving failed symbols.


Modern communication systems mitigate this by employing advanced FEC techniques. Correspondingly, FEC can also serve as a primary feedback signal -- assuming we can measure the *amount* of FEC used for each recommended and applied profile on the network. Future work will include the consideration of levels of FEC, as part of developing a fully closed- or feed-forward loop system, so as to be able to measure the response from the network from the changes in profiles applied, as outlined in this paper. With such a system, the new stated objective would be to maximize channel capacity, while simultaneously minimizing the probability of symbol errors. Such a closed- or feed-forward loop system can consider both traditional control theory solutions and advanced machine learning / reinforcement learning techniques.

All aspects of the algorithm discussed thus far can be modeled based on historical data, producing a reasonable view of the generated profiles, modem assignment, and capacity gains achieved. However, when it comes to FEC, simulating error rates (both corrected and uncorrected) is non-trivial. Lacking a proper estimate of error rates, any analysis assessing the merits of a certain approach against another will be biased towards the approach that results in higher overall modulation orders (i.e. we would be simply ignoring the resiliency dimension in the optimization game).

In contrast with the complexity of modeling FEC, deciding on how FEC will be used within the program is somewhat straightforward: Error rates will be monitored continuously, and upper end thresholds on those rates will be set. If a rate is crossed, a corrective action will be taken (either at a modem level or at a CMTS level). For example, increased error rates across a CMTS or across the entire footprint may indicate too aggressive modulation efficiency thresholds, while increased error rates for a specific modem may indicate a newly formed impairment. We need to be able to separate the two and respond with the appropriate action. An example of an appropriate action would be to adjust the global thresholding offset and/or the  $n$ th percentile hyper-parameters, in response to increased or reduced symbol error rates (in response to overly conservative/aggressive parameter values). To target an action at a cable modem level we will need to modify the existing data model to allow the  $n$ th percentile, and the thresholding offset, to be defined at the device level (rather than globally). We have recently started experimenting with measuring the system response in the lab and correlating the behavior of the system to the generated noise and applied profiles. These are the first steps towards implementing a closed-loop solution.

## 11. Up-Stream Signal Path Implications

Orthogonal Frequency Division Multiple Access (OFDMA, a version of OFDM) can be applied for shared usages scenarios, such as upstream channels -- where multiple access is effectively achieved by assigning subsets of the channel subcarriers to different users, as illustrated in Figure 27.



**Figure 27. Illustration of subcarrier configuration in OFDMA. In OFDMA, contiguous subsets of subcarriers are assigned to different users.**

Given the nature of a combined upstream plant and the nature of OFDMA, additional work will be necessary to adjust the analytic and control elements. We expect to find similar benefits to upstream capacity, including robustness, through the ability to manage and adapt upstream channel modulation profiles. We will update our findings accordingly.

D3.0 upstream channels already carry substantial flexibility, evidenced by variable modulation and error correction. While the network metrics are not yet as granular as those in D3.1, opportunities exist to optimize the performance -- as described in the case study, referenced earlier, and in which different FEC and modulation approaches have been successfully applied to mitigate the ingress impacts of conducted-switching power supplies. PMA for D3.0 US could apply one of a set of pre-defined modulation profiles to a channel. This solution would avoid applying the most robust configurations ubiquitously across the network, by applying them only when and where needed. Initial modeling, based on actual network data and currently conservative modulation profiles, indicates that a 10 to 15% increase in data bandwidth may be achievable. When it comes to the upstream signal path, even a modest improvement can provide significant benefits in deferred capital for node segmentation. These capabilities may also enable additional channels in the upstream spectrum, not currently in use because of challenging group delay or ingress characteristics.

## 12. Near-Real-Time Operation

Noise and ingress are often sporadic, periodic, and/or transient in nature, where the ability to detect and respond to the noise becomes highly dependent on the telemetry capture and sampling frequency. As we push further spectral efficiency and closer to Shannon limits, we'll need to fully understand what optimal sampling and the corresponding frequency of adaptive adjustments that are required to maintain maximum capacity with optimal robustness. Within these efforts, we're currently experimenting with a special on-demand telemetry collection that runs at much higher frequency compared to our standard solution. The new telemetry collection system referred to as "spot-beaming" targets a very small subset of OFDM interfaces in the network and runs on a sub-minute cycle. The current use case for this system is to collect MER data for specific interfaces of interest where gaining deeper knowledge of the MER dynamics is

required. Managing the targeted interfaces is done through an API, such that spotbeaming interfaces can be created or deleted on a needs basis. We still have much work to do in terms of understanding how large of coverage can the spotbeaming system target before running into performance-economic issues. These investigations will also help us understand the possibilities in terms of bridging the gap between our standard collection scheme and the spotbeaming system, with the objective of collecting telemetry in near-real-time across the entire network.

## Conclusion

Implementing PMA across Comcast's D3.1 network holds the opportunity to increase downstream bandwidth by ~30% while enhancing the resiliency of the network. The path to implementing PMA involved addressing a multitude of challenges including the hardening of D3.1, establishing a data architecture that supports capturing telemetry data for ~millions of cable modems at a sufficiently high sampling frequency, and establishing a lab for development & testing of the end-to-end PMA solution. We presented a core algorithm that recommends profiles that optimize the capacity of an OFDM Interface while taking into consideration the recommended mappings between MER and modulation orders, vendor constraints, and changes in MER over time. We are currently in the process of testing the PMA solution in the field in specific locations. In parallel to these efforts, we presented a host of pattern detection algorithms that complement PMA by ensuring that network issues are not masked by PMA, but rather reported to the proper entity within Comcast for taking corrective actions. Lastly, the development of the core algorithm continues with the anticipation of releasing a second version of PMA that operates in a closed-loop configuration—with the measured uncorrectable codeword error rate constituting a feedback signal to the algorithm.

## Abbreviations

CI/CD	continuous integration / continuous development
CM	cable modem
CMTS	cable modem termination system
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
dB	decibel
DOCSIS	data over cable service interface specification
FEC	forward error correction
GIS	Geographic Information System
HFC	hybrid fibre-coaxial
IQR	inter-quartile range
LDPC	low-density parity-check
LTE	long term evolution
MAC	media access control
MER	modulation error rate
N+X	node plus X number of amplifiers in cascade
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PLC	PHY link channel
PMA	profile management application
PNM	proactive network maintenance
QAM	quadrature amplitude modulation

SC-QAM	single carrier quadrature amplitude modulation
SDR	Software-defined radio
SoC	system on chip
SW	software

## Bibliography & References

1. White and Sundaresan, *DOCSIS 3.1 Profile Management Application and Algorithms*. SCTE IBSE, NCTA, CableLabs, Spring Technical Forum Proceedings, 2016.
2. Cablelabs, *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I02-EC*, 2014.
3. Rokach, Lior, and Oded Maimon. *Clustering methods*. Data mining and knowledge discovery handbook. Springer US, 2005.
4. Mutalik, Rice, Subramanya, Wang, *What Gets Measured Gets Done / What Gets Analyzed Gets Transformed, Analytics for a Wider/Deeper Network View*. SCTE IBSE TEC Expo, 2018.

# Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA)

A Technical Paper prepared for SCTE•ISBE by

**Karthik Sundaresan**

Distinguished Technologist

CableLabs

858 Coal Creek Circle, Louisville, CO 80027

3036613895

k.sundaresan@cablelabs.com

**Jay Zhu**

Senior Engineer

CableLabs

858 Coal Creek Circle, Louisville, CO 80027

3036613312

j.zhu@cablelabs.com

**Mayank Mishra**

R&D Wired Intern

CableLabs

858 Coal Creek Circle, Louisville, CO 80027

3036619100

M.Mishra@cablelabs.com

**James Lin**

Lead Engineer

Kyrio/CableLabs

858 Coal Creek Circle, Louisville, CO 80027

3036613871

j.lin@kyrio.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	5
1. OFDM and OFDMA Profiles .....	5
DOCSIS 3.1 Downstream .....	6
2. Noise Characteristics on an OFDM Channel.....	6
2.1. CMs with clean OFDM channels.....	7
2.2. LTE Ingress.....	7
2.3. Interference from a Sweep Generator .....	8
2.4. DS Roll off.....	8
2.5. Suckouts .....	9
2.6. Standing Waves .....	9
2.7. Downstream Tilt .....	9
3. Downstream Forward Error Correction behavior.....	10
3.1. Lab testing of DS FEC behavior on D3.1 equipment .....	10
3.2. Baseline test (no noise) .....	10
3.3. Test to discover failure points, noise across entire channel .....	11
3.4. Test QAM Level tolerance to noise .....	12
3.5. Test tolerance to narrowband noise (AWGN).....	12
3.6. Test narrow band noise (LTE ingress) .....	13
3.7. Understand profile switchover behaviour.....	16
4. OFDM Channel Configuration .....	17
4.1. OFDM Channel Location .....	17
4.2. PLC Location within OFDM Channel.....	17
4.3. OFDM Channel Parameters.....	18
4.3.1. Cyclic Prefix and Roll-Off .....	18
4.3.2. Interleaver.....	19
4.3.3. NCP modulation.....	19
4.3.4. Pilot multiplier .....	19
5. DS CM-STATUS Interactions and Settings.....	20
5.1. CM CM-STATUS State Machine.....	21
5.2. CMTS Management of CM-STATUS Messages.....	21
5.3. CMTS Polling CMs on RxMER and FEC.....	21
5.4. Recommended configurations .....	22
5.4.1. CM Event Thresholds for CM-STATUS Messaging.....	22
5.4.2. CMTS Thresholds for CM-STATUS.....	23
5.5. DS Profile Flapping.....	23
5.6. Suggested CM-STATUS Settings .....	25
6. Downstream Profile Mangement Application.....	26
6.1. DS PMA Software System architecture.....	27
6.2. DS PMA Practical gains.....	29
DOCSIS 3.1 Upstream .....	32
7. D3.1 OFDMA Upstream FEC behavior.....	32
7.1. Noise Characteristics on an OFDMA Channel.....	32
7.2. US FEC behavior.....	32
7.3. Lab testing of US FEC behavior on D3.1 equipment .....	33
7.4. Baseline test (no noise) .....	33
7.5. Test to discover failure points, noise across entire channel .....	33
8. OFDMA Channel Configuration.....	34
8.1. US Channel Location.....	34

8.2.	TaFDM .....	34
8.3.	Ranging location.....	35
8.4.	Minislots .....	35
8.5.	IUC / Profile management.....	36
8.6.	OFDMA Profile Flapping.....	36
8.7.	Upstream channel evaluation tools .....	37
8.8.	US profile change CER based vs RxMER based handling .....	37
9.	US PMA.....	37
9.1.	Baseline test on 256-QAM .....	38
9.2.	US Noise injection test.....	38
9.3.	US PMA test.....	39
	Conclusion .....	39
	Abbreviations.....	40
	Bibliography & References .....	40
	Acknowledgements .....	40

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Examples of Clean RxMER from CMs.....	7
Figure 2 –LTE Interference .....	7
Figure 3 –Sweep Generator Ingress .....	8
Figure 4 – OFDM Channels in the Roll-off.....	8
Figure 5 – Examples of Suckout .....	9
Figure 6 – Examples of Standingwave .....	9
Figure 7 – Examples of DS Tilt .....	9
Figure 8 – Average D3.1 RxMER for a clean channel when carrying traffic .....	10
Figure 9 – Average D3.1 RxMER levels for Correctable and Uncorrectables.....	11
Figure 10 – Alternate view Corrected Codeword percentage vs RxMER.....	11
Figure 11 –D3.1 RxMER levels for 100 % Corrected vs QAM level .....	12
Figure 12 –D3.1 Noise Injection test .....	12
Figure 13 –D3.1 Noise Injection test results .....	13
Figure 14 –D3.1 DS RxMER plots w Noise Injection LTE signal (1 channel).....	13
Figure 15 –FEC behavior w noise injection of LTE signal (1 channel).....	14
Figure 16 – D3.1 DS RxMER plots w Noise Injection LTE signal (3 channels).....	14
Figure 17 – FEC behavior w noise injection of LTE signal (3 channels).....	15
Figure 18 –D3.1 DS Uncorrectables at the instant of Ingress (2 different CMs).....	15
Figure 19 –CM-CMTS DS Profile Switchover at specific.....	16
Figure 20 – Profile for an OFDM Channel in the Roll Off .....	17
Figure 21 – LTE Ingress in cable plant, seen in CM RxMER (729-756 MHz) .....	18
Figure 22 – D3.1 Profile Flapping Behavior.....	24
Figure 23 – D3.1 Profile Flapping Example .....	25
Figure 24 – Profile Management Application Deployment Architecture .....	26
Figure 25 – Profile Management Application Software Architecture .....	27
Figure 26 – Typical PMA Gain vs number of profiles (J-K Correlation).....	29

Figure 27 -PMA capacity gain histogram when using 3 profiles .....	29
Figure 28 -PMA capacity gain histogram when using 5 profiles .....	30
Figure 29 -PMA capacity gain histogram when using 8 profiles .....	30
Figure 30 -PMA capacity gains, 2 profiles vs 10 profiles on a channel .....	31
Figure 31 -Recommended number of profiles .....	31
Figure 32 – Multiple measurements of D3.1 US RxMER from a CM .....	32
Figure 33 –D3.1 US RxMER from a lab CM .....	32
Figure 34 – Average D3.1 US RxMER levels for Correctable and Uncorrectables .....	34
Figure 35 – Time and Frequency division multiplexing .....	35
Figure 36 – OFDMA Channel configuration example.....	35
Figure 37 –D3.1 US Baseline test .....	38
Figure 38 –D3.1 US Noise Injection test 0.2 MHz.....	38
Figure 39 –D3.1 US Noise Injection test 1 MHz.....	38
Figure 40 –D3.1 US Noise Injection test 5 MHz.....	39
Figure 41 –D3.1 US Noise Injection test 5 MHz.....	39

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - DS RxMER to QAM Level mapping.....	10
Table 2 – CM-STATUS Events related to DS Channel Failure.....	20
Table 3 – CMTS CM-STATUS Settings.....	25
Table 4 – CMTS Maximum reports and ACK.....	25
Table 5 - US RxMER to QAM Level mapping.....	33



# Introduction

DOCSIS 3.1 OFDM/A Profiles provide a wide range of modulation choices that can be used to fine-tune the CMTS downstream and CM Upstream transmissions to get the best performance from the current network conditions. A well-designed, optimized set of modulation profiles allows a channel to operate more robustly against ingress noise and also enables an overall higher user throughput.

This paper will discuss the D3.1 Profile Management Application and how it is used by operators in deployment with D3.1 CMTSs and CMs to create Downstream profiles and Upstream IUCs/Profiles. The paper will share our experiences from creating such an application. It will discuss the effect of noise on the choice of a profile assigned to a CM. We share results on how the number of FEC codewords corrected vary with the noise, and at different modulation orders. We also talk about the gain in the network capacity seen by smart Profile creation algorithms versus using a flat profile and simply assigning CMs to the least common denominator profile which fits the CM.

The paper will also describe MAC layer state machine and interactions between the CM & CMTS when a profile fails. The interaction between CM-STATUS messages from a CM (for flagging failures and recovery on a profile) and the CMTS response to that message in changing the profile used, results in “Profile Flapping”. This paper recommends on how to design CM and CMTS MAC Layer settings to make the system be robust when these interactions take place in the D3.1 system.

D3.1 OFDM/A channels can have interference in parts of the channel and different modems experience this differently. Deploying well designed profiles for each channel will decrease the number of errors seen on the channel, reducing trouble calls. It could also unlock a solid 200~400 of Mbps of capacity gain on each channel.

## 1. OFDM and OFDMA Profiles

DOCSIS 3.1 specifications introduced features that leverage the OFDM-based PHY layer, including variable bit loading, and the ability to define multiple modulation profiles on downstream and upstream channels. Other new features include the ability to measure the quality of a downstream channel and test out the profiles in use, and features like upstream probes to measure the quality of the upstream OFDMA signal. The new MAC management messages support this on a per CM basis. There are also extensive additions to important operational items within proactive network maintenance (PNM) which enables measurement of various physical layer metrics and exposes that data to the operators.

The configuration, initiation logic and compute processing needed to optimize some of these functions Downstream Profile creation, Upstream IUC/Profile creation, are not defined in the DOCSIS 3.1 MAC and PHY specifications. This allows such functionality to be moved out of a CMTS and implemented as a Profile Management Application (PMA) running outside the CCAP. Here the idea is to move the profile creation process as an application external to the

CMTS. The PMA can communicate with a data lake and the CCAP to gather the needed information, process the data, and make intelligent decisions to set up the CCAP as needed.

To leverage the new OFDM/A physical layer to its maximum benefit, different subcarriers use different modulation orders. Optimizing the downstream/upstream profiles allows a downstream/upstream channel to be able to operate with lower Signal-to-Noise Ratio (SNR) margin, potentially allowing a channel to operate at an overall higher throughput. The logic to achieve this can be external to a CCAP and enable innovation. For a cable operator, it allows uniform operation of such algorithms across different CCAP platforms.

[D31PMA-INTX16] describes methods for designing OFDM/A profiles and choosing the appropriate modulation orders for a profile. It answers the questions around which profile is appropriate for a CM and what is the optimal set of profiles to use across the an OFDM/A channel for a given set of CMs.

A D3.1 CM supports two or more OFDM channels, each occupying a spectrum of up to 192 MHz in the downstream. The OFDM signal is composed of: Data subcarriers, Scattered pilots, Continuous pilots and PLC subcarriers. A modulation profile consists of a vector of bit-loading values, an integer value for each active subcarrier in the downstream channel. The modulation orders range from 16-QAM to 16384-QAM, the range of bit-loading values is from 4 to 14 (skipping 5); however, it is expected that very low bit-loading values, 7 or less, will be used very infrequently since most plants support 256 QAM today, but those will likely be in use in the roll-off regions.

Each CM will support and can be assigned up to four modulation profiles, including Profile A (used for broadcast frames), an optimized profile for the CM's unicast traffic, and possibly two additional profiles that could be used for multicast traffic.

A CMTS on each OFDM channel, per the DOCSIS specifications, needs to support up to 16 profiles. In the short term, CMTSs support 3~4 profiles per channel and many CMTS equipment assign all the profiles they support to all the modems. Today these profiles are flat, i.e. the same modulation order across the whole channel. In the long term as CMTS vendors support more than 4 profiles per channel, they will assign one profile A and one optimized profile to a CM. There may be a third profile assigned as a fall back profile between the optimized profile and profile A.


## **DOCSIS 3.1 Downstream**

### **2. Noise Characteristics on an OFDM Channel**

Several different types of noise have been identified in the field creating challenges with profile flapping and partial service flapping with the OFDM channel. These includes LTE ingress, Sweep Generator, Suckouts, Channel Roll-offs, Tilt etc. Examples of these are shown in the figures below. Any of these noise ingress can cause uncorrectable codewords on the channel, depending on the severity. All the below graphs show the RxMER level measured at the CM along the y-axis, and the frequency of the OFDM channel in MHz, along the x-axis.

## 2.1. CMs with clean OFDM channels

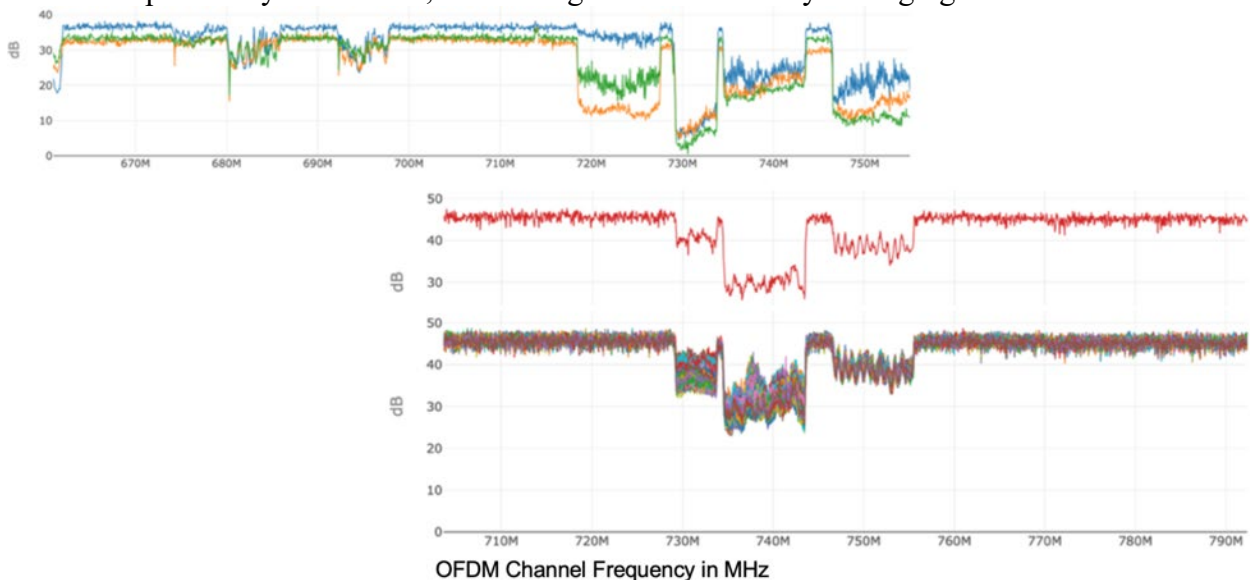
Depending on the plant, many CMs have quite good and clean RxMER levels across the whole D3.1 OFDM channel. This example is a 96 MHz channel starting at 696 MHz (in 850 MHz plant).



**Figure 1 – Examples of Clean RxMER from CMs**

## 2.2. LTE Ingress

The below picture shows two channels, a 96 MHz OFDM channel from 660-756 MHz, and second one from 700-796 MHz. Both of these channels clearly show LTE ingress noise in the CM's DS RxMER, as detailed later in this paper. As seen the LTE ingress is at a specific location in frequency but the ingress level is highly time variant. The bottom graph in this figure shows samples every 30 seconds, and the ingress level is always changing.



**Figure 2 –LTE Interference**

### 2.3. Interference from a Sweep Generator




Figure 3 –Sweep Generator Ingress

### 2.4. DS Roll off

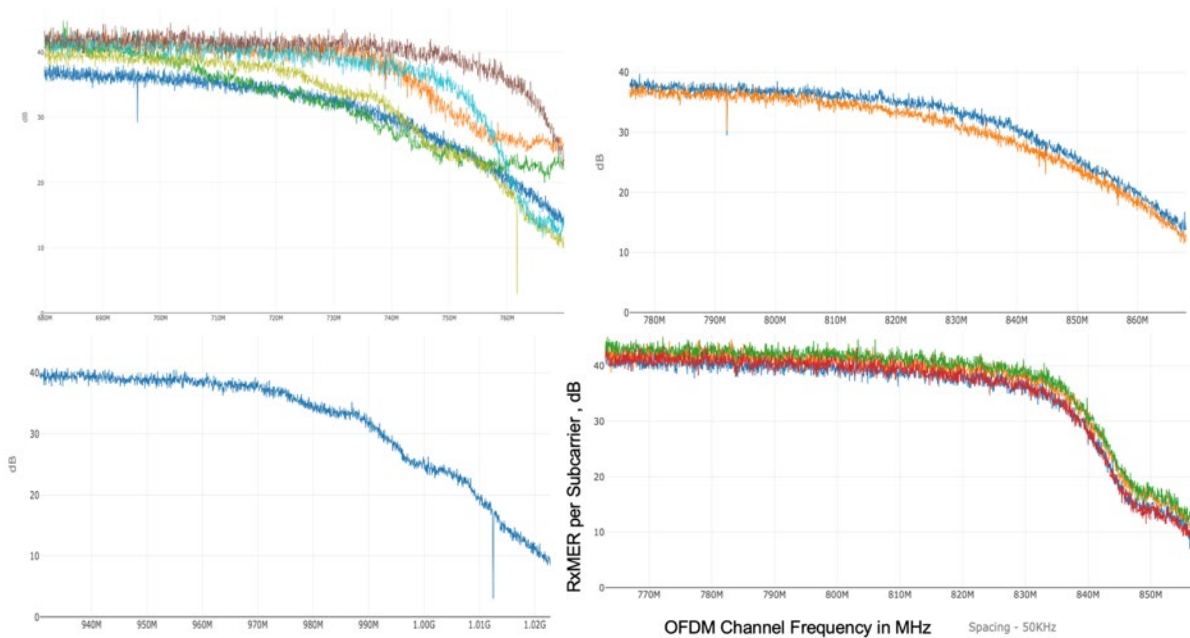


Figure 4 – OFDM Channels in the Roll-off

## 2.5. Suckouts




Figure 5 – Examples of Suckout

## 2.6. Standing Waves




Figure 6 – Examples of Standingwave

## 2.7. Downstream Tilt




Figure 7 – Examples of DS Tilt

### 3. Downstream Forward Error Correction behavior

Forward error correction (FEC) is a way of adding redundancy to messages so that the receiver can both detect and correct common errors. The D3.1 system uses an outer BCH coding and an inner LDPC coding. The devices support the 8/9 code rate for the codeword (16,200 bits). Codeword shortening is used when there is insufficient data to fill complete codewords and to achieve strong burst noise protection.

A PMA needs to be able to map RxMER values in the Downstream, to an appropriate QAM level, when creating a profile. These mappings are defined in [PHYv3.1] and summarized below.

**Table 1 - DS RxMER to QAM Level mapping**

Downstream Constellation/ Bit Loading	DS MER (dB)
16 QAM	15.0
64 QAM	21.0
128 QAM	24.0
256 QAM	27.0
512 QAM	30.5
1024 QAM	34.0
2048 QAM	37.0
4096 QAM	41.0
8192 QAM	46.0
16384 QAM	52.0


#### 3.1. Lab testing of DS FEC behavior on D3.1 equipment

We wanted to understand the performance of the DS FEC on real D3.1 equipment, in an effort to understand at what points will an operator start seeing failures in the system. We are testing these with 2 different CMTSs and 2 different CMs. We hope to expand the testing to include other devices. The test was run on a 96 MHz OFDM channel at 300 MHz and 702 MHz

#### 3.2. Baseline test (no noise)

This test checks RxMER at the CM and runs downstream traffic from the CMTS to the CM:

1. Send traffic to modem starting at 200 Mbps for 30 seconds.
2. Increase traffic in steps of 100 Mbps until Traffic rate of 500 Mbps.
3. Repeat steps 1 and 2 for the profile configured at each QAM level between 256QAM – 512QAM-1024QAM-2048QAM-4096QAM.



**Figure 8 – Average D3.1 RxMER for a clean channel when carrying traffic**

For a particular flat profile (or modulation order), there was little significant difference between the average RxMER values across the channel at different traffic rates both at 300 or 700MHz. However, the difference between the RxMER values measured when using different profiles at different modulation orders varies a bit. This was something we couldn't readily explain as to why the RxMER measurements could be different. In this baseline test there were no corrected FEC codewords for any of the 5 modulation profiles.

### 3.3. Test to discover failure points, noise across entire channel

The next step was to determine at what points would CMs start seeing DS codeword errors. The idea here is to increase the noise floor on the channel (AWGN) and see how the system performs. For each modulation order: QAM 256, 512, 1024, 2048, 4096, the goal was to identify the average RxMER of the channel at which

- the first corrected codeword is seen
- 100% corrected Codewords are seen
- the first uncorrected Codeword is seen




Figure 9 – Average D3.1 RxMER levels for Correctable and Uncorrectables




Figure 10 – Alternate view Corrected Codeword percentage vs RxMER




We identified the different code word corrected codeword failure levels at the CMs, and the corresponding MER values tied to those failure point. As seen in the figure there is an inverted S-curve growth in corrected codewords as noise increases. As expected the lower the modulation order, the more noise is needed to get to the first corrected, 100% corrected and first uncorrectable codewords.

Most CMs stopped receiving data on the profile as soon as the first few uncorrectable codewords are seen. This is because the CM sends out a CM-STATUS message and the CMTS reacts by downgrading the profile or putting the CM in partial service.

### 3.4. Test QAM Level tolerance to noise

This test was to determine if higher modulation orders are progressively more susceptible to noise. The test methodology was to send traffic to the CM at 300 Mbps. We then insert AWGN noise with a MER of 50db over entire channel. The noise was increased in steps of 1db, each noise level was maintained for 15 seconds. The test continued until the modem reported 100% corrected codewords, for each QAM level between 256– 4096.




**Figure 11 –D3.1 RxMER levels for 100 % Corrected vs QAM level**

The RxMER values that yield 100% corrected codewords grows linearly with higher QAM.

### 3.5. Test tolerance to narrowband noise (AWGN)


The test is to determine the performance of downstream FEC when injecting narrow band noise to the OFDM channel. In this test, the downstream traffic is set to a rate of 500 Mbps. Different noise widths including 1MHz, 12MHz, and 25MHz are injected into the channel. The following figure shows the OFDM downstream RxMER with different noise output attenuations when injecting 1MHz noise:



**Figure 12 –D3.1 Noise Injection test**



The noise injection test results are put together with the baseline test result (full-band AWGN noise) for comparison. The following figure shows the results with 256-QAM as the data profile modulation order:



**Figure 13 –D3.1 Noise Injection test results**

As the noise width increases, the LDPC decoder has to correct more codewords. Wider noise injection causes the LDPC decoder to reach 100% corrected codewords at higher RxMER levels (lower noise power) than narrower noise injection. Also, a wider noise width creates uncorrectable codewords at higher RxMER levels and the CM goes into partial channel state on the OFDM channel. This can cause the CMTS to move the CM to a lower modulation order profile or remove the channel from the CM' receive channel set to ensure service quality.


### 3.6. Test narrow band noise (LTE ingress)

We recorded an LTE signal, a single channel at ~10 MHz bandwidth. We injected this into an OFDM channel and gradually decreased the output attenuation. The plot below shows CM downstream RxMER with different noise output attenuation levels that we tested:



**Figure 14 –D3.1 DS RxMER plots w Noise Injection LTE signal (1 channel)**

The following plot shows the test results with 2 different flat profiles with modulation order at 256-QAM and 1024-QAM. The corresponding average RxMER values shown are calculated from CM downstream RxMER values only within the ingress frequency range.



**Figure 15 –FEC behavior w noise injection of LTE signal (1 channel)**

Compare to the previous test result of injecting 12MHz AWGN noise (256-QAM), the LTE ingress noise adds slightly more pressure on the LDPC decoder. The CM starts to have corrected codewords at 38.75dB, which is 6dB higher than where we observed corrected codewords with 12MHz AWGN noise. However, we observed 100% corrected codewords around 27dB with both noise sources at 256-QAM.


We also recorded an LTE signal with 3 adjacent LTE channels. We injected this to an OFDM channel and gradually decreased the output attenuation of the noise signal. The plot below shows CM downstream RxMER with different noise output attenuation levels that we tested:

The following plot shows the test results with 2 different flat profiles with the modulation orders at 256-QAM and 1024-QAM. The corresponding average RxMER:




**Figure 16 – D3.1 DS RxMER plots w Noise Injection LTE signal (3 channels)**

The following plot shows the test results with 2 different flat profiles with the modulation orders at 256-QAM and 1024-QAM. The corresponding average RxMER values shown are calculated from CM downstream RxMER values only within the most powerful LTE channel's ingress frequency range.



**Figure 17 – FEC behavior w noise injection of LTE signal (3 channels)**

When the noise ingress level is large enough, uncorrectable codewords can be observed at the moment the noise injection into the OFDM channel happens. The test injects LTE noise 3 different times. As the ingress noise power increases, an increasing number uncorrectable codeword spikes can be observed (0.18% uncorrectable codewords at the highest spike for one CM, and 0.5% for another CM).



**Figure 18 –D3.1 DS Uncorrectables at the instant of Ingress (2 different CMs)**

The CM also sends CM-STATUS (16-Profile Failure) messages in order to inform the CMTS about its codeword failures on a profile. In the field, ingress/burst noise from the environment can be powerful enough to cause the CMs to have uncorrectable codewords and send CM-STATUS messages to the CMTS. As discussed in detail in Chapter 5, the CMTS acts based on

CM-STATUS messages and changes the CMs' profile, and over time this happens back and forth. The probability of capturing & observing ingress/burst noise in the cable plant is low, since MSOs are capturing RxMER measurements a few times a day, on a CM is low. This can make it hard to diagnose the root cause of data rate issues in the system. Different CMTS/CM vendors can use different algorithms/thresholds to generate & handle CM-STATUS messages, which leaves room for system improvements on this issue.

### 3.7. Understand profile switchover behaviour

This test determines the behavior of the system when using multiple flat profiles and when noise is injected. The goal was to identify the RxMER values at which the system switched profiles (to a lower modulation order profile or back up to a higher modulation profile). Traffic was sent to the CM at 500 Mbps. The CM was ensured to be on the highest profile (Profile 2, Profile 0 is the default). The test introduces noise until the system changes the profile for the traffic to profile 1 and then to Profile 0. The test also removes the noise slowly and notes when the CM returned to profile 1 and then back to Profile 2. When injecting noise, on a downward move from a higher profile to a lower profile, the movement is immediately after the CM sends a CM-STATUS message. On the way back, from a lower profile to a higher profile, there is a built-in hysteresis in the CMTS settings for profile recovery and the time to recover depends on the settings supported by the CMTS.





Figure 19 –CM-CMTS DS Profile Switchover at specific

## 4. OFDM Channel Configuration

There are various items to consider for an operator when Deploying OFDM channels and configuring them on the plant. These include the location of the OFDM channel itself and the location of the PLC within the channel. Care needs to be taken to work around interference/ingress noise which can affect system performance. There are also various physical layer settings like cyclic prefix, Roll-off, Interleaver depth, etc.). There are also many MAC layer settings like CM-STATUS messaging which can affect system stability. Care needs to be taken to avoid problems like profile flapping and ultimately there is a need to build a solution like PMA.

### 4.1. OFDM Channel Location

The CM supports a two independently configurable OFDM channels each occupying a spectrum of up to 192 MHz in the downstream. As the operator opens up spectrum to locate OFDM channels, some operators need to put the OFDM channels in the roll-off region in their plant. Per the results seen above, CMs which are in the roll-off will suffer from FEC codeword errors in the roll off region. At a minimum the operator will need to design profiles which lower the modulation orders used in the roll-off region, to ensure the CMs have a reliable data connection.



**Figure 20 – Profile for an OFDM Channel in the Roll Off**

### 4.2. PLC Location within OFDM Channel

The aim of the PLC is for the CMTS to convey to the CM the physical properties of the OFDM channel. When acquiring the OFDM channel, the CM needs the physical parameters of the channel. The CM first acquires the PLC, and from this extracts the parameters needed to acquire the complete OFDM channel.

The CMTS places the PLC at the center of a 6 MHz encompassed spectrum with no excluded subcarriers and places the 6 MHz encompassed spectrum containing the PLC on a 1 MHz grid. For 4K FFT OFDM, this 6 MHz will contain 56 subcarriers followed by the 8 PLC subcarriers followed by another 56 subcarriers. For 8K FFT OFDM, this 6 MHz will contain 112 subcarriers followed by the 16 PLC subcarriers followed by another 112 subcarriers. The PLC consists of 8 symbols of preamble followed by 120 symbols of data subcarriers, 16-QAM is used for the PLC subcarriers. The CMTS interleaves scattered pilots and data subcarriers, but does not interleave continuous pilots, the PLC, and subcarriers belonging to excluded regions. The insertion of continuous pilots, PLC and excluded regions happens after both time and frequency interleaving. There is no interleaving of the 400 kHz of PLC subcarriers (8 or 16).

So essentially any interference in the same spectrum as the PLC means that the CM could lose the PLC and hence the OFDM channel itself. It is very important that the operator and the CMTS


place the PLC in a part of the spectrum that is less susceptible to noise and interference. The operator with an overriding configuration is expected to place the PLC as appropriate within the channel.

Many operators have built a geographical map of the LTE and other wireless channels in their footprint. This map is then correlated to the location of the OFDM channels in the cable plant in those areas. Specifically, the idea is to find the spectrum with the least amount of overlap with wireless carriers in that region. The best 6 MHz chunk of spectrum within the OFDM channel is then chosen to locate the PLC.

As an example, LTE Band 12,13, 17 are some of the more common bands used by cellular providers in USA [LTEFreq]. These correspond to the following frequencies:

- Band 12: Uplink of 699 – 716 MHz, Downlink of 729 – 746MHz,
- Band 13: Uplink of 777 – 787 MHz, Downlink of 746 – 756 MHz,
- Band 17: Uplink of 704 – 716 MHz, Downlink of 734 – 746 MHz.

A DS OFDM Channel which overlaps an LTE band, say from 648-744MHz or 702-798MHz, could see LTE ingress from these bands. See figure below for an example of how LTE ingress into the cable plant can affect the RxMER of a CM in that region.



**Figure 21 – LTE Ingress in cable plant, seen in CM RxMER (729-756 MHz)**

Similarly, for European deployments in addition to LTE bands, Digital Audio Broadcast (DAB+) uses a wide-bandwidth broadcast technology and typically spectra have been allocated for it in Band III (174–240 MHz). The recommendation is to be aware and keep the PLC out of the DAB+ Region and any other potential source of ingress into the cable plant.

### **4.3. OFDM Channel Parameters**

From many of the initial deployments, the operators have quickly realized that there are many CMTS defaults that need to be revisited, to optimize the channel operation.

#### **4.3.1. Cyclic Prefix and Roll-Off**

The addition of a cyclic prefix (CP) enables the receiver to overcome the effects of inter-symbol-interference caused by micro-reflections in the channel. The time duration of the CP should be chosen to be longer than the time of the longest significant reflection. Windowing maximizes channel capacity by sharpening the edges of the spectrum of the OFDM signal. Spectral edges occur at the two ends of the spectrum of the OFDM symbol, as well as at the ends of internal exclusion bands. The roll off must be integrated within the duration of the CP. The choice of a larger CP from 0.9375 $\mu$ s to 5 $\mu$ s, will affect the efficiency and the capacity accordingly from 1.8 Gbps to 1.5 Gbps, for a 192 MHz OFDM channel. A CP of 1.25 $\mu$ s and a Roll-off period of 0.625 $\mu$ s have given a few operators the stability and performance they are looking for.

### **4.3.2. Interleaver**

The OFDM symbols, comprised of data subcarriers, scattered pilot placeholders, and NCPs, are subject to time and frequency interleaving. Time interleaving mitigates the impact of burst noise, while frequency interleaving mitigates the effect of ingress.

Time interleaving disperses the subcarriers of an input symbol over a set of output symbols, based on the depth of interleaving. Therefore, if an OFDM symbol is corrupted by a noise burst, this burst is spread over the symbols when it is de-interleaved, thereby reducing the error correction burden on the decoder.

Frequency interleaving occurs after time interleaving. Frequency interleaving disperses subcarriers of the symbol along the frequency axis; therefore, OFDM subcarriers impacted by narrowband ingress are distributed between several codewords, reducing the number of errors in each codeword. For time interleaving, the CMTS supports a maximum value of M equal to 32 for 20  $\mu$ s symbol duration (50 kHz subcarrier spacing) and 16 for 40  $\mu$ s symbol duration (25 kHz subcarrier spacing).

Initially running the system with M=16, for 50 kHz spacing looks to protect against many noise sources. Care also needs to be taken that frequency interleaving is enabled, one of the early CMTS implementations turned off frequency interleaving in some cases.

### **4.3.3. NCP modulation**

When the data codewords are mapped to subcarriers within a symbol, a pointer is needed to identify where a data codeword starts, this is the Next Codeword Pointer (NCP). There are a variable number of NCP message blocks (MBs) on each OFDM symbol each pointing to the next codeword. Each FEC encoded NCP MB is 48 bits. The NCP QAM constellation can be QPSK, QAM-16, QAM-64.

The CMTS places the NCP subcarriers beginning from the frequency location of the highest frequency active data subcarrier of the OFDM symbol, and going downwards along active data subcarriers of the OFDM symbols before they are time and frequency interleaved. The CMTS time and frequency interleaves the NCP subcarriers using the algorithm applied to data subcarriers. This essentially allows the NCP modulation order to be QAM-16 or QAM-64 in practice. The QAM-16 looks to be sufficient to handle all variations in the plant noise across CMs in majority of deployments and many deployments are running with QAM-64 as well.

### **4.3.4. Pilot multiplier**

Downstream pilots are subcarriers modulated by the CMTS with a defined modulation pattern that is known to all the CMs in the system to allow interoperability. There are two types of pilots: continuous and scattered. Continuous pilots occur at fixed frequencies in every symbol. Scattered pilots occur at different frequency locations in different symbols.

The Number of Continuous Pilots =  $\min(\max(8, \text{ceil}(M * (F_{\text{max}} - F_{\text{min}}) / 190e6))), 120)$ .

The value of M in equation is as a parameter between ( $120 \geq M \geq 48$ ) that can be adjusted by the operator or the CMTS. The typical value seen in deployments for M is 48, which will give 56 pilots (48 + 8 PLC pilots) for a 190 MHz channel. There is no observed change in behavior/stability by increasing the pilot scale.



## 5. DS CM-STATUS Interactions and Settings

CM-STATUS messages are needed in cases where the CM detects a failure that the CMTS cannot detect directly, or where the CM can send valuable information to the CMTS when an error or a recovery event occurs (for example, the CM can report a T3 timeout to the CMTS). Upon receiving an error indication, the CMTS is expected to act in order to correct the error. A CM transmits a CM-STATUS message on any available channel when it detects an event condition and the reporting of the event type is enabled on the CM. These events are summarized in table CM-STATUS with the Event Type Codes and Status Events and are also defined in detail in the DOCSIS specifications [MULPIv3.1]. Some event types are for a particular downstream channel, a particular upstream channel, or a DSID. For each such event, the CM maintains a separate state variable as to whether the event condition is considered "on" or "off" for each channel or DSID. The CM-STATUS message includes the Event Type Code and a unique Transaction ID for each occurrence of the event (IDs start from 1 and go up to 65535 before wrapping around to 1).

Fault detection and recovery occurs at multiple levels. At the physical level, FEC is used to correct errors where possible. The MAC protocol protects against errors through the use of checksum fields across both the MAC Header and the data portions of the packet. All MAC management messages are protected with a CRC covering the entire message. At the network layer and above, the MAC Sublayer considers messages to be data packets protected by the CRC field of the data packet; any packets with bad CRCs are discarded. Recovery from these lost packets is in accordance with the upper layer protocol.

When loss of lock is detected on an OFDM channel, it is possible that the channel is still partially functional and receives data. So, in case of high FEC errors detected in PLC, NCP or Profile A, the CM attempts to continue using the channel and enters a Partial Channel Mode. If the upstream communication is available, the CM sends out a CM-STATUS message to inform the CMTS of the Lost Lock event. If CMTS becomes aware of an interruption of a CM's Primary Downstream Channel (via a CM-STATUS message from the affected CM or from another CM), the CMTS can take an appropriate action. It can potentially send a DBC-REQ to the CM to reprogram the downstream channel or the profile as appropriate.

**Table 2 – CM-STATUS Events related to DS Channel Failure**

Event Code	Event Description	Channel applicable	Introduced in
2	QAM/FEC Failure	SC-QAM	D3.0
5	QAM/FEC Recovery	SC-QAM	D3.0
16	OFDM Profile Failure	OFDM	D3.1
24	OFDM Profile Recovery	OFDM	D3.1
20	NCP Profile Failure	OFDM	D3.1
22	NCP Profile Recovery	OFDM	D3.1
21	PLC Failure	OFDM	D3.1
23	PLC Recovery	OFDM	D3.1



## **5.1. CM CM-STATUS State Machine**

When the CM detects an event, it creates a CM-STATUS message and it starts a timer for a value between 0 and the max-hold-off with 20ms resolution, before sending the message. When the timer expires the CM-STATUS message is placed sent to the CMTS.

The CM will send the CM status message again if it has not reached the max-reports number of times, or it has not received a CM-STATUS-ACK from the CMTS (with the same transaction number for the same event ID), or the event type has not been disabled. After the first message it will send the message again if applicable on the period of the full max-hold-off timer. The CM transitions back to the IDLE state after the max-reports or a CM-STATUS-ACK is received.

The recovery and failure events are treated as separate events, and as we have seen in the field a recovery message can quickly follow a failure message based on the CM algorithms. If the status transitions from “on” to “off” to “on” while a hold-off is in place in the SENDING state, the max reports is reset and additional CM-STATUS messages are sent as a new transaction-Id.

## **5.2. CMTS Management of CM-STATUS Messages**

When the CMTS receives a CM-STATUS from a CM reporting a Partial Channel Mode, the CMTS can either stop sending data for that modem to the reported profile or move the service flows for the CM to another working profile on that channel. The CMTS attempts to resolve partial channel situations, such as by shifting the service flow to other profiles or other channels. Once the CMTS receives the CM-STATUS message it may react based on any algorithm it chooses. The CMTS may use a DBC message to ask the CM to make a change to its RCS (add or delete a DS channel), or ask the CM to change the DS profile in use. The DBC transaction includes the following: A Request (DBC-REQ) from CMTS to CM, the CM sending back a DBC Response (DBC-RSP) to the CMTS, and the CMTS acknowledging with a DBC-ACK. These message exchanges will require a request and grant cycle and may be on the order of 100s of ms (maybe 300-800 ms). The CMTS may choose to immediately quit forwarding data on the profile or the channel if it believes the CM can receive the data on another channel or profile which the CMTS sends the data.

Some CMTS implementations immediately trigger a profile update upon receiving a CM-STATUS message if the CM is not in profile 0, or an RCS change if the CM is in profile 0. The only way to recover the profile is based on CM-STATUS messaging. This can be used to add stability from the CMTS perspective. If no other CM-STATUS failures occur for that profile during the timer period, then the CMTS will move the traffic back to the higher capacity profile as soon as the OFDM recovery hysteresis time expires.

The CMTS picks the profile order by taking the bps/Hz for each subcarrier and summing all active subcarriers ranking from highest capacity least robust to lowest capacity most robust.

## **5.3. CMTS Polling CMs on RxMER and FEC**

There are a couple of different ways the CMTS can check how the CM is performing on the Downstream. The OPT-REQ MAC Management message is used by the CMTS to cause a CM to test various aspects of an OFDM downstream channel. A single OPT-REQ message can be used to test the CM's ability to receive the specified downstream OFDM profile by checking for FEC statistics, alternatively it can be used to query the CM's RxMER statistics.

CMTSs today follow one or both of the above approaches, they periodically check for RxMER of the CM or the FEC statistics for a particular profile. A CMTS does this periodically by

issuing OPT-REQ messages to every CM on the OFDM channel. The rate at which the CMTS does this can be configured in a vendor proprietary manner.

The CMTS compares the numbers from these measurements with a certain threshold, before it decides it needs to downgrade or upgrade the profile on which the traffic reaches the CM.

#### **5.4. Recommended configurations**

The configuration settings for the CMTS-CM CM-STATUS state machines have the following goals. First is to avoid any unnecessary profile flapping in the presence of transient noise or noise near modulation profile error thresholds. The next is to ensure that data is forwarded to the CM on the correct profile when a more robust profile is chosen. The goal is to avoid unnecessary partial service flapping of the OFDM channel in the presence of noise that impacts profile 0. An operator wants to ensure that when the CM enters the partial service state with the OFDM channel no longer available, data can still be sent over the D3.0 channels.

##### **5.4.1. CM Event Thresholds for CM-STATUS Messaging**

CMs implement different methods to identify the failure of a profile on which it is receiving data. A CM also can implement different failure/recovery thresholds, for each event e.g. failure/recovery of PLC, NCP, Profile 0, Profile 1,2,3 etc. There are different ways of detecting profile codeword errors: using a raw FEC CW error count vs. using a count of time over which errors have occurred.

The first method can be thought of as a count of the Number of Errored Codewords. e.g. if 50 out of 1000 codewords have errors, the CM could declare the profile to have failed and raise the event by sending a CM-STATUS message. Similarly, if 990 codewords out of the last 1000 have no errors, then the CM can declare that the profile is working correctly again and clear the event with another CM-STATUS message.

The second method can be thought of as a count of the time there are errored Codewords on the channel. e.g. if 2 out of last 20 seconds have codeword errors, the CM could declare the profile to have failed and raise the event by sending a CM-STATUS message. Similarly, if 18 out of the last 20 seconds have no errors, then the CM can declare that the profile is working correctly again and clear the event with another CM-STATUS message.

Another method which a failure conditions can be detected would be a combination of the codeword error count and time for which errors are present.

Both the failure threshold and the recovery thresholds can be adjusted to bring out the desired behavior. This then leads to the question of what is the desired behavior, which leads to the question which kind of noise sources are we trying to recover from quickly. If an operator can characterize the noise sources in the plant, then these settings can be fine-tuned for each kind of plant. e.g. a short bursty noise which is present for some amount of time, or a noise which may be present for a longer time durations or strong noise which affects large number of subcarriers and hence codewords, vs a low-grade noise which only causes a small number of errored codewords.

The philosophy which yields the most optimal results is as follows: *Fail fast and recover slowly.*

This means we want the CMs to detect failures in profiles, as soon as possible. This also needs to happen without being unreasonable, i.e. we don't want to declare a profile failure with say just a few codeword errors. So once the errors in a profile are confidently detected then it is better to raise the failure as soon as possible. The idea is to minimize the impact to the customer data

traffic. The earlier the profile failure occurs the sooner the CMTS can move the data to a more robust profile.

Also, on the recovery side the CM needs to make sure the profile is very stable before declaring the profile is good to use. Prematurely moving the traffic to a higher modulation profile will also lead to traffic loss and the profile flapping behavior.

#### **5.4.2. CMTS Thresholds for CM-STATUS**

The CM-STATUS reporting mechanism includes a random holdoff prior to transmission of status report messages. This value is set on the CMTS and makes the CM dampen CM-STATUS messages. The Maximum Event Holdoff Timer indicates the value of that random holdoff timer to be used by the CM when determining when/whether to transmit a CM-STATUS message. This TLV associates a separate hold-off timer value with each CM-STATUS event type code managed by the CMTS.

A Maximum Reports Count Timer value controls how often repeated CM-STATUS messages for the same Transaction Identifier are sent by the CM. It controls how many CM-STATUS messages for the same Transaction Identifier are transmitted by the CM. A Maximum Reports Count of zero signals that the CM continues sending CM-STATUS messages as long as the event condition is "on" and is enabled for reporting. If the CMTS receives a CM-STATUS message from the CM, the CMTS transmits a CM-STATUS-ACK message with the received event type and transaction ID.

When the CMTS receives a CM-STATUS message, it needs to act to remedy the situation. The CMTS action is based on the received event in the CM-STATUS message.

As an example, on Profile Failures: the CMTS could change the profile for traffic, this will be a combination of DBC, OPT, (DPD) etc. If the PLC or Profile 0 failure occurs, the CMTS could reassign the CM receive channel set by removing the affected channel and replacing it with another OFDM channel if available.

### **5.5. DS Profile Flapping**

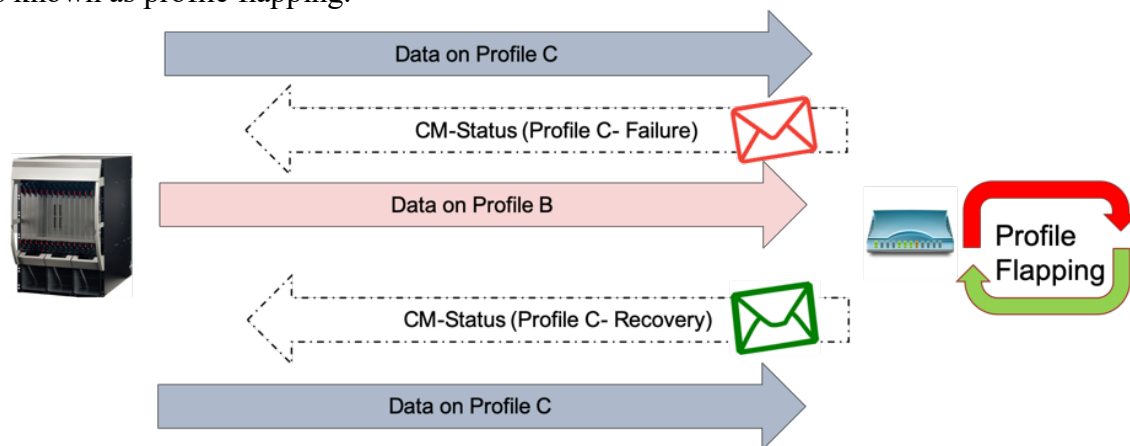
When the CM experiences any kind of intermittent noise on the OFDM channel, e.g. LTE, sweep generator, or other ingress noise, or if the CM is simply at the margin for a profile, the RxMER for the channel falls below the required level for the CM to decode a certain modulation order. The interleaving and FEC are unable to correct for noise and the CM will experience code word errors on the channel on a particular profile assigned to it. Let's say the CM is assigned 3 profiles: Profile A (256QAM), Profile-1024QAM, and Profile-4096QAM. All of these profiles are flat, i.e. every subcarrier has the same modulation order.

If the CM is currently operating on Profile-4096QAM, and it experiences code word errors, then as described above the CM will send a CM-STATUS message to the CMTS, informing the CMTS that it has a profile failure on Profile-4096QAM. Now the CMTS will start sending traffic on the next lower profile, in this case that is Profile-1024QAM. Now let's assume the noise interference is just enough that it can operate on Profile-1024QAM without code word errors. This switchover of profiles happens quite seamlessly on the CMTS. Now the CMTS is still sending traffic on Profile-4096QAM, to other CMs on the profile. So, this CM will continue to try and decode codewords on that profile, though the traffic is not destined to it.

Now if the noise ingress disappears quickly, say within 10 seconds or so, then the CM will be able to start successfully decoding packets on the Profile-4096QAM. If the CM is currently experiencing no more code word errors, (thresholds as described before), the CM will now send a CM-STATUS message to the CMTS, informing the CMTS that it has a profile recovery on Profile-4096QAM. Now the CMTS will move the traffic destined to this CM from Profile-1024QAM back to Profile-4096QAM.

If the noise ingress is strong enough, a CM could drop through multiple profiles all the way down to Profile A, and then back up all the way when the noise disappears.

Now this whole process of receiving traffic on a high profile and then to a lower profile and back is known as profile flapping.




**Figure 22 – D3.1 Profile Flapping Behavior**

In practice depending on the settings on the CMTS and the CM, of the CM-status messaging and of the hysteresis time which the CMTS applies before promoting the CM back to the highest profile, the CM could see significantly lower capacity. The least common denominator profile, Profile A could be 64QAM to accommodate the worst-case user and if others CMs also spent a lot of time on profile A due to the above profile flapping issue, the CM and the network see a much-reduced capacity.

the channel maybe intermittently impaired for some CMs leading to a cycle of profile failure and recovery messages. Each time the CM is forced to a lower profile, the CMTS as per the configured hysteresis settings does not recover the profile fast enough. Profile Downgrade is immediate while the profile upgrade is slow due to the hysteresis in the system reacting to the recovery.

If errors are seen on Profile A, then the consequences are more catastrophic as this means the channel itself is unusable, and the CMTS now considers the OFDM channel inoperable and the CM from partial channel, has moved to partial service as it has lost the use of the whole OFDM channel. Even though the CMTS reacts to both the profile failure and recovery messages, this profile downgrade - profile flapping/ partial service scenarios results in user traffic being dropped by the CMTS-CM. This starts manifesting itself to the user as intermittent connectivity issues and slow speeds in general.

The figure below explains the behavior visually using code word counters over time, as data moves across profiles, along with the timelines for CM-STATUS messages (not drawn to scale).



**Figure 23 – D3.1 Profile Flapping Example**

## 5.6. Suggested CM-STATUS Settings

On the CMTS the CM-STATUS Holdoff Timers need to be configured different from default as below for better performance.

**Table 3 – CMTS CM-STATUS Settings**

CM-STATUS Event	CM-STATUS Description	Setting (multiple of 20 ms)	Time (seconds)
2	QAM/FEC Failure	50	1
5	QAM/FEC Recovery	<b>500</b>	10
16	OFDM Profile Failure	<b>50</b>	1
24	OFDM Profile Recovery	<b>500</b>	10
20	NCP Profile Failure	100	2
22	NCP Profile Recovery	100	2
21	PLC Failure	100	2
23	PLC Recovery	100	2

CMTS Automatic Profile Recovery behavior is not specified in the DOCSIS standards and these are non-Standard CMTS Configurations, which need to be customized per CMTS.

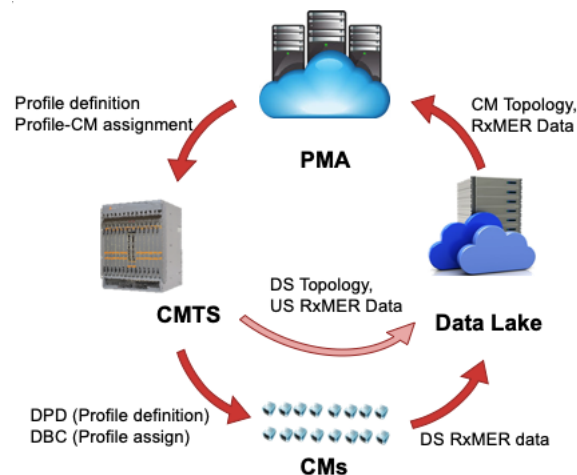
**Table 4 – CMTS Maximum reports and ACK**

CMTS setting	Value
CM-STATUS Maximum reports	5
CM-STATUS-ACK	ON
Vendor Proprietary Mechanism (Hysteresis for profile upgrades). Maybe known as profile guard time, Unfit time, recovery time etc.	5 mins or lower (2 Mins) if supported

## 6. Downstream Profile Management Application

DOCSIS 3.1 introduces the concept of modulation profiles or bit loading characteristics for each subcarrier within the OFDM/A channels. A modulation profile is a list of modulation orders or bit loading configurations, defined for each subcarrier within an OFDM channel, or for each minislot in an OFDMA channel. A CMTS can define multiple modulation profiles/IUCs for use on a channel, where the profiles differ in the modulation orders assigned to each subcarrier or minislot. A CMTS can assign different downstream and upstream modulation profiles for different groups of CMs.

As seen in Chapter 2, the interference patterns/RxMER on the different CMs on the same channel/plant are quite different. The best way to prevent profile flapping and enable robust operation on a D3.1 OFDM channel is to custom design the profiles for the CMs and to the state of the plant. Determining the best modulation profile to use on a channel is difficult, given the number of CMs and the differences in signal quality that they experience. PMA helps operators design the best modulation profiles for each channel, given the channel characteristics seen by each CM on the network.



**Figure 24 – Profile Management Application Deployment Architecture**

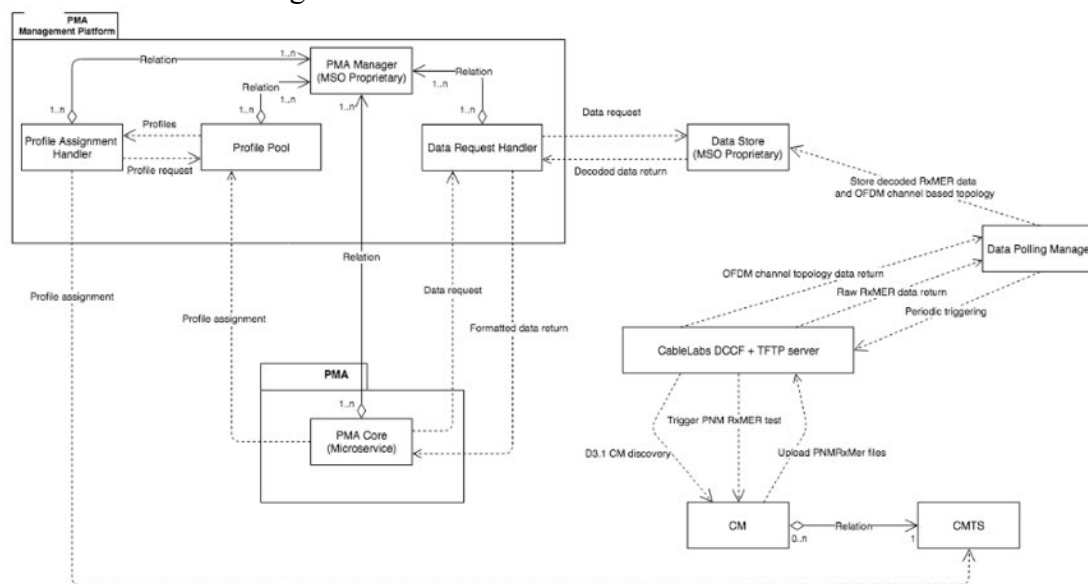
**PMA Goal:** The goal of designing profiles is to increase reliable operation and throughput per CM. PMA essentially consists of intelligent clustering algorithms to group CMs which share similar noise characteristics together: Groups of CMs get assigned a unique custom designed profile, which works around specific ingress issues etc.

The tasks an external PMA performs for both downstream and upstream profiles are as follows:

1. Create a set of optimized modulation profiles for use on an OFDM or OFDMA channel by selecting the best modulation order for each subcarrier based on the channel quality measured at the CMs/CMTS using the channel profile test or probes. (For all CMs)
2. For a new CM joining the network and periodically for all active CMs, find the best fit among existing modulation profiles and recommend modulation profile usage. (Per CM)
3. Create backup profiles or downgrade a CM based on errors on a certain profile. E.g. based on CM performance and SNR margin, provide a better modulation profile for a CM. (Per CM)

**PMA Benefits:** Well-designed profiles minimize codeword errors which lead to data loss and sub-optimal throughput. Designing profiles around noisy areas in the plant makes the system operation more robust. CMs can get downgraded/ upgraded on profiles when the noise is intermittent, PMA can prevent this profile flapping. PMA can maximize network capacity by optimizing the bit loading of every subcarrier. The bandwidth gains in running a well-designed set of profiles can be anywhere from 20% to 40% capacity increase on a channel, compared with running the whole channel at 256-QAM. This can translate to a solid 200 to 400 Mbps extra capacity on each OFDM channel. This enables an operator to match growing bandwidth demands and defer potential node-splits and new equipment costs. PMA is in full scale deployment with one large operator, and in field trial with another. 10 other operators are trialing out PMA in their labs (for Downstream and upstream) as they ramp up their D3.1 deployments.

### 6.1. DS PMA Software System architecture



A DOCSIS data collection framework is necessary to collect RxMER data from the CMs so that a PMA can create the appropriate profiles for the channel. DCCF, the DOCSIS common collection framework, developed by CableLabs along with industry partners is being used by

many operators and vendors. Here DCCF is used as an example of a data collector that can handle CM discovery, DOCSIS 3.1 PNM (& SNMP) data polling, and preprocessing for OFDM/A channel-based topology discovery, basic data storing and retrieving functionalities, and polling job management, etc. While DCCF is one of the frameworks for data collection from the cable plant, alternatives of DCCF can be built and chosen by vendors and operators.

In the architecture diagram, the following components perform the tasks described below.

- PMA Manager is the PMA Management Platform that handles profile assignment scheduling, policy input management, CMTS management, MER data redirection, and profile output management.
- PMA Core is the micro-service of the PMA core algorithm. It's a light-weight server that can handle multiple profile calculation requests efficiently.
- DCCF (or XCCF) with TFTP server (or MSO's own remote TFTP server) is a tool that will help MSOs setup and collect data. It attaches to one or multiple CMTSs and automatically discovers existing DOCSIS 3.1 CMs in operation. It's also able to discover the OFDM/A channel-based topology for all these DOCSIS 3.1 CMs, which is an essential function for PMA, which calculates profiles for each OFDM channel and needs the CM MAC address list to be hosted in CMTS\_IP/slot/port/channel fashion.
- Data Polling Manager periodically triggers DCCF to start a new round of OFDM DS RxMER data polling, and also triggers OFDM channel-based topology discovery on the CMTS. It could be an MSO specific implementation with an ability to pre-process the raw data from the data collector (DCCF in the system).
- Data Store is where the PMA gets the RxMER data and MAC addresses of CMs that are on each OFDM channel. It consists of two parts, a Data Lake and a thin service layer that serves data with restful APIs.
- Profile Translator is a shim layer that translates PMA's profile output to actual assignment information and automation process to CMTSs through available interfaces such as RESTCONF or CMTS CLI.
- Data Request Handler is the service that handles the PMA Core's data request and reply with data from the Data Store.
- Profile Pool is where original profile assignments are stored.
- Profile Assignment Handler handles actual profile assignment to the CMTS by requesting profile data from the Profile Pool.

The recommended functions of a Data Polling Manager are:

- PNM RxMER test triggering functions, including periodic triggering scheduler functions
- Manage DCCF API calls and translate DS RxMER data to a PMA understandable format
- Understand and translate returned OFDM channel-based topology from DCCF
- Perform highly efficient Data Store API calls for storing pre-processed data in place

The recommended functions of a Data Store are:

- The restful API service of the Data Store must comply the APIs defined by the PMA and must serve the data in formats that PMA can understand
- The Data Lake of the Data Store could be a database or an HDFS, per actual needs

PMA interacts with the Data Store on the link (data request and decoded data return) below using APIs standardized by CableLabs in the PMA YANG model [PMA-TR-CL]. The advantages of this model are that it is highly scalable. Request don't hit the network (CMTSs/CMs) directly when there's demand, and it can also serve plant data for other applications without querying the devices multiple times.




## 6.2. DS PMA Practical gains

We have collected and analyzed field RxMER data from D3.1 CMs from a few different cities from 3 different operators. This data includes approximately 142 unique OFDM channels with about 25000 D3.1 CMs spread across these D3.1 channels. In most cases, each channel had anywhere from 100 to 300 CMs (average of 174 CMs per channel), and in a few of the cases the channel had 40~80 CMs. We ignored channels which had less than 20 D3.1 CMs.


We design optimal profiles for the set of CMs on each channel, and then calculate the capacity gain for each of these channels, when using a set of profiles, as described in the paper on profile management algorithms [D31PMA-INTX16].

This capacity gain metric is essentially the capacity gain of a channel when using multiple profiles for unique sets of CMs and comparing that to all CMs using the single 256-QAM profile. The gain in capacity (J) increases as the number of profiles increases. For the sample channel below, one can see that once we get past ~7 or 8 profiles there is only minor incremental benefit to the overall capacity of the channel




**Figure 26 – Typical PMA Gain vs number of profiles (J-K Correlation)**

Using this method, we calculated the capacity gain for each channel in our field data set of 142 channels and for a range of 3-8 profiles. Below are the plots of the gains seen across this sample population of OFDM channels, shown as a histogram. Each histogram is for a certain number of profiles. As the number of profiles supported on the D3.1 CMTS increases, the gain of each channel and hence the gain of the collective of the channels increases.




**Figure 27 -PMA capacity gain histogram when using 3 profiles**




**Figure 28 -PMA capacity gain histogram when using 5 profiles**

Again, the PMA algorithm here has created robust profiles which work around the noise seen by each CM, and at the same time it also groups CMs to profiles so that the overall capacity of the channel is maximized.




**Figure 29 -PMA capacity gain histogram when using 8 profiles**

The next histogram below shows a comparison of the capacity gains when using 2 customized profiles on a channel vs 10 customized profiles on a channel. This comparison gives a visual of how the capacity gains for 2 profiles (in blue) shift to the right for 10 profiles (in light green). A higher number of well-designed profiles unlocks significant capacity gains across the plant.



**Figure 30 -PMA capacity gains, 2 profiles vs 10 profiles on a channel**

The next histogram below shows a comparison of the recommended number of profiles. When you add another profile to the system, the capacity gain(J) will increase. The recommended number of profiles for that channel is when that improvement in capacity is less than an incremental thresholds value. The below plot shows 3 histograms each with the recommended number of profiles with 3 different threshold values.



**Figure 31 -Recommended number of profiles**

As the increment threshold decreases from 1% to 0.5% to 0.25 %, then recommended number of profiles increases. For a minimum gain threshold of 1% majority of the recommended number of profiles is under 10. This is a good number of profiles for a D3.1 system to support, and many CMTS vendors are planning to increase the support for the number of DS profiles from ~3-4 to 8-10. The operators deploying D3.1 are coming to a consensus around this number of profiles.

# DOCSIS 3.1 Upstream


## 7. D3.1 OFDMA Upstream FEC behavior

The D3.1 equipment supports a minimum of two independently configurable OFDMA channels each occupying a spectrum of up to 95 MHz in the upstream. The systems support upstream transmissions from 5 to at least 204 MHz and agile placement of the OFDMA channels within that range.

### 7.1. Noise Characteristics on an OFDMA Channel


US spectrum is in general noisier and more susceptible to interference than the DS spectrum. We are only beginning to get US RxMER data from field CMTSs. A few examples of these are shown in the figures below.

Depending on the plant, many CMs have relatively clean RxMER levels across a 20-55MHz OFDMA channel. (In the figure below, we are seeing some MER reporting issues from the CMTS every 5 MHz or so).



**Figure 32 – Multiple measurements of D3.1 US RxMER from a CM**

In the figure below, we are seeing some changes in the measured RxMER when one CM is present on the channel vs when multiple other CMs are also using the same OFDM channel.



**Figure 33 –D3.1 US RxMER from a lab CM**

### 7.2. US FEC behavior

The choice of codeword sizes to be used in any given burst is based on the grant in the MAP message. The grant indicates which minislots are assigned to a given burst and which upstream profile is to be used. 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. Codewords are filled and transmitted in the following order, Full, Medium, Short, with codeword shortening when needed.

The ability of the system to support a given QAM level depends on the RxMER values and the mappings to an appropriate QAM level, when creating a profile. These mappings are defined in [PHYv3.1] and are summarized in the Table below.

**Table 5 - US RxMER to QAM Level mapping**

Upstream Constellation/ Bit Loading	US MER(dB)
QPSK	11.0
8 QAM	14.0
16 QAM	17.0
32 QAM	20.0
64 QAM	23.0
128 QAM	26.0
256 QAM	29.0
512 QAM	32.5
1024 QAM	35.5
2048 QAM	39.0
4096 QAM	43.0

### 7.3. Lab testing of US FEC behavior on D3.1 equipment

We wanted to understand the performance of the US FEC on D3.1 equipment, in an effort to understand at what points will an operator start seeing failures in the system. We are testing these with 1 CMTSs and 2 different CMs. We hope to increase the testing to include other CMTSs and CMs. The test was run on a 5-45 MHz OFDMA channel

#### 7.4. Baseline test (no noise)

This test checks the RxMER at the CM and ran downstream traffic from the CM to the CMTS:


1. Send traffic from the modem starting at 150 Mbps for 30 seconds.
2. Repeat for the US IUC/profile configured at each QAM level between 16QAM – 64QAM-256QAM. We will extend the testing to other QAM level as time permits.

In this baseline test there were a few corrected FEC codewords for all of the 3 profiles.

#### 7.5. Test to discover failure points, noise across entire channel

The next test determines at what points would the CMTS start seeing codeword errors. The idea here is to increase the (AWGN) noise floor on the channel and see how the system performs. For each modulation order: QAM 16, 64, 256 and each with 2 different pilot patterns, the task was to identify the average RxMER of the channel at which

- the first corrected/ uncorrected codeword is seen
- 100% corrected/uncorrected codewords are seen



**Figure 34 – Average D3.1 US RxMER levels for Correctable and Uncorrectables**

We identified the different corrected codeword failure levels at the CMTS, and the corresponding RxMER values tied to those failure point. As seen in the figure there is an inverted S-curve growth in corrected codewords as noise increases. As the number of corrected codewords reaches 100%, the number of uncorrectable codewords start increasing. As expected the lower the modulation order, the more noise is needed to get to the first corrected, 100% corrected and first uncorrectable codewords. In the upstream, we can see the number of uncorrectable codewords go up to high levels, as this CMTS has been configured to operate only on this single OFDMA US. In addition, the CMTS does not change the affected IUC at this time.

## 8. OFDMA Channel Configuration


### 8.1. US Channel Location

Though the supported upstream frequency range starts at 5 MHz on a D3.1 plant, the few European operators we have been collaborating with, have seen noise issues in the lower part of the spectrum. Rather than debug one too many things at once, the operators have made a choice to start their OFDMA channels at 20 MHz or 23 MHz and higher. The European operators have the luxury of a higher split at 65 MHz, so can still fit in a 40 MHz OFDMA channel. The channel also has excluded sub carriers for the return sweep generator. Also, typically the lower frequencies are excluded subcarriers or forced to be unused carriers.

### 8.2. TaFDM

DOCSIS 3.1 also supports simultaneous Time and Frequency Division Multiplexing (TaFDM) between SC-QAM and OFDMA channels. This means both OFDMA and SC-QAM can

simultaneously operate on the same frequencies, divided in time. This allows for the use of OFDMA across the entire spectrum, while maintaining backward compatibility with legacy DOCSIS SC-QAM channels. The figure below from [MULPIv3.1] provides an example of how TaFDM can operate with an OFDMA channel sharing the same spectrum as four SC-QAM channels.



**Figure 35 – Time and Frequency division multiplexing**

In some initial testing, when using the TaFDM feature between SC-QAM & OFDMA channels, it has been observed that the throughput varies depending on the set of channels in use. In one example configuration, there was one 40 MHz OFDMA channel and 4 SCQAM channels in between as per above diagram. The throughput for the time with OFDMA-only portion in use is the highest, the SCQAM-only is the lowest and when they are time sharing the throughput is somewhere in between. To achieve a good OFDMA throughput, consecutive spectrum is needed without the use of SC-QAM channels. The CMTS tends to schedule OFDMA traffic first in the OFDMA area before it schedules OFDMA bursts in the TaFDM area. If that area is in the lower part of the US spectrum, then the OFDMA also has to deal with the ingress noise which is more typical at the lower frequencies up to 20 MHz. Some operators have turned off TaFDM as that feature is not quite mature on the CMTS implementations and has not gone through enough CMTS-CM system debug.

### 8.3. Ranging location

On each CMTS, there needs to be a configuration for the ranging zone within the OFDMA channel. Initial field testing has shown that that moving this ranging zone above the 18-20 MHz range yields better stability for the D3.1 CMs. Moving the ranging zone to a cleaner and stable part of the spectrum ensures that the CMs come online and stay online on the OFDMA channel.



**Figure 36 – OFDMA Channel configuration example**

### 8.4. Minislots

Minislots are 8 sub carriers or 16 subcarriers wide. Number of symbols in time for an OFDMA frame(K) is in the range 6-36 symbols wide. When a single subcarrier is excluded, the CMTS needs to readjust the Minislot locations, as no excluded subcarriers are allowed within a minislot.

Exclusion-Bands or Zero-Modulated sub-carrier could be easier to configure in steps of whole minislots. At the time of testing earlier this year, some CMs were unable to handle complex UCDs. Along with the modulation order, the minislot also needs to be configured with the appropriate choice of pilot patterns. There is a choice of 7 pilot patterns (for each minislot size), among the pilot patterns tested, pattern 4 looked to be more robust and had relatively fewer lost traffic compared to pattern 1.

## **8.5. IUC / Profile management**

It is intended that the burst descriptor associated with the data profile IUC 13 be configured as a robust OFDMA profile usable by any DOCSIS 3.1 CM served by that upstream channel. The CMTS uses data profile IUC 13 for all OFDMA data grants to modems which have not completed registration. The CM transmits data using the OFDMA Burst Descriptor for IUC 13 prior to registration.

During or after modem registration, the CMTS has the option of assigning the modem to use any data profile specified in the UCD. Typically, the Burst Descriptors for data profiles other than IUC 13 will be configured for higher performance than IUC 13, although not all of these Burst Descriptors will be usable by all modems. IUC 13 is the lowest common denominator profile, it used by all CMs before registration and after registration for sending mac management messages. Data Profile IUCs (IUC # 5, 6, 9, 10, 11, 12, (and 13)) can use the following modulation orders BPSK, QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM.

Initial CMTS implementations have a restriction on the number of profiles available on a channel. This is somewhat restrictive. Some operators have started with 2 IUC profiles, as an example: IUC-13 (Default/Fallback): using 16-QAM, and an IUC-12 (Data): using a mix of 64-QAM in low frequency, 256-QAM in higher frequency area

## **8.6. OFDMA Profile Flapping**

After registration, the CMTS grants OFDMA bandwidth for data transmissions to a CM using one of the CM's assigned OFDMA Upstream Data Profile (OUDP) IUCs. The CMTS cannot grant data bandwidth to a CM using an IUC not specified as one of that CM's assigned OUDP IUCs. Upon successful completion of a transaction assigning one or two assigned OUDP IUCs to a CM, that CM needs to be ready for transmitting data using the assigned IUCs.

A CM supports 2 US Profiles/IUCs at a time. A CM starts on the OFDMA channel with IUC 13 (say for example 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) via DBC messages. The CMTS continues to use the default IUC-Profile 13 to forward traffic to avoid packet loss during IUC change, when the DBC is in process. In practice, this means is that the US capacity for the CM is changing intermittently quite significantly which leads to a degraded performance and user experience. We observed this in multiple MSO lab trials and also in our testing.



## **8.7. Upstream channel evaluation tools**

Because it is expected that not all upstream data profiles will be usable by all modems, a CMTS might wish to evaluate a modem's performance using a particular profile before assigning that profile to be used. There are two tools to aid the CMTS in gathering information about upstream profile performance: upstream probes, and upstream Data Profile Testing bursts.

A CMTS uses upstream probes for ranging-related functions such as determining transmit pre-equalizer coefficients and additionally using an upstream probe to take an RxMER measurement. To do this, the CMTS grants P-IEs in a P-MAP message with the "MER" bit set. When the CMTS receives the probe transmission corresponding to such a grant, it performs the RxMER measurement.

Some types of upstream profile performance measurements cannot be performed using probe bursts, like FEC performance or count CRC errors for a particular profile. Probe bursts cannot be used for these purposes since they carry no information. D3.1 systems support sending/receiving upstream Data Profile Testing bursts. The CMTS first assigns a Data Profile Testing SID to the modem on one or more upstream channels. (Transmit channel set encodings can be sent as part of a DBC transaction.) The CMTS then sends a grant to a Data Profile Testing SID. The CM responds to this grant by sending a Data Profile Testing burst in the grant.

## **8.8. US profile change CER based vs RxMER based handling**

There are a few different ways in which the CMTS can handle US profile issues for a CM.

The first method is using the codeword error rate, or the number of uncorrected codewords seen from a CM. In this case the CMTS knows that the IUC (modulation and pilot patterns) is not good enough for the CM to successfully transmit on the US. The CMTS can then chose to move the CM to a lower modulation/higher pilot pattern profile.

The second method is using the US RxMER, or the measured upstream RxMER seen from transmissions from a CM. In this case the CMTS based on the RxMER levels, makes a judgement based on thresholds that the signal is not good enough for the CM to successfully transmit on the US using that IUC. The CMTS then choses to assign an appropriate IUC/profile.

A third method is to combine the two metrics and decide on when the profiles need to change. As seen in this discussion the need for US RxMER is paramount, to design IUC/profiles and creating profiles across US channels across the plant. Currently US RxMER support on CMTSs is limited by accuracy issues and also not having any TFTP support to move the data to an external data lake.

## **9. US PMA**

Similar to the downstream, we want to understand how effective a custom designed profile/IUC will behave on the upstream and the kind of benefits customized profiles will give us.

## 9.1. Baseline test on 256-QAM

This test runs upstream traffic from the CM to the CMTS and checks the RxMER at the CMTS

1. Send traffic to modem starting at two different packet sizes and rates.
2. Repeat for different QAM orders.

We notice a few US packets being dropped, we are still working through the reasons why we have these packet drops even in the baseline case. The figure below shows the US RxMER measured at the CMTS and the traffic rates and number of packets sent and received.




Figure 37 –D3.1 US Baseline test

## 9.2. US Noise injection test

The test here was to determine behavior of single IUC/profile after noise is injected at specific location within the channel. A narrow band noise was injected (at widths of 0.2Mhz, 1 MHz, 5MHz into the US channel. As expected the packet drops increase for additional noise width. The figures below show the US RxMER measured at the CMTS, each of which show the noise ingress, and the traffic rates and number of packets sent and received.

For all of these tests, we are still continuing to analyze the relationship between packet sizes, packet rates, minislot sizes/ number of symbols in a frame etc. (Also having an SC-QAM US as a primary channel makes it a bit hard to tease out the distribution of load across the channels.)




Figure 38 –D3.1 US Noise Injection test 0.2 MHz






Figure 39 –D3.1 US Noise Injection test 1 MHz



**Figure 40 –D3.1 US Noise Injection test 5 MHz**

### 9.3. US PMA test

This test repeats the 5 MHz noise injection test, but this time with a customized profile, which works around the noise ingress. This profile allows the channel to operate at the same level in the presence of the ingress, as the system did without any ingress. Upstream profile design gives the same basic benefits as it does in the downstream, it has the ability to allow the channel to operate with uncorrectable codewords, which means a more stable usage of the upstream. The potential gains and benefits of the US PMA will be likely more than the Downstream, given that the upstream suffers from a lot more noise ingress and funneling effects from the plant. As we get more US field data from D3.1 CMs, we will gain a better understanding of the kind of profiles we need to create and the number of profiles needed etc.



**Figure 41 –D3.1 US Noise Injection test 5 MHz**

## Conclusion

A D3.1 network is a highly capable access network. The system features including the LDPC FEC, time and frequency interleaving and other signal processing enhancements have made a D3.1 network very robust to noise ingress. When the noise is severe custom modulation profiles using a Profile Management Application is absolutely need. The customized profile creation and configuration increases the reliability of the network (upstream and downstream) and in addition maximize the capacity. Support for 8-10 profiles on the downstream will be important in the years to come, as OFDM becomes the cornerstone of DOCSIS technology. MAC layer settings around CM-STATUS needs to be optimized to get the best traffic connectivity and experience for the customer. With increased size of the OFDM channels, the impact of optimizing them is also huge. Some good engineering is need in the configuration of the DS and US channels and their settings. Support for upstream signal quality data from a CMTS will become more important as operators start deploying more OFDMA.

## Abbreviations

bps	bits per second
CM	cable modem
CMTS	cable modem termination system
DOCSIS	Data over Cable System Interface specification
FEC	forward error correction
HFC	hybrid fiber-coax
Hz	hertz
ISBE	International Society of Broadband Experts
PMA	Profile Management Application
RxMER	receive modulation error ratio
SCTE	Society of Cable Telecommunications Engineers

## Bibliography & References

[PHYv3.1] DOCSIS 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I16-190121, January 21, 2019, Cable Television Laboratories, Inc

[MULPIv3.1] DOCSIS 3.1, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I17-190121, January 21, 2019, Cable Television Laboratories, Inc.

[PMA-TR-CL] DOCSIS 3.1 Profile Management Application Technical Report CM-TR-PMA-V01-180530

[D31PMA-INTX16] DOCSIS 3.1 Profile Management Application and Algorithms (2016), Greg White and Karthik Sundaresan, CableLabs. <https://www.nctatechnicalpapers.com/Paper/2016/2016-docsis-3-1-profile-management-application-and-algorithms>

[LTFreq] [https://en.wikipedia.org/wiki/LTE\\_frequency\\_bands](https://en.wikipedia.org/wiki/LTE_frequency_bands)

## Acknowledgements

*Thanks to different operators across the world for sharing their DOCSIS 3.1 data with CableLabs. Thanks to Comcast, Shaw, Videotron, Vodafone Germany, NOS Portugal, for graciously sharing their data, and involving CableLabs in their D3.1 trials and deployment efforts. Thanks to Dan Rice, for inviting CableLabs to be part of the Comcast efforts for D3.1 testing and analysis, and thanks to Ray Hammer on running countless tests to understand the CM-STATUS behavior and thanks to Paul Schauer for collecting and sharing data. Thanks to Nader Foroughi, Shaw, for involving CableLabs with their PMA trials. Thanks to Peter Wittman at Vodafone and thanks to João Fernandes at NOS for involving CableLabs with upstream OFDMA testing.*

# **Containers – To Use It or Not to Use It**

## **Understanding New Technologies to Evaluate Need for Containers**

A Technical Paper prepared for SCTE•ISBE by

Avinash Raghavendra  
Solution Architecture, NFVi and Cloud  
Ericsson North America  
6300 Legacy Dr, Plano, TX 75024  
469-266-3077  
avinash.raghavendra@ericsson.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
Conclusion .....	8
Abbreviations.....	9
Bibliography & References .....	9

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Deployment Model .....	4
Figure 2 – Container Model Expanded .....	5
Figure 3 – Virtual Model Expanded .....	7
Figure 4 – Evaluation Flow Chart .....	8

# Introduction

With the proliferation of Cloud Computing into all major segments of the society, be it private, public or hybrid, Network Function and Application owners now have a multitude of choices to deploy their Network Functions/Applications. Previously it was deployed directly on a bare metal, now the applications can be virtualized or containerized.

One of the most popular and rapidly growing technology is Containers. The containers with their cloud friendly features and functionality have rapidly made strides to be in the forefront of the next big wave of deployments be it Edge, 5G or any Network Function.

But is it a panacea for all that ails a virtual or physical world? In this paper, we shall look to answer that important question by evaluating the needs of the network function to identify the best technology fit, be it a Virtual Machine or a Container for the Network Function/Application.

There are many different aspects that are considered as part of this evaluation. These include Infrastructure, Networking, Application Design & Deployment. Other Operational considerations including scalability, availability, monitoring, and life cycle management also play a very important role in the evaluation. Finally, one must not ignore the security considerations that play an important role in this evaluation.

## Content

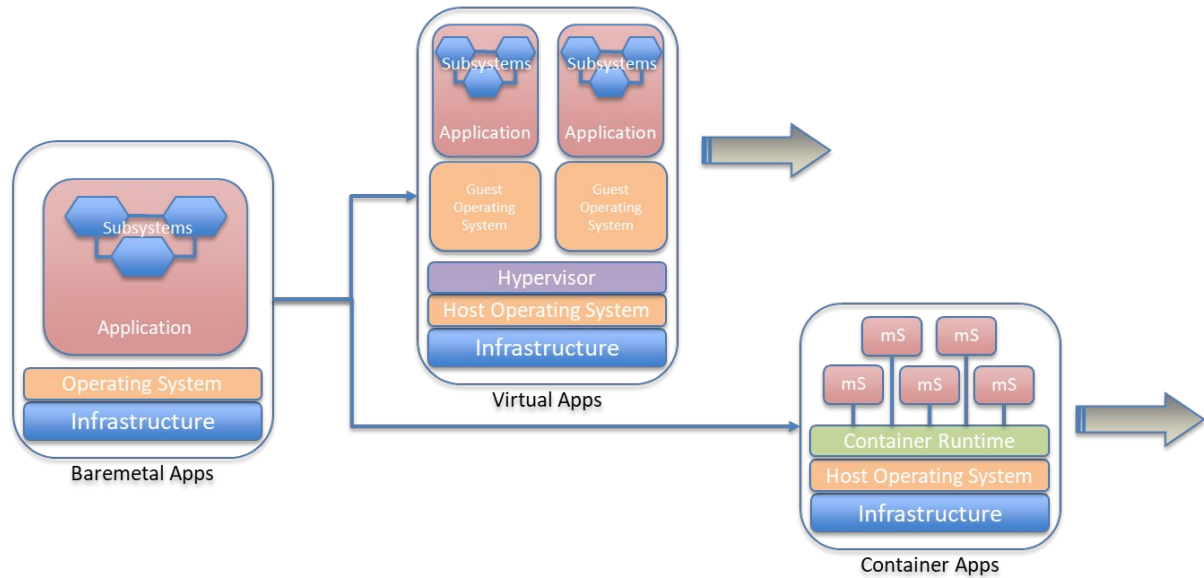
The dawn of general-purpose computing spurred the application revolution. Application were developed and over time optimized to run on various compute architectures. As applications became more specific in purpose, the resource utilization climbed higher. The resource utilization also became harder to predict for application placement. In medium and large enterprises, this brought an unexpected result, underutilization of compute resources.

The first step taken by the industry to solve this problem was virtualization. Applications were either uplifted or rewritten to be virtualized. Virtualization meant the applications had well defined boundary and were clearly quantifiable at definition and at runtime. This method caught on rapidly and with more adoption we started to see the benefit of virtualization, which were better utilization of hardware and network while providing better orchestration of deployment and life-cycle.

Specialized applications that were CPU, RAM and Network bound were virtualized quickly, saw the limitations of virtualizations. The limitation was of wide range, from virtual layer overhead for CPU and Network to more static deployments. These issues resulted in development of customization for dedication of CPU, RAM and Network to virtual workloads which in turn caused more static deployments. Static deployments resulted again in underutilized compute resources and islands in workloads.

Solving these issues meant bringing the workload of application back to bare metals. Containers provide a packaging for application with access and resource reservation while sharing the underlying kernel. Containers eliminates the need to have guest operating system there by removing the overhead caused by virtualization. By contrast, containers require container runtimes that does not add to the

continuous bulk overhead like in virtualization but just facilitate the initiation of the container and lightweight runtime linkage to the host OS (Operating System) and its common libraries.



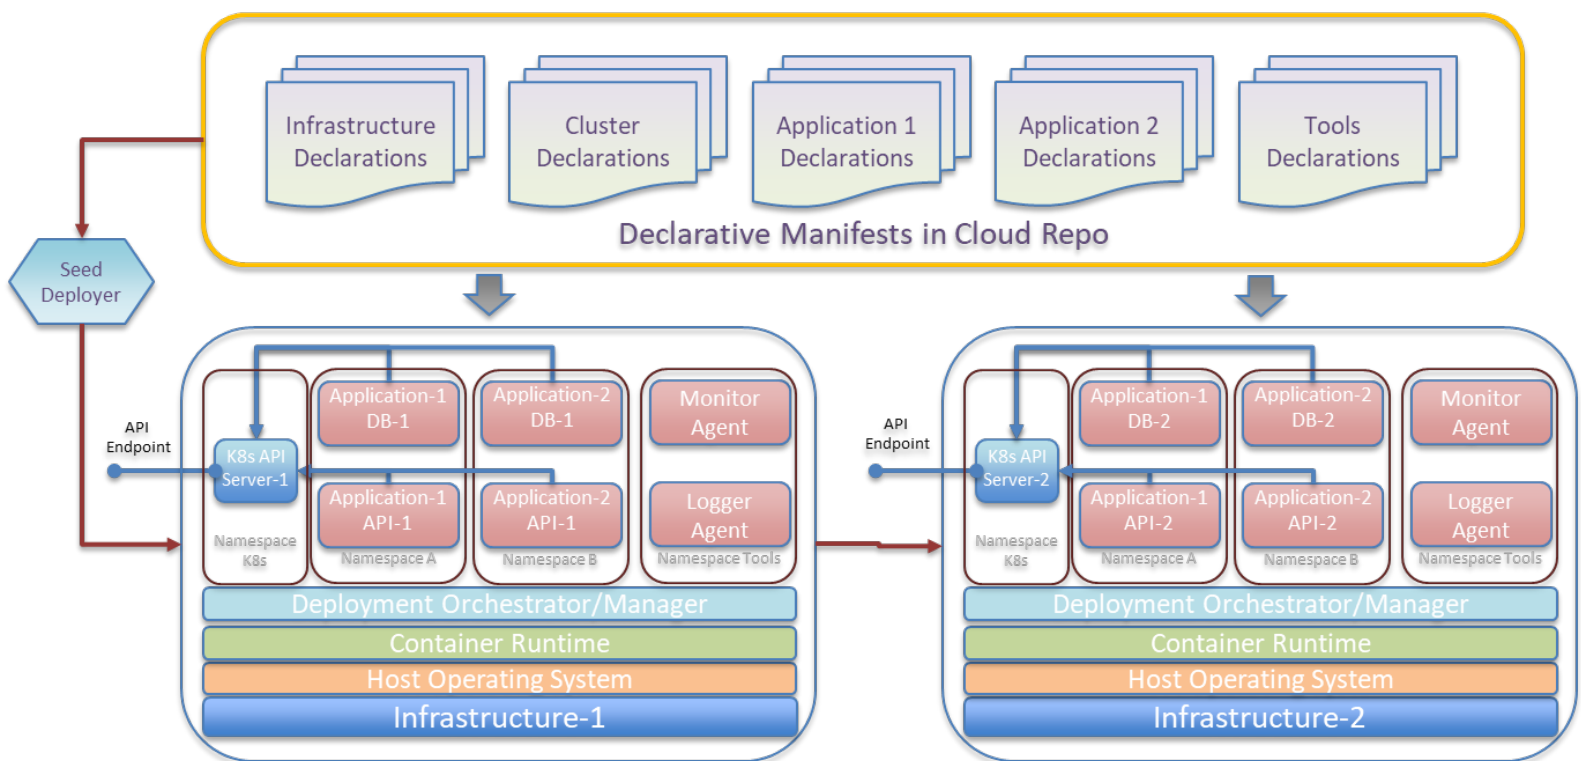
**Figure 1 – Deployment Model**

## A Case for Containers

A declarative model is where the definition and configurations of an application is specified ahead of its deployment in a structured way. An application that is modeled on declarative way is the best fit to move to containers.

Historically applications have been designed and built to be monolithic, but this is changing. Application that are designed to be its smallest unit of work that interact using REST (Representational State Transfer) interfaces are called microservices. Microservices isolate a single application and its dependencies, all the external software libraries the app requires to run both from the underlying operating system and from other containers. Applications that are designed as microservices are best suited to run and managed as containers. Another benefit of microservices are that they are extremely portable since all its dependencies are packages together without the common libraries and operating system. Applications that are microservices should also callout their individual CPU, memory and storage requirements.





**Figure 2 – Container Model Expanded**

Existing and planned infrastructure also plays important role in deciding for containers. Containers work well on small but distributed deployments of infrastructure. This is because microservices that interact closely to deliver a solution are usually deployed with multiple instances over a geo-redundant infrastructure to load balance and to provide resiliency. The benefit of such deployments is on-demand scalability, when the need for those services increases, more application instances can be quickly deployed to meet the demand just at the required region and later scaled down.

The benefit of the portability of well-defined and de-coupled containers is high resiliency and scalability. Since it is well defined, it can be automated for runtime deployment. With successful and repeatable automation, multiple instance of the same piece of software can be quickly deployed on demand, geographically with little change.

Networking is another aspect that requires consideration for container technology. Containers work well with simple networking layout and standard network protocols. Container, though in use in IT enterprises for many years are still nascent in advanced networking. New developments are being made at fast pace, but it is best suited for application that need or have requirements with simple networking model TCP(Transmission Control Protocol), HTTP (HyperText Transfer Protocol)and HTTPS (HTTP Secure) and technologies like OVS (Open vSwitch), SR-IOV (Single Root I/O Virtualization) and DPDK (Data Plane Development Kit).

Life-cycle management and Operation is critical for any application. The operations and the team behind it must be very well versed on container technologies like Docker, Containerd, Kubernetes etc. to take advantage of the benefits of containers. DevOPS model is where the development and

operations are part of the same process chain that seamlessly delivers and operates the application and solutions. Application as containers requires their features and capabilities to be developed and deployed iteratively. This combined with multiple instances of the same microservice means little or no impact when new features are delivered. This methodology is called rolling deployment.

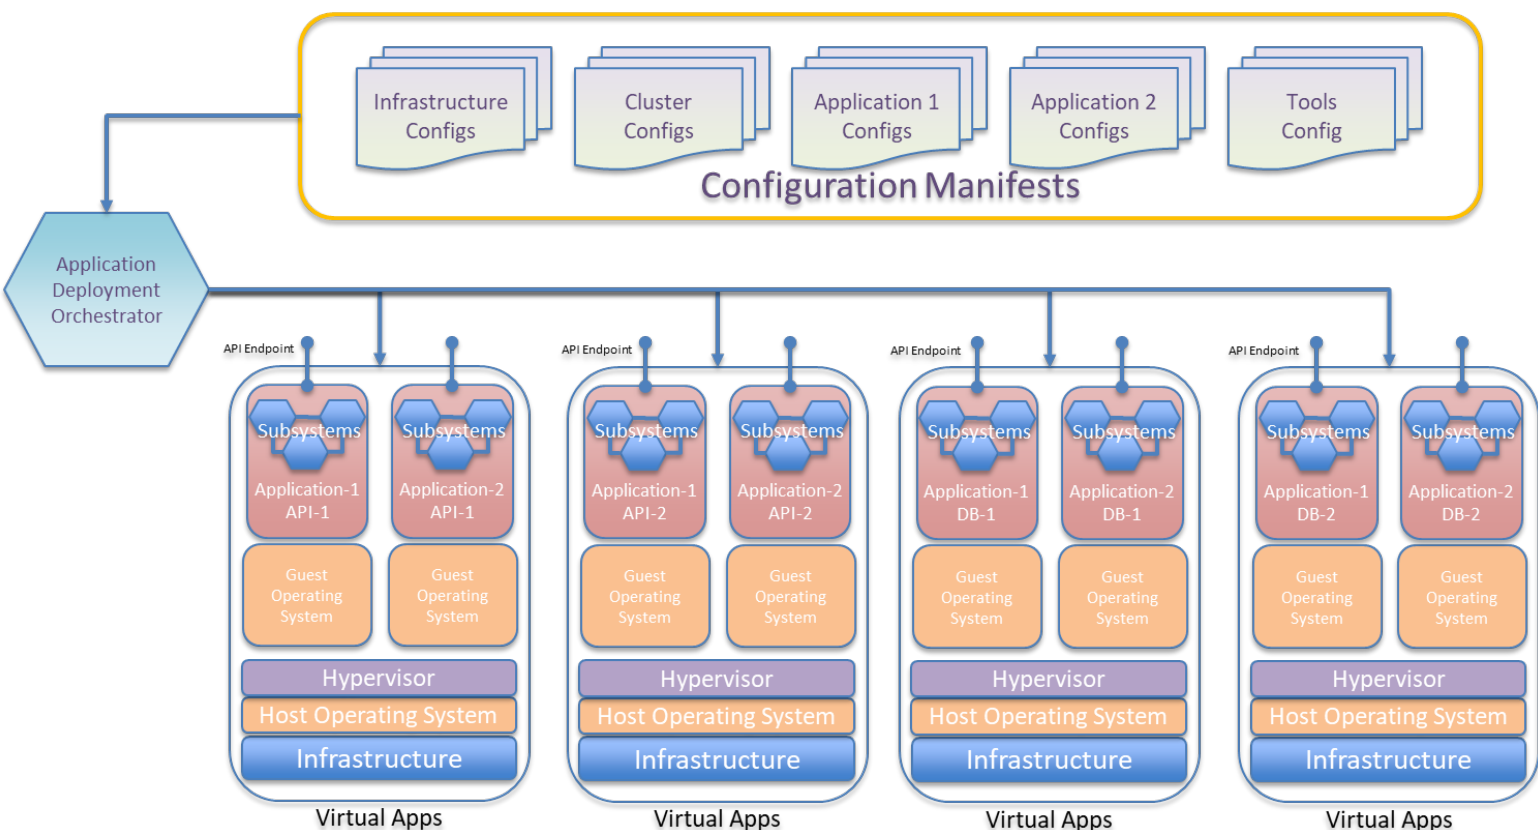
Monitoring and alerting mechanisms will change with container as well. It is no more a single entity alerting on its various subsystems. The monitoring is now on each microservice that consumes the resources independently to operate while sharing the common infrastructure. This means the monitoring thresholds must be declared again and configured at deploy time. The monitoring software must also understand and differentiate each unit of software and their dependencies on each other. Clear visualization of alerts and dependencies helps in better life cycle management of the whole solution.

Each container or microservice logs independently. This means the log capturing mechanisms must now be capable of handling the several pieces of logs, roll and archive them. It should also provide mechanisms to search and visualize the various logs that are generated by each microservice.

## **A Case against Containers**

Security is the most important yet least understood components of applications. Applications that handle single channel security policy enforcements expect policies are applied to the front-end services while the rest of the internal applications are hidden behind façade. This is true in case of virtualized workload where frontend application that accept ingress traffic are usually secured with TLS and firewall etc. while internal applications like database etc. do not need such policies. Such application is a least fit for container or microservices methodology. Each container deployed on an infrastructure introduces a small but definite surface area that can be compromised or used to compromise the infrastructure. The DevOPS (Development-Operations) model that is not equipped to scan all container images that need to be deployed also weakens the security of the solution. In a virtual world, a privilege of an application is contained within a virtual machine, so a rouge application at worst can bring down a VM (Virtual Machine), but a rouge container can now access the bare metal and compromise the whole infrastructure or the solution.

The design and architecture of software and application is the foundational in deciding deployment and runtime model used. An application that is designed as top down with relatively full analysis requirements sees less benefit from the container model. Containers benefit from iterative introduction of small features seamlessly rather than massive changes that disrupt the solution. Not all applications perform better in container model, applications that are tightly coupled to provide a specific service in efficient way may not be suitable to be broken down in to containers. Such applications are better suited either as virtual machine or application specific hardware.



**Figure 3 – Virtual Model Expanded**

The CI (Continuous Integration)/CD (Continuous Deployment) model for traditional software tend to look at application as a black box that has overall input code and overall output application and test based on that. Such models cannot fit in container model as each container now represents a unit of software providing its services decoupled from other containers. This means the CI/CD model should not cater to individual container, its code and its dependencies and package it, scan it and test it based on the container specific interfaces. The CI/CD model also must tie all the individual containers together in the sequence represented by the solution to test the end-to-end flow and be able to deploy it. The conversion of traditional CI/CD model to true DevOPS model is in most cases not feasible if the traditional model is already successful.

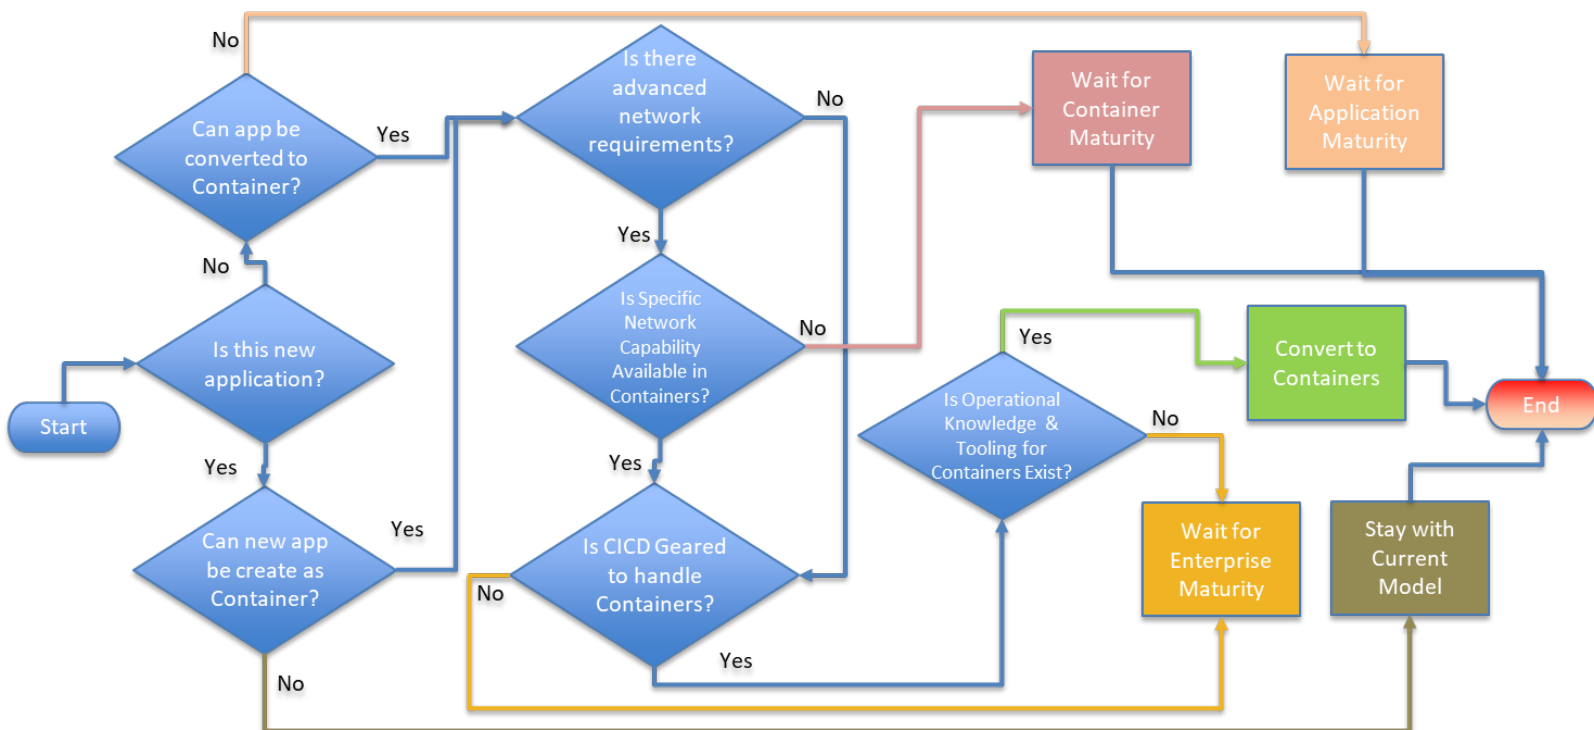
Existing operational and support model plays a crucial role too. Enterprises and operators who have a working and successful operational tooling and support process in traditional software and application either on bare metal or as virtualized would experience an immense surface area to cater in a container model. Each container requires individual logging, monitoring and altering capabilities and appropriate seamless action or pro-action from the operations. Monolithic applications that define, operate, monitor and alert its subsystems efficiently does not see a benefit by moving to container model, on the contrary the software might become more complex and less efficient. This is also true for application that does not rely on declarative model.

The birth of containers was in IT and Enterprises which had little or no specializations in terms of workloads like NUMA (Non-Uniform Memory Access) awareness, usage of advanced networking like SR-IOV and DPDK. So, application that do not demand specializations work well in the container model. On the other side most TELCOs and Content Providers have applications that requires very specific technologies to function and meet the required efficiency and performance. It would be better for such applications to wait till advanced technologies are integrated, evaluated and mature before adoption.

## Conclusion

Containers and its ecosystem offer a tremendous advantage to runtime capabilities and efficient utilization of underlying infrastructure. Every person or organization planning to create a piece of software must look at these emerging technologies and evaluate themselves to adopt to it. Operators of existing software, application and deployments must consider ways to incrementally add tooling, knowledge and process to convert to container model if their evaluations can conform container model.

A high-level flow chart represented in figure 4 can be used to evaluate a specific application based on factors discussed in this paper.



**Figure 4 – Evaluation Flow Chart**

## Abbreviations

NFV	Network Function Virtualization
NFVi	NFV Infrastructure
Edge	Closest to consumer
NUMA	Non-Uniform Memory Access
REST	Representational State Transfer
TCP	Transmission Control Protocol
HTTP	HyperText Transfer Protocol
HTTPS	HTTP Secure
OVS	Open vSwitch
SR-IOV	Single Root I/O Virtualization
DPDK	Data Plane Development Kit
DevOPS	Development-Operations
TLS	Transport Layer Security
VM	Virtual Machines
CI	Continuous Integration
CD	Continuous Deployment
API	Application Programming Interface
E2E	End-To-End
OSS	Operations Support System
VNF	Virtual Network Function
QoS	Quality of Service

## Bibliography & References

OCI: Open Container Initiative <https://www.opencontainers.org/>

Kubernetes Projects in Special Interest Groups: <https://github.com/kubernetes-sigs>

Kubernetes: <https://kubernetes.io/>

Docker: <https://www.docker.com/>

Containerd: <https://containerd.io>

# The Evolution Of Cellular IoT

A Technical Paper Prepared for SCTE•ISBE by

**Hani Beshara**

Director, MANA Radio Network Solutions  
Ericsson Inc.  
6300 Legacy Drive, Plano, Texas 75024  
(214) 906-8231  
hani.beshara@ericsson.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	4
1. Overview of cellular IoT market segments .....	4
1.1. Massive IoT segment.....	4
1.2. Broadband IoT segment .....	6
1.3. Critical IoT segment.....	7
1.4. Industrial Automation IoT segment.....	8
2. IoT Use cases – additional technology considerations .....	9
2.1. Security .....	9
2.2. Identity .....	10
2.3. Machine Intelligence .....	11
Conclusion .....	12
Abbreviations.....	12
Bibliography & References .....	13

## List of Figures


<b>Title</b>	<b>Page Number</b>
Figure 1 Ericsson Mobility Report – June 2019 .....	3
Figure 2 Cellular IoT Segments .....	4
Figure 3 EAP-TLS ID management.....	10

# Introduction

The IoT market, from networks to devices to use cases and applications, is evolving at a record speed. It is set to unleash a major transformation that has not been seen since the industrial revolution. IoT promises to change how we live and interact with the world around us, and to transform business productivity measures. It will help to finally unwire the factory and to provide the next level of automation such that industries can operate more efficiently and offer optimal products & services.

Cellular IoT is not a new concept. It is widely adopted across the globe, with 2G and 3G networks providing Low Powered Wide Area Networks (LPWAN) connectivity enabling many early IoT applications. As per Ericsson's mobility report, approximately 400 million 2G cellular connected devices have been in operation since early 2016. 4G LTE helped provide greater bandwidth, lower latency and increased support for large volumes of IoT devices per cell. By the end of 2018, the number of connections reached 1 Billion cellular connection. These will be enhanced further with the arrival of 5G networks, initially enabled by the 5G New Radio (NR) standard, which will enable Ultra-Reliable Low Latency Communications (URLLC) that support increasingly critical applications.

The remarkable cellular IoT growth rate is expected to continue, and the number of devices connected by Massive IoT and other emerging cellular technologies is forecast to reach 4.1 billion by 2024.



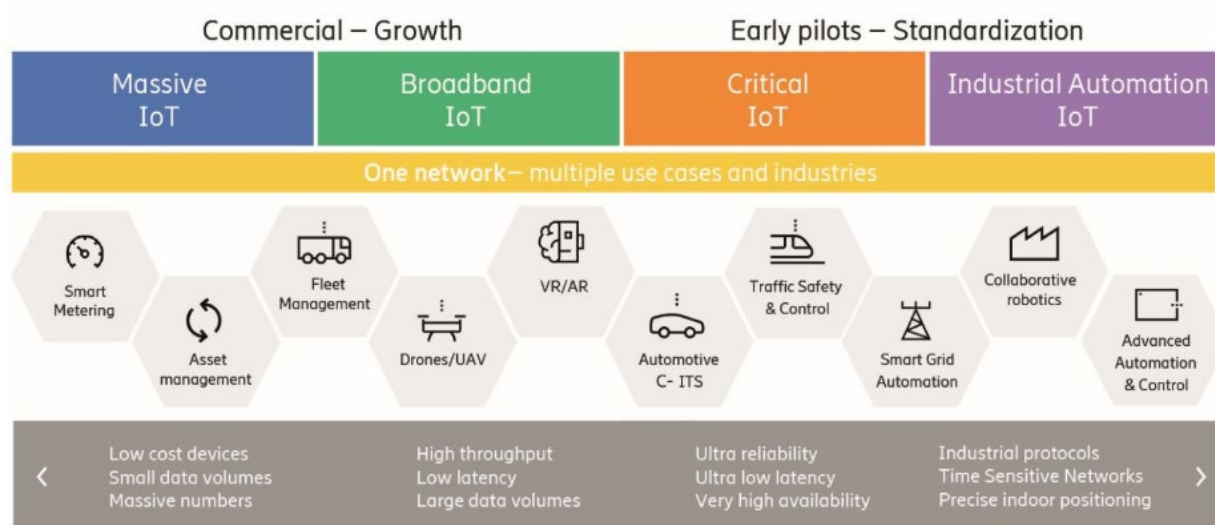
**Figure 1 Ericsson Mobility Report – June 2019**

Cellular IoT is supported by a rapidly growing ecosystem based on 3GPP global standards, offered by an increasing number of mobile network providers, and developed by a large number of device, chipset, module, and network infrastructure vendors. It offers unmatched global coverage in virtually every country in the world, Quality of Service, scalability, security and the flexibility to handle a varying set of requirements for a comprehensive range of use cases.



Four cellular IoT market segments are emerging with different requirements. These are:

- Massive IoT: low cost devices, small data volumes, massive numbers
- Broadband IoT: high throughput, low latency, larger data volumes
- Critical IoT: ultra-reliability, ultra-low latency, very high availability
- Industrial IoT: industrial protocols, time sensitive networks, precise indoor positioning



**Figure 2 Cellular IoT Segments**

This paper will review the emerging segments and the current and future support for the specific requirements within those segments through LTE and NR wireless technologies.

## Content

### 1. Overview of cellular IoT market segments

Massive IoT is a commercial cellular IoT market segment that is enabled by the 3GPP Release 13 standards which introduced the low complexity IoT variants; LTE-M and NB-IoT. As of early 2019, there are 80+ massive IoT networks in operation around the world. The Massive IoT segment has been growing rapidly with a compounded annual growth of 27 percent expected between 2018 and 2024 (Ericsson Mobility Report).

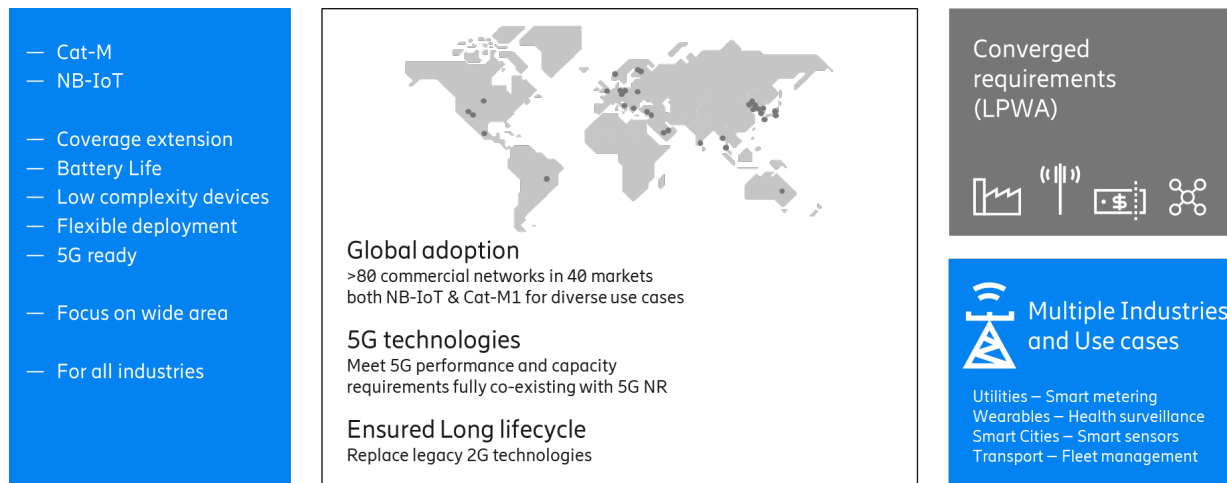
The Broadband IoT cellular segment refers to the existing IoT use cases that are enabled using 4G LTE, as well as the future use cases that would require 5G NR high throughput networks.

Critical IoT and Industrial IoT are emerging segments that will be made possible with enhancements associated with 5G NR and 3GPP releases 14/15/16+ by enabling Ultra reliable low latency communications (URLLC) and industrial protocols used for manufacturing.

#### 1.1. Massive IoT segment

Massive IoT connectivity targets huge volume of low-complexity devices that infrequently send or receive small messages. The traffic is tolerant of delay and typical use cases include low-cost sensors,

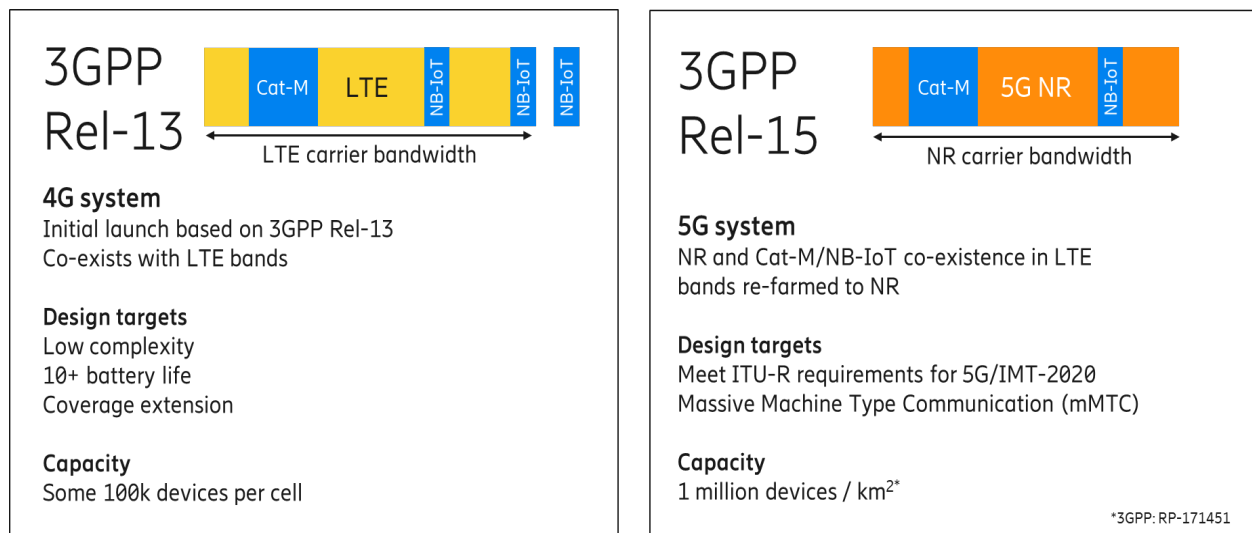
meters, and trackers. Such devices are often deployed in challenging radio conditions such as in basement of a building. Therefore, they require extended coverage and may rely solely on a battery power supply which puts extreme requirements on the device's battery life.



To support these attributes, new technology variants have been developed in the 3GPP standards Release 13; LTE-M and NB-IoT. Of the two, LTE-M (or CAT-M) supports greater bandwidth and complexity and is suited for wearables, trackers, alarm panels and customer support buttons, and hundreds of other use cases. LTE-M can support a VoLTE voice connection in addition to data, and supports bandwidth up to 1.4 Mhz.

NB-IOT is a stand-alone radio access technology that is designed for ultra-low complexity devices. NB-IoT only devices cost below \$5 and support a 180 KHz bandwidth (similar to GSM technology). NB-IoT does not support voice, but supports superior coverage extension, up to 100 Kms. As GSM networks are sunset, it is expected that its IoT traffic will migrate to NB-IoT.

LTE-M and NB-IoT are designed as future proof technologies that will coexist with 5G network as 4G cellular networks, either evolve to, or are totally replaced by 5G networks. Both massive IoT technologies are being added to 3GPP Releases 15/16.



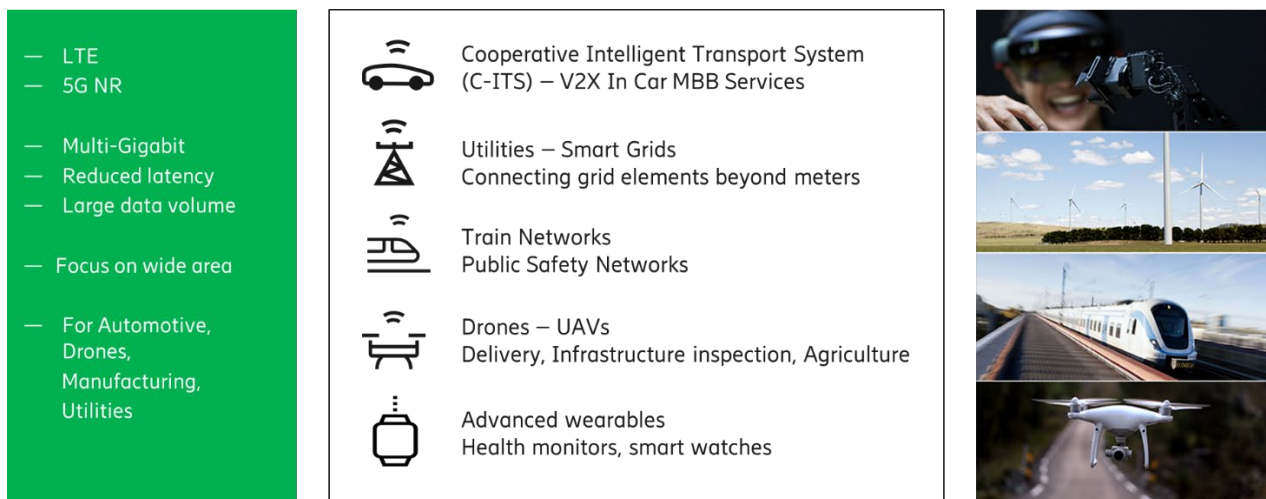
3GPP IoT technology uses licensed spectrum comprising of pre-existing LTE bands which provides superior capacity and channel management, lower inter-device interference, reliability, and QoS. 3GPP service providers can provide scalable IoT services to their customer to meet their IoT device requirements in any geographical area.

Also in the context of LPWAN, other competing technologies, such as SigFox and LoRa, provide data service for IoT devices using the unlicensed spectrum. Such technologies enable new low cost networks and devices, however, they could suffer from added interference from other technologies using the same frequency resulting in limited capacity, limited bandwidth, and is usually under strict government regulations for both Tx power and the duty cycle. In Europe, SigFox and Lora regulations limit device transmission power to 25mW (14dBm) and the duty cycle of channel access to 0.1 percent or 1 percent of channel time. The unlicensed technology is expected to be used for the low cost, low bandwidth, and low reliability use cases. For a detailed comparison between licensed cellular IoT technologies and unlicensed IoT options, please refer to 5G America's white paper "LTE and 5G Technologies Enabling the Internet of Things".

## 1.2. Broadband IoT segment

Broadband IoT connectivity use cases require superior performance with low latency and high throughput. Typical applications include advanced wearables, aerial and ground vehicles, AR/VR enabled devices and sensors that require greater capabilities than what CAT-M or NB-IoT can provide. LTE has a range of device categories well-suited for such applications including a lower cost CAT-0 device. For example, LTE is already providing cellular connectivity to millions of modern cars for infotainment as well as preventive maintenance. There are LTE capable smart watches in the market today and LTE-connected drones are in the proof of concept stage. In support of this IoT market segment, LTE offers high spectral efficiency and data rates, low latencies and has been enhanced with extended device battery life and improved coverage. With advanced multi-antenna solutions and carrier aggregation, LTE enables peak rates in excess of 1 Gbps. Added to this, there are mechanisms for fast connection establishment and data delivery. With instant transmission schemes, the radio interface latency can be as low as 10ms. LTE scheduler can also support advanced priority handling mechanisms to provide superior performance to a selected group of users.

With 5G NR, the segment will expand with even greater bandwidth availability and lower latency of ~5ms over the air that can be suited for AR/VR and connected vehicle type applications.



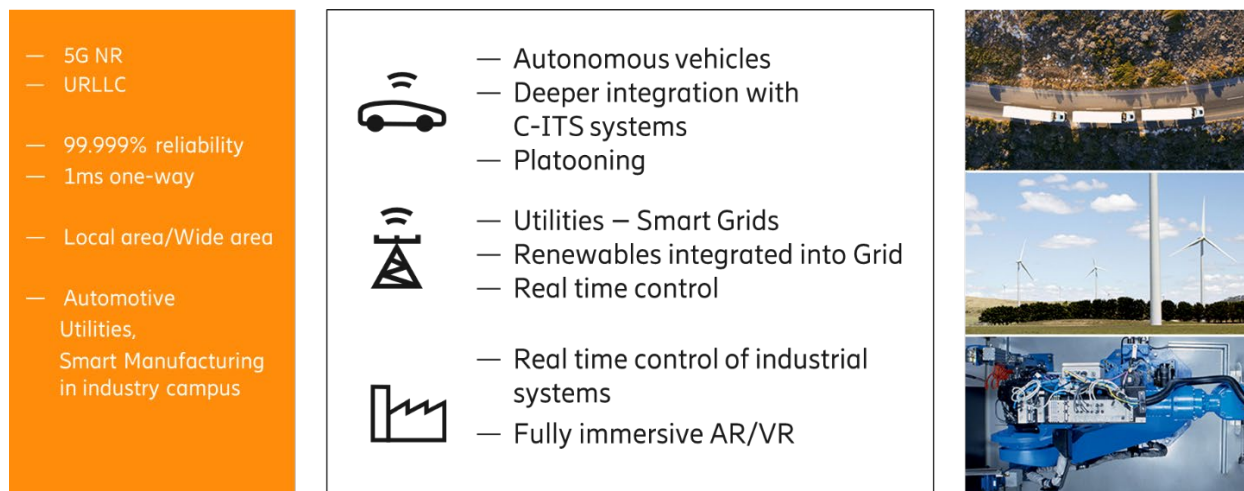
### 1.3. Critical IoT segment

Ultra Reliable Low Latency Communications (URLLC) is a particular focus of the 5G NR standards in releases 15, 16, and 17 in support of the critical IoT market segment. Critical IoT use cases require extreme reliability with down time of ( $10^{-5}$  to  $10^{-9}$ ) mins per year and low latency (1ms and lower) over the air. Examples include smart grid control, fault restoration, real time control of machinery, intelligent transportation systems, remote surgery, and fully immersive AR/VR.

Early releases of 5G NR reduce latency and increases reliability over LTE. NR spectrum with flexible numerology supports 15 KHz to 120 KHz subcarrier spacing which will in turn reduce frame alignment delay and help reduce NR latency up to 1/8 msecs. Also new features, such as, Instantaneous DL assignment & UL dynamic scheduling, ultra-short duration transmission have been standardized to achieve the ultra low latency requirement. From redundancy perspective, 3GPP releases 16 and 17 are adding features for increased reliability including multi-connectivity, diversity, Robust coding and modulation, rapid retransmission protocols, prioritization mechanisms, and multiple signal transmission formats to achieve a reliability goal for 99.99999% and higher.

Critical IoT segment is expected to also take advantage of end to end network slicing and 5G NR QoS template to provide industries and private networks with its performance guarantees.

While wireless network price per bit has been on the decline, URLLC use cases is expected to have a much higher cost and price level.

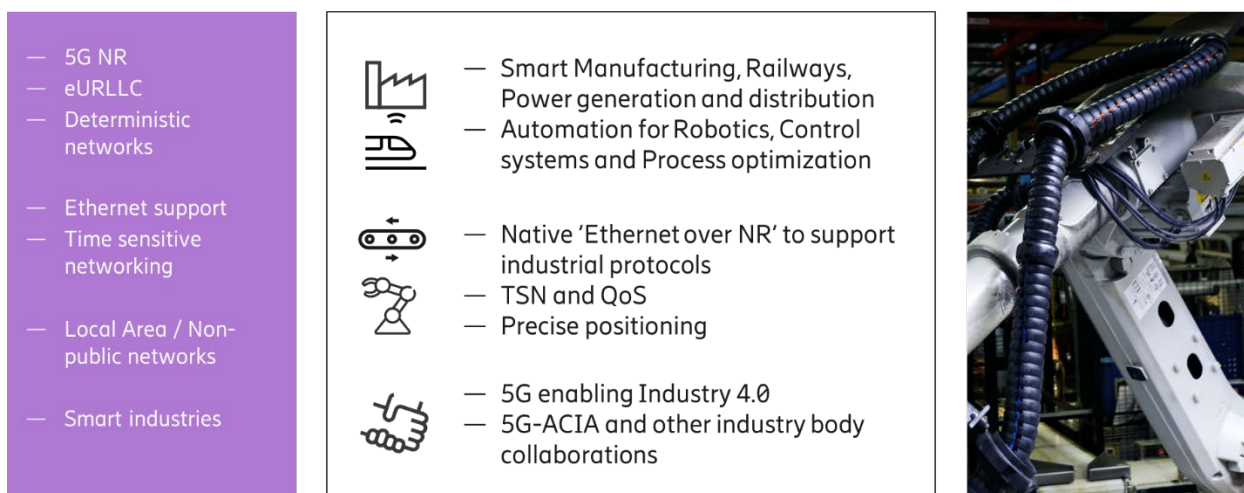


## 1.4. Industrial Automation IoT segment

Industrial Automation IoT builds on the capabilities of the three previous segments to address the specific needs of manufacturing and control systems for railways, power generation and power distribution.

This segment tailors the most demanding requirements from Manufacturing and Industrial Campuses and will support deterministic (and time sensitive) networks, industrial protocols such as PROFINET, EtherCAT, POWERLINK, and Modbus TCP/IP natively running over ethernet, together with very precise positioning.

The functionality and definitions for this segment are currently being defined in 3GPP, heavily influenced by Industry 4.0, and key industrial players and bodies such as 5G ACIA. It will be a 5G specific segment valid for local area and, almost certainly, non-public network deployments.



The ambition of this segment is to take that final step in the digitalization of the factory and move from today's mix of wired and wireless connectivity using multiple technologies into a wireless 5G network which will consolidate all requirements.

To fully support this segment, the standard must merge 5G NR and industrial control domains. 3GPP standards work is ongoing to enable precise indoor positioning, native support for Ethernet over NR, scheduling and QoS adaptations that will enable transparent Time Sensitive Networking (TSN) and allow reuse of existing industrial devices and control systems. With these additional capabilities, 5G NR have the potential to become the one common network that can support all existing use cases as well as fill the unfulfilled gaps in industrial communication and control networks.

## **2. IoT Use cases – additional technology considerations**

In addition to the IoT connectivity aspects discussed above, there are few additional key considerations that are important for the success of all IoT use cases in all segments and for managing IoT devices over their complete life cycle; Security, Identity and Machine Intelligence.

### **2.1. Security**

While securing a low price and a low complexity sensor may not seem as a critical activity, it has been proven that it can act as a trojan horse that opens a gate to the underlying network. Hence, IoT security needs to be taken seriously, even in what may seem as a simple application. Generally, Security of an IoT use case is a function of using secure communication, application security, and device security. Together, these functions protect the device, the network, guarantee data ownership, and ensure that trusted devices remain secure throughout their entire operational life.

Communication protocols like TLS, DTLS, and OSCORE are meant for IoT devices communications, however, not all supported algorithms are equally secure. Low complexity devices also have limited memory and processing power, hence it is even more important to select the optimal communication protocol. Newer protocols like TLS v1.3 are more secure and in many cases, are also more efficient.

In addition to secure communications, cryptographic keys and algorithms play an important role in securing IoT devices. IoT devices often only support the simpler symmetric key cryptographic algorithms (vs. public key cryptographic functions). However, with proper design (such as IETF Authentication and Authorization for Constrained Environments/OSCORE), it is possible to use public-key cryptographic functions in small IoT devices. The power consumption of complex computations can be reduced by using optimized hardware acceleration of cryptographic functions. It is therefore likely that future small IoT devices will have certain dedicated cryptographic hardware.

Persistent cryptographic keys must be stored securely and kept isolated from application software and physical interfaces. IoT devices can use the “Isolation” mechanism of Trusted Execution Environments (TEEs) to achieve this objective. Recently, ARM’s “TrustZone” TEE technology was brought to constrained devices. For more powerful devices, there are alternatives such as Intel SGX. Also, dedicated security components like Trusted Platform Modules or proprietary ASICs (application-specific integrated circuits) can be used. The goal is to develop solutions that can achieve a high level of security, albeit at higher cost and power consumption levels and in many use cases, integrated TEEs will be sufficient and more cost-effective.

To maintain security during their operational life, IoT devices should support secure software/firmware upgrade through “root of trust” mechanism. Secure upgrades are realized by having the software signed prior to release and having a trusted subsystem in the device perform a verification of the software before it is programmed/loaded into the device. New standardization work for securing updates for software/firmware. Procedures for secure device life-cycle management have been initiated. It is however complex and may have to be tailored for each specific use case.



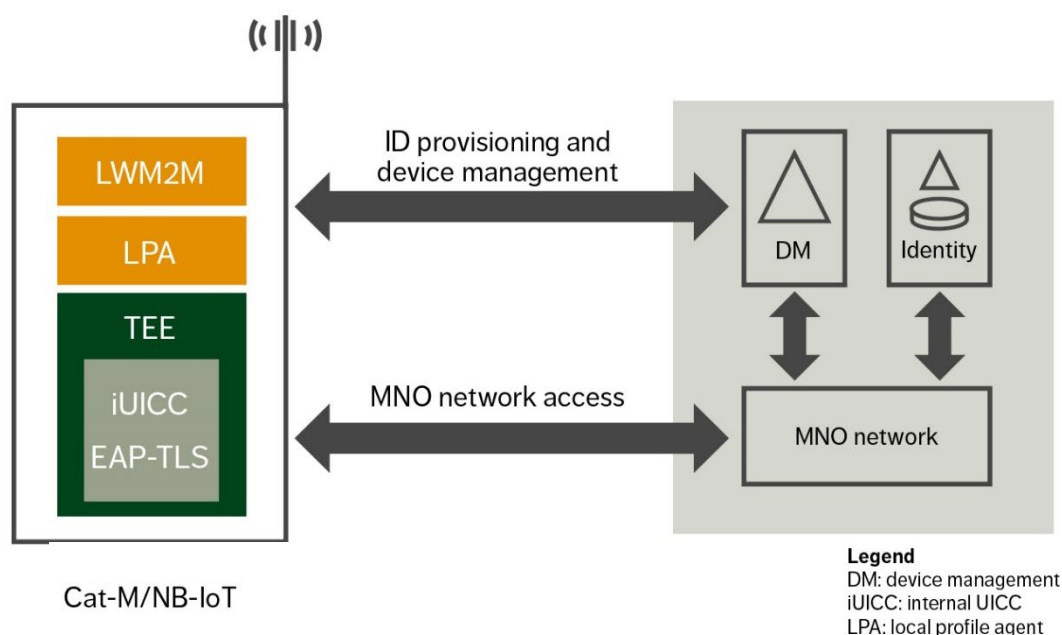
The criticality of device security is growing in the industry, however, much more efforts are needed to develop trusted systems that cover the needs of life-cycle management and applications security. Supporting secure software update is crucial to the creation of trustworthy IoT devices.

## 2.2. Identity

A digital identity is crucial in identifying device trustworthiness and in the overall network security. By trustworthiness, we mean that a device can prove that it has been produced by a legitimate manufacturer through an initial identity. A digital identity can be used for authentication, to maintain data ownership, or for software origin verification.

An identity consists of a securely stored secret and an assigned link between this secret and an identifier or name. This can be accomplished using a public key infrastructure (PKI), where the device holds a private key and the identity is a certificate that links this key to an identifier written into the certificate. For IoT devices, traditional PKIs have their problems. Their cryptographic operations can be cumbersome for highly constrained low complexity devices, the certificates can be large, and the certificate revocation management is usually so complex and error prone that it is hardly used. Furthermore, traditional PKIs have privacy issues. These issues can be addressed, using Enhanced Privacy ID, but at significantly higher complexity costs.

As an alternative to PKIs, it is possible to use identities based on symmetric key cryptography. This method is already in use for the 2G, 3G and 4G mobile network systems that use SIMs to hold the authentication credentials. SIMs use dedicated hardware chips and are relatively complex, mainly for legacy reasons. More cost-effective solutions are on their way, such as the integrated Universal Integrated Circuit Card (iUICC), in which the SIM hardware is integrated into the device processors. For 5G mobile network systems, symmetric key based identities for network access will remain in use, but in 5G it is also possible to use PKI-based identities via Extensible Authentication Protocol (EAP)-TLS. Figure below illustrates EAP-TLS ID management and use for network access.



**Figure 3 EAP-TLS ID management**

Blockchains can play even a more central role in the distributed approach to handling the trust in device identities. These options make it possible to link device lifecycle management with that of the device identity in a common framework.

### **2.3. Machine Intelligence**

MI technologies are essential to building IoT systems that can improve their own performance as more data becomes available and more knowledge is inferred and retained. With IoT and its associated large volume of data and billions of devices, MI is required to intelligently automate data collection and processing. Distributed MI (DMI) concerns the deployment, dynamic composition and life-cycle management of multi-node MI services, which can be chained for provisioning an intelligent system.

With distributed MI, it is possible to move the intelligence toward the device end (or the edge), which will minimize E2E latency, enhance data privacy and lower bandwidth requirements. Such on-device MI (ODMI) enables horizontal connectivity of devices to edge infrastructure that hosts DMI services. To accomplish this goal, IoT devices should be capable of performing low-power computation at or close to where the data is generated, or where the resulting action is needed. From a software perspective, it is possible to offload MI computation to hardware accelerators at the edge. In this layer, compilers and schedulers break down MI workflows, and distribute it into smaller tasks, which can help optimize the overall process.

Investment in scalable and flexible MI IoT systems will play a key role in evolving networks, as new devices, sensors, and actuators are being added and removed. Edge compute and ODMI are hence important in providing flexible deployment option and for managing new and complex Service Level Agreements.



# Conclusion

IoT use cases and technology are evolving rapidly. 3GPP standards provide a holistic, optimal, and flexible solutions covering a wide range of use cases. 3GPP's LTE-M and NB-IoT are optimized to serve a variety of use cases with global reach where small amount of data are sent infrequently, with high latency tolerance, low cost, extended coverage, and low battery consumption.

Today's 4G LTE broadband networks serve a large number of existing IoT use cases that require low latency and/or high throughput (e.g. Vehicle OEM communication, security cameras, etc.). 5G NR and 3GPP releases 14/15/16+ will enable an emerging set of use cases in critical IoT that require ultra reliable low latency communications, as well as the support of industrial IoT protocols that are used in the manufacturing sector.

When considering an IoT solution, careful consideration of security, device identity and IoT data management are key to success. With regards to security, the implementation of cryptographic functions on the device is the optimal approach to achieving strong device security. TEEs will soon be applied to IoT devices to support use cases in which secure storage and isolation are required. Secure identities are important to identify the origin of data and to ensure secure connectivity. New cost-efficient solutions for LPWAN access will emerge, leveraging the device's built-in security capabilities. Machine Intelligence will make it possible to move processing toward the device end and the mobile which will minimize E2E latency, enhance data privacy, and lower bandwidth requirements.

## Abbreviations

3GPP	Third Generation Partnership Project
AP	access point
ASIC	Application Specific Integrated Circuit
bps	bits per second
EAP	Extensible Authentication Protocol
EtherCat	Ethernet for Control Automation Technology
FEC	forward error correction
HFC	hybrid fiber-coax
HD	high definition
Hz	Hertz
iUICC	Integrated Universal Integrated Circuit Card
IoT	Internet of Things
ISBE	International Society of Broadband Experts
LPWAN	Low Power Wide Area Networks
LTE	Long Term Evolution
MI	Machine Intelligence
ML	Machine Learning
NR	New Radio – 3GPP 5G Radio Protocols
OEM	Original Equipment Manufacturer
OSCORE	Object Security for Constrained RESTful Environments
PROFINET	Portmanteau for Process Field Net
PKI	Public Key Infra-structure
QoS	Quality of Service
SCTE	Society of Cable Telecommunications Engineers

SGX	Software Guard Extensions
TEE	Trusted Execution Environment
TLS	Transport Layer Security
TSN	Time Sensitive Networking
URLLC	Ultra Reliable Low Latency Communications

## Bibliography & References

Ericsson white paper, January 2019, Cellular IoT Evolution for Industry Digitalization

Ericsson Technology Review, January 9, 2019, Key technology choices for optimal massive IoT devices

Ericsson Mobility Report – June 2019

5G America's white paper - LTE and 5G Technologies Enabling the Internet of Things

# Extended Spectrum DOCSIS®: A Pragmatic Approach

A Technical Paper prepared for SCTE•ISBE by

**Steve Condra**

Senior Engineering Director and Product Manager  
Teleste Intercept LLC  
440 Forsgate Drive, Cranbury NJ 08512  
+1 (678) 641 8099  
steven.condra@telesteintercept.com

**Kari Maki**

Manager, R&D  
Teleste Corporation  
Telestekatu 1, 20660 Littoinen, Finland  
+357 405540506  
kari.maki@teleste.com

**Arttu Purmonen**

VP, Systems Marketing  
Teleste Corporation  
Telestekatu 1, 20660 Littoinen, Finland  
+357 405540506  
kari.maki@teleste.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	4
1. Introduction .....	4
2. Tests.....	4
2.1. Hybrids.....	4
2.2. Amplifiers .....	6
2.3. Amplifier cascades .....	7
3. Methods and implications.....	8
3.1. Two alternative methods cope with high TCP.....	8
3.1.1. Method 1.....	9
3.1.2. Method 2.....	9
3.2. 3.2 Throughput.....	9
3.3. Practical guidelines.....	12
Conclusion .....	12
Limitations.....	12
Abbreviations.....	14
Bibliography & References .....	15

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Downstream Load.....	5
Figure 2 - Performance of the best hybrid .....	6
Figure 3 - An example load leading to 70 dBmV TCP.....	7
Figure 4 - Three cascaded amplifiers, theoretical MER versus measured MER .....	8
Figure 5 - Method 1 .....	8
Figure 6 - Method 2 .....	9
Figure 7 - Throughput, RPD Node and three cascaded amplifiers .....	10
Figure 8 - Throughput, RPD Node and four cascaded amplifiers .....	11

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - CM Minimum CNR performance (CM-SP-PHYv3.1-I11-170510) .....	10
Table 2 - MER, N+3 network.....	11
Table 3 - MER, N+4 network.....	12

# Introduction

Extended Spectrum DOCSIS (ESD) has been a topical subject due to ever-increasing broadband speeds in the digital landscape. While cable industry veterans have competed over who boasts the highest frequency, very little, if anything, has been published about what moving towards 1.8 GHz or even beyond that means in practice. We have performed real 1.8 GHz full spectrum measurements and, in the process, have revealed what it takes to offer new services using frequencies of up to 1.8 GHz. Our focus is on the amplifiers that are often needed even after distributed access roll-outs. The results of our measurements are enriched by cable operators who have contributed to and grounded our research by providing feedback and real network challenges. Our study covers variables that are expected to limit ESD implementations in North America, such as 1) length of cables, 2) length of amplifier cascades, 3) existing taps, 4) performance of the state-of-the-art amplifiers equipped with the latest hybrids and 5) capabilities of the latest Remote PHY (RPD) products. The results of our study provide pragmatic proposals for how DOCSIS OFDM frequency blocks are placed above currently employed frequencies and what kinds of limitations these proposals have. Our objective is to offer the latest information and unbiased practical proposals that can help cable operators obtain the most out of their networks with minimal changes. Although some changes will be crucial, significant costs can be mitigated through careful planning. Careful planning is not limited to the choice of amplifiers and taps, given that managing the interplay between RPDs and amplifiers must be considered to reach the rising broadband speed expectations of subscribers. Please take note of the following instructions for your technical paper or operational practice:

# Content

## 1. Introduction

It might be fairly easy to convert some cable networks to support N+0 architectures, while other networks will utilize amplifiers and amplifier cascades even when the next decade arrives. These amplifiers should work up to 1.8 GHz, which is not a walk in the park for engineers. In comparison to 1 GHz networks, the attenuation of coaxial cables at 1.8 GHz is over 40% more. Besides this challenge, also taps and splitters, even when designed for 1.8 GHz networks, cannot have the same attenuation at 1.8 GHz as their predecessors had at 1 GHz. To cope with these challenges, higher amplifier output levels or alternative workarounds are needed. Real tests in real conditions reveal to us what can be expected when the latest amplifier technologies enter the markets in 2020.


Before the tests, we wanted to be sure that the measurements describe a new reality and unlearn old parameters that would falsify our results. Existing amplifier cascades in the field today were built when old-school cable experts used to discuss many parameters, including composite second order (CSO) and composite triple beat (CTB) intermodulation distortion. However, new indicators are needed when services are digital and advanced modulation methods overtake cable networks. These new indicators, Modulation Error Ratio (MER), Bit Error Ratio (BER), Total Composite Power (TCP) and Carrier to Interference Noise Ratio (CINR), are valid scales for analyzing how networks and devices perform when the load is digital. While BER is the only thing that matters for end users, MER is faster for measuring and can be used to indicate BER. MER is also a better indicator than BER as new Forward Error Correction (FEC) methods introduced along DOCSIS 3.1 are extremely effective. The combination of Bose-Chaudhuri-Hocquenghem (BCH) and low-density-parity-check (LDPC) coding is so strong that BER values are either perfect or inferior, but seldom something between, while MER offers a more comprehensive overall snapshot of how tight margins networks have. It should be noted that although DOCSIS standards define MER, they still define carrier to noise ratio (CNR) as well. However, we can use MER and CNR analogously in calculations because network noise can be assumed to have a Gaussian distribution, as we will see in later chapters.

## 2. Tests

We started our tests by exploring 1.8 GHz hybrids. After these tests, we proceeded to amplifiers and amplifier cascades. Amplifier cascades were also modelled theoretically to understand if theoretical models and real test results are consistent with each other. In such a case, theoretical models could be used to complement and overcome possible uncertainties raised by real measurement results.


### 2.1. Hybrids

In February 2019, we tested the latest 1.8 GHz hybrids in our R&D, which were still prototypes. We used a various mixed single carry QAM (SC-QAM) and OFDM loads (an example is shown in Figure 1). The used frequency ranges were 602 MHz ... 1218 MHz and 1218 MHz ... 1794 MHz, respectively. The source MER was 51 dB over the whole frequency range. This source MER was selected because it is a realistic, perhaps even pessimistic, portrayal of performance that current RPD products are capable of if we ignore their amplifier stages. However, it is enough to fulfil DOCSIS specification for OFDM (*CM-SP-PHY*) and SC-QAM



**Figure 1 - Downstream Load**

(*CM-SP-DRFI*) when 51 dB MER is deteriorated by the amplifier stages that are integrated to the RPD nodes. Figure 2 reveals the performance of the very best hybrid model that we tested in several TCP points. The lines have some angularity as values between the TCP points are approximations. The figure shows interesting spots at 609 MHz (SC-QAM) and 1.7 GHz (OFDM). The worse OFDM MER was clearly caused by the higher frequency, not by OFDM. Indeed, tests with various load mixes revealed that the load caused by SC-QAM and OFDM does not differ if the level and frequency are the same. We increased the full spectrum load until MER at 1.7 GHz in the output of the hybrid reached 40 dB, while Pre-FEC BER was better than 1E-9. At this point, the impact of the source MER was negligible (less than 0.5 dB) and SC-QAM MER at lower frequencies was sound. The best performing hybrid model was able to produce 72 dBmV TCP, while the worst (not in the figure) hybrid reached 70 dBmV. In June 2019, we had improved 1.8 GHz hybrids in our R&D. Now the highest performing hybrid model was able to produce 74 dBmV TCP in the similar setup under the same criteria as in February 2019. Our current estimate is that once hybrids are available in volumes their performance will reach 76 dBmV TCP under the same conditions, and amplifiers equipped with these hybrids are available in 2020.



**Figure 2 - Performance of the best hybrid**

## 2.2. Amplifiers


Amplifiers have signal losses after hybrid components as additional components such as feedthrough current chokes, diplexers and connectors are needed in them. In total, these components have around a 6 dB loss at 1.8 GHz. As higher frequencies carry higher power and cause worse non-linearity, we can estimate that a 1.8 GHz amplifier equipped with the state-of-the-art hybrid (76 dBmV TCP) has at least 70 dBmV TCP in the output port. Figure 3 illustrates what it means in practice if the lowest downstream channels are around 500 MHz and the full load burdens an amplifier up to 1.8 GHz. The output is sloped up to 1.2 GHz, meaning that the virtual level at 1.8 GHz is 56 dBmV, while the practical level is around 48 dBmV for channels above 1.2 GHz. The virtual level can be used to calculate the needed gain. Before calculating the gain we must know the lowest allowed downstream input level that does not lead to poor amplifier CNR impacting MER negatively. Our target is to reach 57 dB CNR, while the noise figure (NF) is 10 dB.

$$CNR = \text{The lowest input level} - NF - \text{Thermal noise}$$

$$\text{The lowest input level} = CNR + NF + \text{Thermal noise}$$

$$\text{The lowest input level} = 57 \text{ dB} + 10 \text{ dB} + (-57 \text{ dBmV}) = 10 \text{ dBmV}$$





**Figure 3 - An example load leading to 70 dBmV TCP**

On the other hand, the output level is limited to 56 dBmV as Figure 3 shows. Although the 56 dBmV limitation is virtual, it matters as it defines the needed gain through the needed slope illustrated in Figure 3. Thus, we can calculate that a 46 dB gain is needed as we know that the lowest input level is 10 dBmV.

$$56 \text{ dBmV} - 10 \text{ dBmV} = 46 \text{ dB}$$

The 46 dB gain covers a loss of 1300 ft. cable (P3 500) or, alternatively, it can compensate a feedthrough loss of seven 1.8 GHz taps (3.5dB@1.8GHz) and 90 ft. cable between the taps.

### 2.3. Amplifier cascades


A theoretical analysis of amplifier cascades can be done using the following equation:

$$MER_{out} = -10 \log_{10} \left( 10^{\frac{MER_{in}}{-10}} + 10^{\frac{CNR_{amp1}}{-10}} + 10^{\frac{CNR_{amp2}}{-10}} + 10^{\frac{CNR_{ampn}}{-10}} \right)$$

The equation points how the CNR performance of amplifiers reduces MER in the output. The equation can be simplified when all amplifiers have the same CNR performance.

$$MER_{out} = -10 \log_{10} \left( 10^{\frac{MER_{in}}{-10}} + n \times 10^{\frac{CNR_{amp}}{-10}} \right)$$

The theoretical analysis of three cascaded amplifiers is shown in Figure 4, illustrating how theoretical MER decreases in function of the output level as amplifiers have a lower CNR when the output level increases. Besides the theoretical calculation, we measured real cascades. Less surprisingly, real cascades behave almost according to the theory. However, even when distortions start to dominate, the theoretical analysis holds, although it should apply only when noise can be assumed to have the Gaussian distribution.

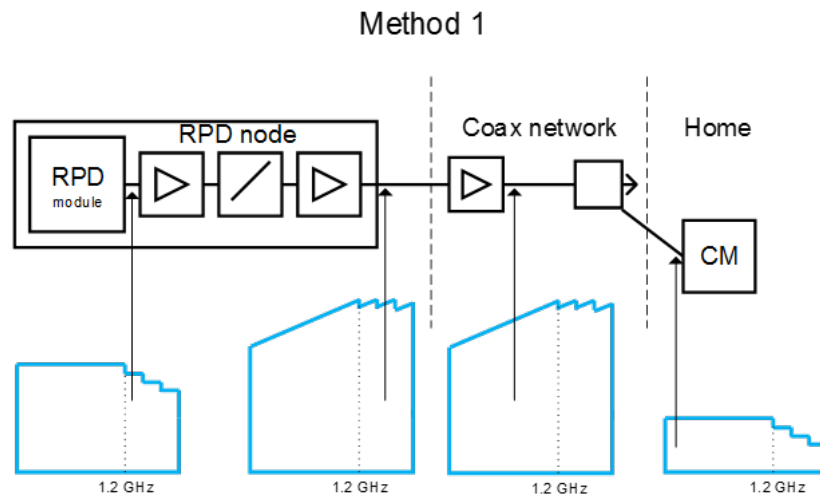


**Figure 4 - Three cascaded amplifiers, theoretical MER versus measured MER**

### 3. Methods and implications

#### 3.1. Two alternative methods cope with high TCP

The linearity of amplifiers declines at higher frequencies. Therefore, the maximum TCP that 1.2 GHz amplifiers are able to produce is not available with 1.8 GHz amplifiers. As RF load on higher frequencies limits TCP more than the same load on lower frequencies, the solution could reduce the RF load above 1.2 GHz. We propose two alternative methods that can also be done in practice.



**Figure 5 - Method 1**

### 3.1.1. Method 1

In Method 1, the reduction of RF power is performed by a remote PHY device (RPD) node using back-off for OFDM signals as Figure 5 demonstrates. The RPD node is equipped with amplifier stages in combination with high pass filters. First, the RPD staggers OFDM blocks and after the amplifier stages and filters the output of the RPD node is sloped up to 1.2 GHz. While every OFDM block has the same slope their signal level is reduced. Due to the tilt of coaxial cables, the Cable Modem (CM) sees a flat level until 1.2 GHz and the staggered OFDM blocks up to 1.8 GHz.




Figure 6 - Method 2

### 3.1.2. Method 2

Figure 6 describes Method 2 that uses a flat top above 1.2 GHz. This is achieved by filters before the last amplifier stage. Channels below 1.2 GHz are sloped in the output of the RPD node in the same way as in Method 1. Due to the tilt of coaxial cables, the cable modem sees a flat top until 1.2 GHz and every received OFDM channel has around -3 dB negative slope. The flat top approach can use, for instance, a 1 GHz verge frequency instead of 1.2 GHz if it is seen as more appropriate for the existing network.

## 3.2. 3.2 Throughput


Based on measurements demonstrating performance of 1.8 GHz amplifiers, we built two example cases demonstrating throughput of the 1.8 GHz network by using method 1. Both cases use 492/602 MHz split and similar upstream load but different downstream modulations and cascade lengths. The throughput and used modulation methods in the cases are presented in Figures 7 and 8. Networks in both cases employ frequencies of up to 108 MHz for legacy services and frequencies between 108 MHz and 492 MHz for four OFDM blocks (each 96 MHz).

**Table 1 - CM Minimum CNR performance (CM-SP-PHYv3.1-I11-170510)**

Constellation	CNR (dB) Up to 1 GHz	CNR (dB) 1 GHz to 1.2 GHz	Min P6 <sub>AVG</sub> dBmV
4096	41.0	41.5	-6
2048	37.0	37.5	-9
1024	34.0	34.0	-12
512	30.5	30.5	-12
256	27.0	27.0	-15

It should be noted that in both cases even the RPD node includes amplifiers as we expose in Figures 5 and 6. Both cases employ cabling and taps but they do not impact MER as they are passive. However, as passives elements attenuate, we made sure that cable modems received enough high signal levels specified in the DOCSIS standard extract presented in Table 1. In both cases, the TCP was in line with limits discussed in the section Amplifiers (2.2).

In the first case (Figure 7), we use frequencies between 602 MHz and 814 MHz for 37 SC-QAM channels, frequencies between 814 MHz and 1402 MHz for 1024 OFDM and frequencies between 1402 MHz and 1794 MHz for 512 OFDM. The setup consists of one RPD node and three cascaded 1.8 GHz amplifiers. Table 2 shows MER over four different frequencies, in the RPD node output and after 3 amplifiers. With the given values, the example can be used to reach 9.7 Gbps downstream capacity.




**Figure 7 - Throughput, RPD Node and three cascaded amplifiers**

**Table 2 - MER, N+3 network**

Frequency	MER (N+3 network)	
	RPD Node	N+3
830 MHz	48.0 dB	43.0
1.1 GHz	47.5 dB	42.5
1.3 GHz	45.5 dB	39.0
1.7 GHz	41.5 dB	34.5

In the second case (Figure 8), we stress the network even further through higher order modulation methods and a longer amplifier cascade. Now we use frequencies between 814 MHz and 1218 MHz for 4096 OFDM, frequencies between 1218 MHz and 1410 MHz for 2048 OFDM, frequencies between 1410 MHz and 1602 MHz are used for 1024 OFDM, while frequencies above 1602 MHz are used for 512 OFDM. In this case, we have an RDP node followed by four cascaded amplifiers, the MER is reported in Table 3 over different frequencies in the RPD node output and after four amplifiers. With the given values the example leads to 10.6 Gbps downstream capacity being around 1 Gbps higher than in the first case, although the cascade of the amplifiers is longer. The difference is explained by more effective use of frequencies and lower MER margins than in the first case.



**Figure 8 - Throughput, RPD Node and four cascaded amplifiers**

**Table 3 - MER, N+4 network**

Frequency	MER (N+4 network)	
	RPD Node	N+4
830 MHz	48.0 dB	42.0
1.1 GHz	47.5 dB	41.5
1.3 GHz	45.5 dB	38.0
1.7 GHz	41.5 dB	33.5

### 3.3. Practical guidelines

Certain practical details must be considered when 1.8 GHz amplifier cascades are built.

1. Amplifiers must have a cable equivalent frequency response apart from the used input equalizer values. This eliminates cumulating errors that a linear frequency response would cause. As these amplifiers will compensate preceding cables, amplifier outputs would have the same linear frequency response that exists in the RPD node output.
2. The accuracy of up and downstream alignments becomes paramount when MER margins turn narrow. Our proposal would be to use automatic adjustments performed by the amplifiers as manual “roughly right” will not be enough when the last decibels matter.
3. Automatic Level and Slope Control (ALSC) must support flexible pilot frequencies if networks will first employ 1.0 GHz or 1.2 GHz frequencies and later on are upgraded to employ frequencies of up to 1.8 GHz. Otherwise, operators are forced to change pilot detection units during the upgrade.
4. 1.8 GHz amplifiers will need more power even if the state-of-the-art technology is used. Not only because of higher downstream frequencies but also because of the higher upstream frequencies. Even if future hybrids become more effective, alternative ways to mitigate increased power consumption should be investigated. Currently available adaptive power methods and active power factor correction are examples, but research producing even more effective methods should continue.

## Conclusion

As we have discussed, amplifier cascades will be a reality in the coming years even though N+0 networks are on the horizon. However, as we substantiate by theoretical and practical methods, even four amplifiers in a cascade can carry the magical 10 Gbps capacity. Our paper presents two methods that cope with high TCP and provides examples of how the methods could be exploited. More importantly, these methods are based on technologies that are commercially available in 2020. Nonetheless, to harvest the full potential of HFC networks, 1.8 GHz amplifiers should perform automatic adjustments or alternatively cable technicians should define rigorous methods to test that amplifier cascades are tuned to perfection even when outdoor conditions such as temperature change.

## Limitations

Our study used a 492/606 MHz split, although it is only one of the options. Moreover, tighter guard bands are possible if more complex diplexer technologies are used, but we wanted to stay pragmatic and use a method that has been widely tested, namely robust but changeable diplexer plug-ins. While FDX amplifiers offer significantly tighter guard band, our study did not cover their use. As industry has discussed their

benefits, we encourage future studies to address their limitations, such as increased complexity, higher power consumption and lower CNR performance.

# Abbreviations

ALSC	automatic level and slope control
BCH	Bose-Chaudhuri-Hocquenghem
BER	bit error ratio
CINR	carrier to interference noise ratio
CM	cable modem
CNR	carrier to noise ratio
CSO	composite second order
CTB	composite triple beat
dB	decibel
dBmV	decibel millivolt
DOCSIS	Data-Over-Cable Service Interface Specifications
ESD	extended spectrum DOCSIS
FDX	full duplex
FEC	forward error correction
Gbps	gigabits per second
GHz	gigahertz
HFC	hybrid fiber coax
LDPC	low-density-parity-check
Mbps	megabits per second
MER	modulation error ratio
MHz	megahertz
NF	noise figure
OFDM	orthogonal frequency division multiplex
QAM	quadrature amplitude modulation
RPD	remote PHY device
SC-QAM	single carry QAM
SNR	signal to noise ratio
TCP	total composite power



## Bibliography & References

- (1) Chapman, Emmendorfer, Howald and Shulman. 2012. Mission is Possible: An Evolutionary Approach to Gigabit-Class DOCSIS. <http://www.bowe.id.au/michael/isp/DOCSIS/collected-references/mission-possible-evolutionary-approach-to-docsis-whitepaper.pdf>. Accessed 10 June 2019.
- (2) Werner, Tony. 2015. World's First Live DOCSIS 3.1 Gigabit-Class Modem Goes Online in Philadelphia. <https://corporate.comcast.com/comcast-voices/worlds-first-live-docsis-3-1-gigabit-class-modem-goes-online-in-philadelphia#.VnlhdFcTuHJ.twitter>. Accessed 10 June 2019.
- (3) CableLabs®. Remote PHY specifications. <https://specificationsearch.cablelabs.com/?query=&category=DOCSIS&subcat=MHAV2&doctype=&content=false&archives=false&currentPage=1>. Accessed 25 June 2019.
- (4) Association of German Cable Operators. 2015. ANGA COM Exhibition & Conference program. [http://www.angacom.de/fileadmin/user\\_upload/presse/pdf-downloads/2015/ANGA\\_COM\\_2015\\_Congress\\_Brochure.pdf](http://www.angacom.de/fileadmin/user_upload/presse/pdf-downloads/2015/ANGA_COM_2015_Congress_Brochure.pdf). Accessed 25 June 2019.
- (5) Teleste Corporation. 2015. Teleste introduces new DOCSIS® 3.1-compliant nodes to celebrate the milestone of 400.000 delivered units for the Teleste AC family. <https://www.teleste.com/news/2015/teleste-introduces-new-docsis%C2%AE-31-compliant-nodes-celebrate-milestone-400000-delivered-units-teleste-ac-family>. Accessed 25 June 2019.
- (6) Baumgartner, Jeff. 2018. Xfinity Mobile Starting to See ‘Real Momentum’. <https://www.multichannel.com/news/xfinity-mobile-starting-see-real-momentum-417693>. Accessed 12 June 2019.
- (7) CableLabs®. 2017. CableLabs Completes Full Duplex DOCSIS Specification. <https://www.cablelabs.com/cablelabs-completes-full-duplex-docsis-specification>. Accessed 27 June 2019.
- (8) Breznick, Alan. 2018. Here Comes DOCSIS 4.0. <https://www.lightreading.com/cable/docsis/here-comes-docsis-40/d/d-id/743285>. Accessed 12 June 2019.
- (9) Finkelstein, Jeff. 2018. No Matter Where You Go, There You Are... <https://broadbandlibrary.com/no-matter-where-you-go-there-you-are/>. Accessed 12 June 2019.
- (10) Cloonan, Tom. 2018. Technologies to Help Solve Future Broadband Issues (RMD, RMC, Extended Spectrum DOCSIS & Others). [https://www.scte.org/SCTEDocs/Expo/DAA2018/SCTE18\\_DAA\\_TomCloonan\\_Arris\\_FINAL.pdf](https://www.scte.org/SCTEDocs/Expo/DAA2018/SCTE18_DAA_TomCloonan_Arris_FINAL.pdf). Accessed 10 June 2019.
- (11) Baumgartner, Jeff. 2019. CableLabs Kicks Off Pursuit of DOCSIS 4.0. <https://www.lightreading.com/cable/docsis/cablelabs-kicks-off-pursuit-of-docsis-40/d/d-id/752355>. Accessed 24 June 2019.

# Connected Independence

A Technical Paper prepared for SCTE•ISBE by

**Kevin Alcox**  
CTO  
iEldra  
Littleton, CO  
303-886-8415  
k.alcox@ieldra.com

**Thomas Priore**  
Vice President Product and Architecture  
iEldra  
Jim Thorpe, PA  
215-948-2890  
t.priore@ieldra.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Technology and Process .....	4
1. Architecture.....	7
1.1. IoT Sensors.....	7
1.1.1. Sensor Types.....	8
1.1.2. Sensor Selection.....	8
1.2. IoT Gateway.....	9
1.3. Broadband Gateway .....	10
1.4. Cloud Infrastructure.....	10
2. User Experience .....	10
2.1. Subscriber.....	11
2.2. Family Caregiver .....	11
2.3. Professional Caregiver .....	12
2.4. Institutional Caregiver .....	12
2.5. Call Center Agent .....	12
2.6. Healthcare Provider.....	12
2.7. Case Manager.....	13
3. Key Metrics.....	13
3.1. Sensors Metrics.....	14
3.1.1. Sensor Battery Levels .....	14
3.1.2. Other Sensor Failures .....	15
3.2. IoT and Broadband Gateway availability .....	16
3.3. Escalation and Incident Response .....	16
Conclusion .....	16
Abbreviations.....	18
Bibliography & References .....	18

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1: Care Crisis Evidence .....	3
Figure 2: Representative Health and Wellness System Deployment.....	6
Figure 3: Multi Sensor Battery Life .....	14
Figure 4: Motion Sensor Battery Life .....	14

## List of Tables

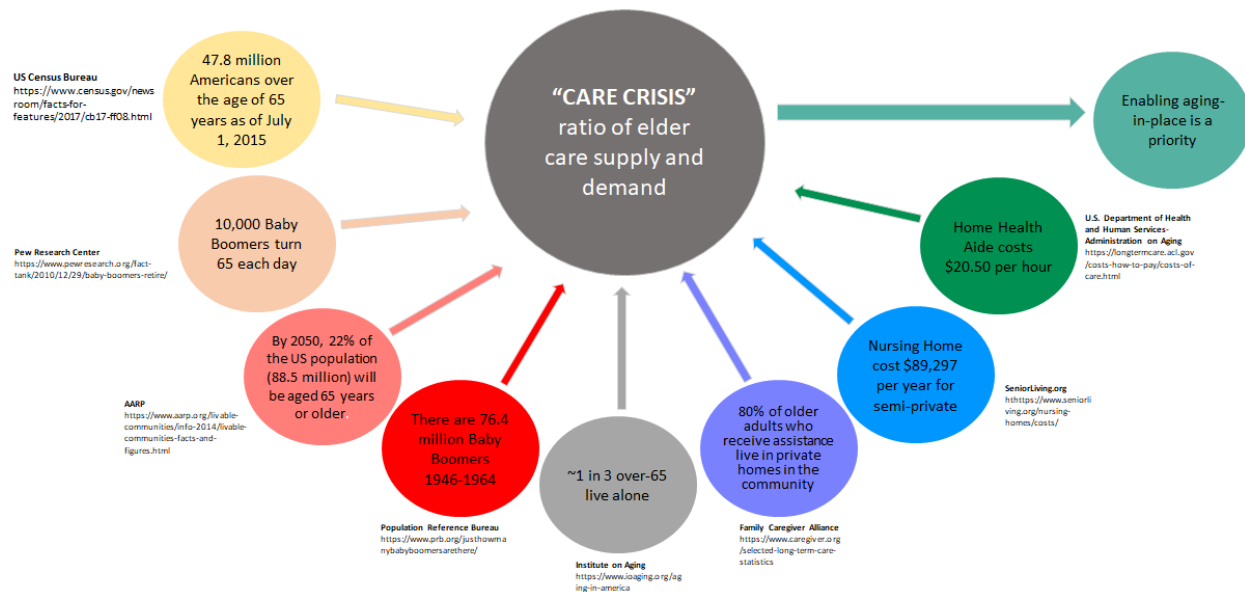
<b>Title</b>	<b>Page Number</b>
Table 1:ADLs and IADLs .....	5

# Introduction

This paper discusses how connectivity, IoT devices, and AI can address a growing worldwide problem: caring for our elders. By connecting the latest IoT technology and remote caregivers, aging in place or aging in place longer becomes a viable option for any older adult or loved one that has access to a broadband network.

While this paper recognizes that the world's older population is booming and the number of caregivers for that population is declining, creating a "care crisis", it does focus on aging, healthcare and insurance in the United States. Figure 1 below, along with the bullets above it emphasizes the scale of this problem. The burden on the healthcare industry to provide better care and better health at a lower cost is mounting. Technology in the form of Passive Remote Patient Monitoring (PRPM), Remote Patient Monitoring (RPM) and Telehealth is emerging as the go-to-resource multiplier to improve the way providers and caregivers take care of these individuals.

- By 2020 about 45 million Americans will be caring for 117 million seniors, taking responsibility for food delivery and health monitoring.
- Three out of four caregivers want to use technology to make their duties easier, but only 7% have done so, according to a 2016 study sponsored by AARP.
- By 2030, 20% of the US population will be aged 65 years or older
- As of Q4 2016, 79.5% of householders age 65 and older owned their homes
- Assisted living facilities cost an average of \$43,200 per year
- A Home Health Aide costs an average of \$20 per hour
- The average length of stay in a skilled nursing facility is 835 days (2.28 years)



**Figure 1: Care Crisis Evidence**

"As the number of people over the age of 80 increases in the next 20 years, the number of people in the primary caregiving years will remain flat," states AARP's 2013 report. Meanwhile, in 2050, there will be three times as many people age 80 and older as there are today. As a result, by 2050, the caregiver support ratio which was 7.2 in 2010 when Boomers were in their peak caregiving years, is projected to drop to 2.9

percent when the boomers will have reached their eighties. According to AARP's report, "In just 13 years, as the Baby Boomers age into their 80's, the decline in the caregiver support ratio will shift from a slow decline to a free fall."

In summary, it is not uncommon for someone to receive care at home for several months or longer, followed by a two and a half year stay in an assisted living facility, with almost 60% then requiring a nursing home stay of somewhere between nine months and a little over two years. All combined, this is a total of approximately 4-5 years of long-term care. In this scenario, the total cost of care could easily exceed \$300,000. This is daunting considering that it would be in addition to the approximately \$245,000 that Fidelity Investments estimates the average retired couple will spend on healthcare- other than assisted living or nursing care expenses- during the span of their retirement years.

This reality is being recognized by many, including the Centers for Medicare & Medicaid Services (CMS). CMS' latest transformation initiative is called the Meaningful Measures framework which identifies the highest priorities for quality measurement and improvement and looks to drive change in those areas.

Healthcare transformation has existed since the beginning of medicine, driven by culture, beliefs, values and technology. The latest evolutionary phase began in the early 2000's driven by reimbursement and technology changes and is yet to reach a stable, functional model. Coordinated, evidence-based accountable care is now a key driver to healthcare transformation to deliver better care and better health at lower costs. Connectivity in the home and in the form of wearables is the platform to enable this. Technology becomes both a safety net to support independent living and a lever to multiply the amount of finite care available from a dwindling caregiver population. To accelerate this change, CMS is authorizing Medicare Advantage beneficiaries to be reimbursed for technology assistance starting in calendar year 2019.

A health and wellness platform offers MSOs a unique opportunity to participate in this latest healthcare transformation. With reimbursement change being driven by the Centers for Medicare and Medicaid and that change being fundamentally dependent upon health and wellness information, better coordinated care and teamwork among the caregivers, connectivity and technology providers will be an integral part of a successful and stable transformation. The connected home market is estimated to be about \$28 billion, growing roughly 20% annually. Markets include broadband companies, home healthcare providers, assisted living, nursing homes & the consumer.

With HIPAA-compliant wireless IoT-based sensors and smart speakers placed in the home, connected to secure communication links between the home and cloud-based AI engines, the technology foundation for transformation is established. Proactive monitoring and measurement lead to improved risk management, the primary goal for CMS and insurance providers and an opportunity for MSOs.


## Technology and Process

Activities of Daily Living (ADLs) are standards that describe routine activities that people tend to do every day without needing assistance. There are six basic ADLs: eating, bathing, dressing, toileting, transferring (walking), continence, and several others that contribute to the ability for a person to live independently. Table 1 presents the full, standard list of ADLs and IADLs. The performance of these ADLs is important for determining what type of long-term care is required. The commonly accepted activities of daily living are listed in the Roper-Logan-Tierney Model of Nursing (RLT). There are also Instrumental Activities of Daily Living (IADL) defined by Lawton & Brody. The IADLs are a more complex measure of a person's ability to live independently.

**Table 1:ADLs and IADLs**

ADL	Source	Type
Breathing	RLT	ADL
Communication	RLT	ADL
Continence	RLT	ADL
Controlling body temperature	RLT	ADL
Dressing	RLT	ADL
Drinking	RLT	ADL
Eating	RLT	ADL
Elimination (Toileting)	RLT	ADL
Maintaining a safe living environment	RLT	ADL
Mobilization - Body Movement	RLT	ADL
Sleeping	RLT	ADL
Bathing	RLT	ADL
Working and playing with a sense of purpose	RLT	ADL
Handling Transportation (driving or navigating public transit)	Lawton-Brody	IADL
Housework and Basic Home Maintenance	Lawton-Brody	IADL
Managing Finances	Lawton-Brody	IADL
Managing Medications	Lawton-Brody	IADL
Preparing Meals	Lawton-Brody	IADL
Shopping	Lawton-Brody	IADL

Today’s technology supports a platform solution that provides peace of mind to families by enhancing the “Aging at Home” experience with an economical and comprehensive, independent-living activity measuring and monitoring solution. This technology-based solution uses unobtrusive (i.e.; no cameras) sensors to track, measure, and assess the ADLs and IADLs of a subscriber. Connecting this monitoring and measuring with timely human acknowledgment, interaction, and resolution through a caregiver network, including a 24/7 customer care center will provide the guardrails to allow more people to age at home. In a majority of situations, adult children have assumed the burden of caring for their parents, but still have the responsibility of their own families and frequently live in a different local than their parents. This health and wellness system enables an older adult to stay in their home and maintain a level of independence, yet creates a safety net and means for caregivers to actively care for an aging adult without being physically tied to the older adult’s home will revolutionize the well-being of older adults. With nearly 48 million seniors in America and 90% wanting to stay in their homes, technology is the enabler to make this happen.



**Figure 2: Representative Health and Wellness System Deployment**

IoT devices, that when properly configured and deployed within a living space, along with pattern analysis can help identify the fulfillment or lack of fulfillment of an ADLs/IADLs. A minimum viable product would be comprised of on-premise IoT sensors and a custom voice assistant configured to identify ADL event occurrences, assess IADL and ADL competency, provide 2-way audio or physical communications and securely communicate event attributes to an IoT gateway. The IoT gateway provides local infrastructure management, event data caching and secured event data transport to a cloud-based operational support system (OSS). The secured OSS includes artificial intelligence and machine learning to automatically configure and optimize the ADL event monitoring algorithms, sensor

configuration actions, and IADL assessment actions. Figure 2 presents a conceptual architecture of this platform.

It is extremely important that the system is tuned to avoid “alarm fatigue”. A variety of modern communications methods, chosen by each individual caregiver to suit their own preferred communication method should be utilized to engage remote support assistance in a recursive model to assure an event response.

This entire communication path should be secured using a combination of FIPS 140-2 compliant modern cryptography methods such as TLS, RSA-1024 certificates, AES-128, AES-CCM, and ECDH key exchange to produce the highly secure and HIPAA-compliant health and wellness measurement and monitoring system.

The platform is built upon the tenet that technology can be leveraged to provide a safety net for those wishing to maintain their independent living lifestyle. The system’s second key tenet is that humans will establish a normal and measurable pattern of physical and cognitive behavior, deviations from which indicating areas of concern for the caregivers of that individual. Care plans must be unique for every independent living person. The care plan establishes a normal pattern of behavior, defines the thresholds for specific monitorable actions and a set of users. The activity engine receives the raw events and applies a series of care plan rules with their custom thresholds via a machine learning model. This statistical data-to-decision process generates an alert when a deviation from normal is detected and also automatically adjusts the monitoring action thresholds to optimize the alerting to the independent living person’s normal patterns.

## **1. Architecture**

### **1.1. IoT Sensors**

Critical to the assessment of the ADLs/IADLs are the sensors that pick up the activity of the person being monitored. The sensors pick up the physical and environmental cues of the individual’s activity that defines a pattern of daily activity.

The IoT sensors may be connected to each other and/or other resources via wired or wireless connections. Today’s industry standards for these connections may be implemented through the use of any known wired or wireless communication standard, including but not limited to Ethernet, 802.11a/b/g/n/x, universal serial bus (USB), Bluetooth, Bluetooth LE, cellular, near-field communications (NFC), Z-wave, ZigBee or even a proprietary standard.

Sensors can relay events to an IoT gateway in one of two ways. The sensor can be polled by the IoT gateway and the state of the sensor can be read, or the sensor can be set either by the manufacturer or via a configuration to transmit an event when a sensor threshold is met. For example, a temperature sensor can be polled every five minutes to generate a stream of events. Alternatively, the temperature sensor can be sent to transmit an event for every 0.1° of change. Both methods have benefits and drawbacks. Polling unchanged data risks reducing battery life. Using thresholds requires thresholds to be set at a sensitive enough level to collect the needed information for analysis.

When a threshold is met, or a sensor is polled, it creates an event that is transmitted to the IoT gateway. The IoT gateway converts the signals from the IoT network nodes into a communications protocol that is capable of transferring the events to a cloud-based computing infrastructure that supports complex algorithmic data analysis and event handling workflow.



### **1.1.1. Sensor Types**

There many types of IoT sensors and many sensor packages are capable of collecting multiple types of data simultaneously.

- Motion Sensors are used to ensure that a person is active and moving around when they should be, and not active at times they should not be. They can be used to monitor waking and sleeping, trips to the bathroom, and movement around the house or rate of movement. Motion sensors can be infrared-based, radar-based, laser-based or ultrasonic.
- Temperature Sensors are used to ensure that the environment is safe for a person. It can be used to confirm that climate control is working properly and is being used. Localized temperature changes can also be an indicator of stove or oven use or a personal hygiene event (i.e.; taking a bath) and the duration of that event.
- Light and UV Sensors are used in conjunction with motion sensors to help establish activity, especially at night. Motion at night along with the presence of light can indicate that a person is awake and intend to stay awake for a period of time. UV sensors can help identify if the light source is natural (the sun) or man-made.
- Contact (aka; Door/Window) Sensors can be used to determine if a person is arriving or leaving the home, has left a door open or closed for too long or used a door at a time that they shouldn't be. These contact sensors can also be used to determine if food preparation is occurring by monitoring the opening and closing of refrigerators, microwaves, and ovens or be used to know when a particular cabinet or drawer has been accessed (i.e.; a medicine cabinet).
- Pressure sensors can be used to determine a more precise location of a person. They can identify is a person is standing at a specific location (i.e.; near a toilet), lying in bed, or sitting in a chair.
- Smoke and Gas Sensors can be used to identify if there is a safe environment for a person. It can be used to identify fire, or gas leaks, especially where stoves are involved.
- Moisture Sensors can be used to identify environmental issues (i.e.; water leaks or flooding) or personal hygiene incidents.
- Biomedical - There are many biomedical sensors on the market for measuring medically relevant parameters such as heart rate, respiration, oxygen levels, glucose levels, weight, temperature, and many others.

### **1.1.2. Sensor Selection**

Sensor selection is critical for the accurate and reliable measurements of ADLs. There are many considerations in the selection of sensors. Poor sensor selection can have a very detrimental effect on the systems operations or the ability to support deployments. Below are some of the parameters of sensors that should be examined when making a selection.

- Ease of installation - The installation of sensors should be easy to do with common tools or no tools at all. The sensor should be designed so that the installation should be reliable; the sensor should be capable of being installed in a way so that it will stay in the location at which it is installed and stay orientated properly.

- **Powering** - Considerations must be made on how the device is to be powered. Batteries must be monitored and changed when low. Some devices may be difficult to access to change the batteries. And some sensors use an uncommon type of battery that is not readily available via normal retail outlets. If a device is powered via a wall outlet, the length of the power cable and location of outlets limit installation options.
- **Sensitivity / Accuracy / Detection Zone**- Sensors should be capable of monitoring the environment they are deployed in. They should have the appropriate sensitivity to measure indications of the presence of people, motion, temperature, and light. For example, a motion sensor should be able to detect motion at a reasonable distance in a room. It is not necessary for the sensor to detect motion across an entire room as sensors can be placed in choke points or areas people are forced to travel through as they move from room to room.
- **Communication Range** - Sensors need to communicate back to the IoT gateway. All wireless sensors have a rated range but in actual real-world deployments, that range tends to be much smaller than what is advertised. Sensors placed out of range for direct communication with the gateway can still be communicated with by placing repeaters throughout the home or using a mesh network technology. Repeaters do add cost, installation complications, and complexity to the deployment that could affect reliability.
- **Tamper Detection** - Even though the primary use is not as a security system, it is still important to know if a sensor is being tampered with. Knowing whether a sensor has been touched helps diagnose detection issues or other failures of the sensor and its event reporting.
- **Communication Method** - Whether it's wired ethernet, or a wireless standard, the communication method of the sensor has a large impact on other selection considerations. It affects powering options and battery life, installation options and the ease of installation, the distance at which the sensor can be placed from a gateway, security of the system, the cost of the sensor itself and the IoT gateway cost and functionality. Deployments with a mix of communication methods can complicate installations and drive up the cost of gateways.
- **Cost** - As there will be multiple sensors per deployment, the cost of sensors adds up very quickly. Every other consideration has an effect on cost. Sensors that are powered by batteries tend to be cheaper than sensors that include power supplies, but servicing batteries also has a cost. More reliable communication methods tend to drive up cost, but may simplify installation. Cost is always a balancing act and it's important to factor in the long-term operational costs of maintaining a sensor, and not just examining the per unit cost of acquiring a sensor.

## 1.2. IoT Gateway

The IoT Gateway provides connectivity to the mesh of sensors and the local area network. It's responsible for receiving events from the sensor, adding additional context and metadata to the event then passing that event on to the cloud infrastructure via the Broadband gateway.

If the IoT gateway does not have internet connectivity, it can cache the events so that when connectivity is restored, cached events can be transmitted.

The IoT gateway may support one of many communication standards like Ethernet, 802.11 a/b/g/n/x (WiFi), universal serial bus (USB), Bluetooth, Bluetooth LE, cellular, near-field communications (NFC), Z-wave, ZigBee or other proprietary standards.

If WiFi or ethernet is used care must be taken to allow for the segmentation of data via routing of the sensor mesh network traffic from other broadband traffic in the home. Many IoT gateways include their own WiFi APs which can be used to establish a hidden, secure SSID for the sensor mesh network.

The IoT Gateway is also responsible for enabling the provisioning and configuration of the sensors, along with the monitoring of the sensors. The gateways have physical or API mechanisms to enable the pairing process of sensors that use NFC, Bluetooth, Z-Wave or ZigBee protocols. The pairing process needs to include a way to name and identify the location of the sensors during the pairing process. It must keep track of the sensor identification and serialization and include that data with the event messages it passes on to the cloud services. The Gateways should collect and report the battery level and online status of sensors and report this information up to the cloud monitoring services.

### **1.3. Broadband Gateway**

The Broadband Gateway provides internet connectivity to the IoT Gateway. As communication from the IoT Gateway is secured with encryption the IoT Gateway can communicate on the customer's local area network. If quality of service (QoS) is desired, it may be necessary for the IoT Gateway to communicate to the Broadband gateway via an isolated network.

It is important that the availability of the Broadband Gateway is monitored closely. If the Broadband Gateway becomes unavailable, real-time monitoring the customer is not possible. Ideally if there is an outage that affects the Broadband Gateway, all stakeholders related to the monitoring of the customer should be notified.

### **1.4. Cloud Infrastructure**

All the events created in a home are collected and analyzed in the cloud infrastructure. Each event is correlated to a specific customer in a specific location to a specific sensor at a specific point in time. This event stream is collected and analyzed with machine learning to establish ADL patterns. It is important that a large sample set of normal behavior is collected so that baseline patterns can be established. Machine learning will continue to adjust the base line patterns but flag deviations that are out of the ordinary.

Not only can out-of-the-ordinary deviations be flagged, but hard rules can be set such as if somebody is not up and moving by 9 a.m., an alert condition is reached.

When an alert condition is reached, the infrastructure can automatically begin notification escalations based on a real set that's established for the customer. This escalation can start with notifying local Caregiver, move on to notify distant loved ones and if there is no response from the escalations ultimately result in a notification to Emergency Services to perform a wellness check.

It is imperative that the cloud infrastructure is highly reliable and continually monitored. Services should be deployed across multiple data centers and geographically diverse locations, with load balancing, fail over and disaster recovery.

## **2. User Experience**

While the functionality of the aging in place platform depends upon cutting edge technology for sensing, artificial intelligence and machine learning, all of that innovation is wasted unless it can be translated into a compelling user experience. One that is intuitive and appropriate for the target user. One that is informative enough to become part of the user's daily routine, or efficient enough to be used all day long.

or intuitive enough to be used weekly, monthly or even less frequently and yet, not so “noisy” as to drive the user to a point where the platform becomes a source of irritation for the user.

Above all, the user experience needs to be flexible. Plain old voice telephony must be covered, SMS messaging must be supported and mobile push notifications must be available. Smart speakers should be available to support voice-first interaction with the system where appropriate.

In the independent living ecosystem, there are a variety of aging in place system user types to be addressed:

## **2.1. Subscriber**

For purposes of this paper, the term “subscriber” refers to the person for whom the aging in place platform is being used to establish a safe and healthy independent living situation. In the eyes of the cable operator this is a subscriber. In the eyes of a medical provider, this is a patient. In the eyes of the insurance company, this is a beneficiary. In the eyes of a home healthcare agency, this is a client. In the eyes of the senior community, assisted living facility, skilled nursing facility, or memory care facility, this is a resident.

The subscriber interacts with the system in both active and passive manners. Active interaction includes verbally communicating with the caregiver network and call center agents, responding to system-initiated queries, setting home/away status and engaging in cognitive analysis activities. Passive interaction happens by the subscriber going about their normal daily routine.

An important item to consider is the subscriber’s technology aptitude and the technology that they have access to. They may or may not have a landline. They may or may not have a mobile phone and if they have a mobile phone, it may or may not be a smart phone. They may or may not have a tablet and they may or may not have a computer.

The subscriber themselves may have hearing, seeing or speaking impairments which will also impact their relationship with the system. Both a mobile app and a web page user interface should be available to a subscriber to support their personal needs.

## **2.2. Family Caregiver**

A family caregiver is someone who provides emotional, financial, nursing, social, homemaking, and other services on a daily or intermittent basis for a family member, a friend or a neighbor. Most family caregivers volunteer their time, without pay, to help with the care needs of a loved one.

The family caregiver user experience is always an active interaction and involves managing their preferred communication method, responding to alerts, managing the rule thresholds (for those with administrative privileges), requesting and reviewing activity reports, limited infrastructure management and managing caregiver priorities (for those with administrative privileges).

In most situations, it is expected that the family is involved in the aging in place care plan. In fact, patient and family involvement is one of CMS’s Quality Strategy goals. A robust aging in place system will support event management and event escalation through a series of caregivers. After all, this is a volunteer position and Family Caregivers typically have their own immediate family responsibilities which may prevent them from always responding to an alert.

Both a mobile app and a web page user interface should be available to a Family Caregiver to support their personal needs.

### **2.3. Professional Caregiver**

Professional caregivers are hired to provide care for the subscriber. These caregivers can provide medical or non-medical care in the home. Frequently, payment for services is the only difference between a Professional Caregiver and a Family Caregiver. Therefore, the user experience is the same between a Family Caregiver and Professional Caregiver.

Both a mobile app and a web page user interface should be available to a Professional Caregiver to support their personal needs.

### **2.4. Institutional Caregiver**

The Institutional Caregiver is a Professional Caregiver working in a commercial care facility environment. Independent living, assisted living, skilled nursing and memory care are the typical commercial care facility types. In these environments, the user experience at the individual resident level is the same as for a Family Caregiver; however, there is also a dashboard to support a centralized monitoring and dispatch capability. In telecom terms, this dashboard looks much like a display you might see in a Network Operations Center (NOC). This dashboard allows the user to see a high-level status of each resident and comprehensive view of the active alerts.

Institutional Caregivers are typically equipped with tablets.

### **2.5. Call Center Agent**

The Call Center Agent serves as the backstop to the aging in place care escalation model. In the event that none of the caregivers in the subscriber's call escalation tree are able to respond to an alert, it will be routed to a 24/7 staffed call center. This call center must be staffed with agents that are more health-oriented than what an MSO employs. Here the Call Center Agent will attempt to reach the subscriber one more time before engaging first responders. Because of the potential financial charges, the subscriber needs to opt-in for the type of response that the call center is to take with respect to first responders.

Call Center Agents would be considered super users which means they would have the same capability as a Family or Professional Caregiver and also to serve as a help desk for system users. Additionally, they would have responsibilities with respect to system administration to address equipment issues, billing issues and network outages.

Call Center Agents would be desktop users.

### **2.6. Healthcare Provider**

One of the great benefits of this technology is the additional insight into patient physical and cognitive wellness that is available to a healthcare provider. The system provides a time-series of evidence for the healthcare provider to corroborate the observations from the patient assessment. With the additional insight afforded to them, Healthcare Providers are better able to catch symptoms early on before they develop into something more serious. Continued interaction with the patient also allows the Healthcare Provider a chance to evaluate how well treatments are progressing, and to change medication and visits as needed.

While it is not expected that a Healthcare Provider would be an active member of the patient's escalation call tree, that decision should be left up to the patient and their family. In this scenario, the user experience would be just like a Family or Professional Caregiver. Rather the expected Healthcare Provider interaction with the system will be through graphs and reports showing event occurrence and trends.

The subscriber or Family/Professional Caregiver would need to permit the Healthcare Provider to access the subscriber's information and this could be delivered in a push or pull model. Certain health plans reimburse healthcare providers for reviewing digitally collected data as part of a treatment plan. Scheduled reports could be delivered to the Healthcare Provider. Alternatively, the Healthcare Provider could pull up health and wellness reports on-demand during an office visit or as part of the burgeoning house call services for instance.

Aging in place data can also be fed into the patient's electronic health record for long-term trend analysis.

## **2.7. Case Manager**

Insurance companies, hospitals, and home health providers may all assign a Case Manager to monitor a patient. Case management is designed to provide for a patient's needs while controlling costs. A Case Manager evaluates what services are considered medically necessary, and works with different service providers to ensure that the required services are being given in the proper setting.

Once the patient is discharged from the hospital, Case Managers often must navigate complex care needs in a remote setting, which may involve different services being offered at different times by different providers. Ensuring that procedures and services fall under insurance policy coverage and will be paid for is a critical function that requires evidence to make proper decisions. It also involves educating patients on lifestyle adjustments, how to take medication, and when to come in for follow up appointments.

Chronic illness management, post-acute care and readmittance reduction are all activities of a Case Manager that depend upon evidence-based ADL and IADL measurement. The aging in place health and wellness measuring and monitoring system is designed to provide that information. As with the Healthcare Provider, scheduled reports could be delivered to the Case Manager or they could request reports on-demand.

## **3. Key Metrics**

Because the platform is dealing with subscriber health and wellness, availability of the platform is critical. Making sure the sensors, IoT Gateway, Broadband Gateway are operating correctly, that systems have internet connectivity, and the cloud platform's services are fully functional is critical to providing a safe service.

MSOs are quite familiar with monitoring their infrastructure and outside plant to proactively address outages. However, when problems arise with customer premise equipment and services usually contact initiated by the customer drives service calls.


When providing a platform that deals with health and wellness a proactive approach needs to be taken to ensure that things are working correctly on premise. In fact, unlike other services MSO offer the subscriber is unlikely to know if part of the service is not operating correctly.

### 3.1. Sensors Metrics


As the majority of sensors are battery-powered, one of the most important metrics to track is battery level. It is important to understand what metrics reflects the imminent failure of a battery. It is also important to understand the average lifespan of a battery of a device based on the number of events generating so the projection can be made as to when the battery should be serviced. Servicing batteries infrequently can result in sensor failure. Servicing batteries too frequently is a waste of money and can be an annoyance of the subscribers.

#### 3.1.1. Sensor Battery Levels

Each sensor model and vendor may report battery levels differently. They may return voltages, a percentage of 'life' remaining, or sometimes values that are unclear or not well documented. It is important to characterize battery performance of a specific type and model of sensor to understand how to set thresholds for battery replacement. Below are graphs from two different sensors made by two different vendors showing battery life:



**Figure 3: Multi Sensor Battery Life**



**Figure 4: Motion Sensor Battery Life**

In both graphs the x-axis represents time and the y-axis represents battery level. In the first graph you'll notice that when battery life reaches about 50% there are spikes showing an increase in battery level. This is because the sensor is shutting off due to low voltage. Once the cutoff voltage is reached, there is no current draw on the battery and its voltage level increases high enough over the cutoff threshold so that the sensor starts operating again. Soon after that, the sensor again reaches a low voltage cutoff and the sensor shuts off. This cycle continues until the battery does not have enough voltage to raise above the cutoff voltage and restart the sensor. When sensors behave this way it's important to track sensor battery level history, and not to use individual samples to determine when the battery should be replaced. It is possible that a single sample can be misleading as it can indicate a high battery level while the sensor is actually in a failure mode.

The graph on the right shows a sensor with a more predictable behavior. You can see that there's a relatively linear decrease of battery level overtime until the battery level reaches about 15% at which point there's a quick drop off.

In both cases battery levels will reach a point where the sensor is inoperative and battery levels cannot be collected. It's important to be able to identify if a sensor has become inoperative due to battery failure or due to some other reason. Not having battery history on a per sensor basis makes that determination impossible.

It is also important to be able to predict battery failure and ensure that batteries can be delivered and installed on premise either by a technician or by caregiver before the sensor fails.

### **3.1.2. Other Sensor Failures**

Sensors may stop operating properly for reasons other than battery failure. They could have been moved from where they were installed, or even reoriented so that they are not monitoring the correct area. It is hard to know if events are not being received because of a sensor failure or because there is no activity the area the sensor is intended to detect.

Many sensors are not polled for events, but generate events when thresholds set at the sensor are met. The lack of events could be due to either sensor failure, or that there is nothing in the environment making events. If a motion sensor is pointed at the ceiling, it will still respond to queries from the IoT, report battery level and appear healthy, but it will not generate events. This is effectively a sensor failure.

It is possible to correlate multiple events across multiple sensors related to a common action in the environment. Machine learning can be used to pick up these patterns. But generally, if someone moves from the bedroom to the kitchen one would assume they would see events in both bedroom and kitchen. If events are being received in the kitchen but never being received in the bedroom it is possible that the bedroom sensor has failed. Is important to generate a set of metrics and to understand the general premise layout to correlate activities across multiple sensors to help determine sensor failure.



### **3.2. IoT and Broadband Gateway availability**

Should an IoT Gateway or Broadband Gateway become unavailable. It is important to quickly assess whether the outage is due to a problem in premise or due to an event on the operator's network. Metrics on gateway availability must be continually tracked and when the gateways become unavailable the cause of the failure must be determined quickly. This so that the correct resources can be dispatched to re-establish communication. It is also imperative to notify caregivers when outages occur, to allow them to take actions to reduce risk to the subscriber.

### **3.3. Escalation and Incident Response**

It's useful to track the number of escalations needed before a satisfactory response is achieved on a subscriber by subscriber basis and by population. The main point of escalations is to have someone check on the subscriber to ensure that they are safe and healthy. This may be a phone call or an in-person check. It could even be an escalation to emergency services. If over the course of several separate events, it takes multiple escalations before subscriber safety is determined, it may be an indication that the subscriber may not have the correct escalation order or the correct people involved in escalations. Metrics can help make recommendations to the subscribers to examine and adjust their escalation paths.

It is also important to understand how quickly people in the escalation chain take to respond. A balance must be struck between giving a person a chance to respond to an escalation versus how long is it safe to wait before the next person in the escalation path is contacted. Escalate too fast and there is a risk in repeatedly causing stress to the people responding to the escalation (i.e.; caregiver fatigue). Escalate too slow and there is risk that the subscriber is not getting the help they need. It's important to understand escalation and response timing for specific users and for populations to help set the correct thresholds.

## **Conclusion**

### **“Necessity is the mother of invention”**

With more aging adults needing care than there are caregivers to provide it for them, technology must be leveraged to provide that care. Broadband providers have long sought to be a part of this solution and the timing is finally right for that to happen with the maturation of IoT sensor technology and cloud computing. This paper has presented the concept of a cloud-based system of unobtrusive monitoring and measurement that is aligned with Medicare's value-based programs and quality strategy. These initiatives seek to financially reward those that deliver better care, better health and lower cost. Put more simply, Medicare is moving towards reimbursement of a service that utilizes MSOs' broadband service.

By starting with a simple wellness system, adoption by the subscriber will be easier to overcome, technical and logistical deployment challenges will be easier to overcome and reimbursement models will be easier to understand. Passive Remote Patient Monitoring is designed to advance to Remote Patient Monitoring and then to full Telehealth as needs require.

MSOs are an ideal delivery mechanism for health and wellness monitoring systems. Older adults, and their caregivers, are already comfortable and familiar with basic broadband technology. Additionally,

those caring for older loved ones tend to embrace new technology and are increasingly open to using it in their personal health and wellness care.

The model PRPM platform is non-intrusive, protects patient privacy, provides caregiver peace of mind, is accessible, easy to use, interactive, facilitates better outcomes, optimizes patient/provider interaction, and reduces overall health and wellness costs. It does not require MSOs to turn themselves into telehealth providers.

# Abbreviations

ADL	Activities of Daily Living
CMS	Center for Medicare and Medicaid Services
HIPAA	Health Insurance Portability and Accountability Act of 1996
IADL	Instrumental Activities of Daily Living
IoT	Internet of Things
MSO	Multiple System Operator
NFC	Near-Field Communications
NOC	Network Operations Center
OSS	Operational Support System
PRPM	Passive Remote Patient Monitoring
RLT	Roper-Logan-Tierney Model of Nursing
RPM	Remote Patient Monitoring

## Bibliography & References

*Roper-Logan-Tierney Model of Nursing*; Winifred W Logan, Alison J Tierney, Elsevier Health Sciences, 2000. ISBN 0702041076, 9780702041075

*Instrumental activities of daily living*; MP Lawton, EM Brody, U Médecin. The Gerontologist, 1969.

<https://www.aarp.org/content/dam/aarp/home-and-family/personal-technology/2016-01/2016-Caregiving-Innovation-Frontiers-Infographics-AARP.pdf>

<http://westviewnursing.com/by-2030-u-s-demographic-milestone-begins-to-lower-caregiver-ratio>

<https://www.census.gov/newsroom/facts-for-features/2017/cb17-ff08.html>

<https://www.mylifesite.net/blog/post/so-ill-probably-need-long-term-care-but-for-how-long/>

<https://longtermcare.acl.gov/costs-how-to-pay/costs-of-care.html>

<http://www.aarp.org/home-family/caregiving/info-08-2013/the-aging-of-the-baby-boom-and-the-growing-care-gap-AARP-ppi-ltc.html>

<https://www.mylifesite.net/blog/post/so-ill-probably-need-long-term-care-but-for-how-long/>

<https://www.catman.global/how-best-buy-could-soon-dominate-a-28-billion-market>

# Opting-In: Designing Privacy Tracking for Consumer Confidentiality & Cryptographic Assurance for Enterprises

A Technical Paper prepared for SCTE•ISBE by

**Brian A. Scriber**

Distinguished Technologist and VP of Security Technologies  
CableLabs  
858 Coal Creek Circle, Louisville, CO 80027  
@brianscriber  
b.scriber@cablelabs.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Privacy and International Policy Initiatives.....	3
1. Defining Protected Data .....	3
2. De-Identification of Data.....	4
3. Opt-In, Opt-Out, and Explicit Consent.....	4
4. Data Portability.....	4
5. Right to be Forgotten and Right to Deletion .....	4
6. Data Provenance .....	5
7. Privacy and the Network Operator.....	5
8. Costs of Non-Compliance .....	5
9. Revenue Opportunities.....	5
Data Modeling .....	5
10. Databases .....	5
11. Adding Data Protection .....	6
12. Data Model Impact of Adding Protection.....	7
Incurring Costs Related to Data Model Changes.....	8
13. Technical Transition.....	8
14. Operational/Procedural Transition .....	9
15. Privacy Impact Areas .....	10
Solution Options .....	10
16. Don't Collect Private Info in the First Place .....	10
17. Modify the Existing Data Models .....	10
18. Add Separate Data Model for Privacy.....	10
19. Restrict Access to a Centralized Data Store .....	11
20. Restrict Access but in a Decentralized Store with User Data Self-Sovereignty .....	11
Conclusion .....	12
Abbreviations.....	12
Bibliography & References .....	12

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 GDPR Definition of Personal Data.....	3
Figure 2: HIPAA Privacy Rule .....	3
Figure 3 Protected Data in PIPEDA .....	4
Figure 4 Conventional Database Entity Relationship Data Model .....	6
Figure 5 Extending the Data Model for Sensitive Data.....	7

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 Technical Transition Costs .....	8
Table 2 Operational/Procedural Transition Costs .....	9

# Introduction

Privacy, particularly consumer privacy, has lived a dynamic existence over the last decade. Consumer views have changed as has the regulatory environment. With GDPR in the EU, PIPEDA in Canada, and with CPA in California, there is an immediate need for technical solutions to efficiently run our businesses in this strong regulatory environment. What if you and your department could help enable a revenue opportunity in this space as opposed to just mitigating regulatory risk? Join us to explore how self-sovereign identity and opt-in/opt-out tracking can work hand in hand using some of the nascent tools in cryptography and software development. We will explore transaction signing, distributed verification, collaborative acknowledgments, time synchronization, user-directed sharing of protected information, the Right to be Forgotten, and cryptographic key distribution systems enforced by smart contracts. A key part of this paper investigates how data analytics and privacy can coexist without one working opposite the other through the application of advanced key management techniques, zero knowledge proofs, and smart contracts.

## Privacy and International Policy Initiatives

Privacy has regional meaning and with different practices, the legislative and compliance directives differ as well. There are some themes that can be extracted from the California Consumer Privacy Act (CCPA) in the United States, the General Data Protection Regulation (GDPR) in the European Union, and PIPEDA in Canada. Each is unique, but a policy that addresses several of these could be successful for companies deploying capabilities in each of these regions. Those key attributes include the following:

### 1. Defining Protected Data

As mentioned in the intro to this section, protected data has differing definitions based on region.

“Personal data is any information relating to an individual, whether it relates to his or her private, professional or public life. It can be anything from a name, a home address, a photo, an email address, bank details, posts on social networking websites, medical information, or a computer’s IP address” -- GDPR

**Figure 1 GDPR Definition of Personal Data**

The United States has a couple differing places to look for data protection; with state-level initiatives, look for these definitions to remain a bit fluid over the near future.

“Individually identifiable health information held or transmitted by a covered entity or its business associate, in any form or medium, whether electronic, on paper, or oral” – Privacy Rule, HIPAA

**Figure 2: HIPAA Privacy Rule**

In Canada, PIPEDA is the prevailing legal doctrine related to privacy and it defines the protected data

“Data that contains any factual or subjective information, recorded or not, about an identifiable individual. This consists of not only personally identifiable information (PII) such as name, age, ID number and ethnicity or medical records, employee files, credit records and so on, but also opinions, evaluations, comments, social status and disciplinary actions.” -- PIPEDA

**Figure 3 Protected Data in PIPEDA**

thus (it goes into further detail around sensitive data not protected explicitly by PIPEDA as well):

## **2. De-Identification of Data**

According to the GDPR Recital 26, “Pseudonymised data is still considered personal data” and the US Department of Health and Human Services goes into great detail over de-identification of data and the complexities therein<sup>1</sup>. The undercurrent of all of this is that simply taking the name out of a data set, and optionally replacing it with a number/key/identifier, is insufficient for privacy protections because the collection of data held may still be enough to identify participants and to enable the ability to re-identify them.

## **3. Opt-In, Opt-Out, and Explicit Consent**

The ability for consumers to opt out of the use of their personal data and to be most generally compliant, a policy of opting in for use of that data is a path some are taking as explicit consent, such as in a clear affirmative action, for data use is a requirement of GDPR.

## **4. Data Portability**

Data portability is a requirement that looks for consumers to be able to extract the data related to them, including the personal information, but to do so using a “standard” mechanism. The implication of this requirement is that other systems should also be able to consume this format as well. Personal information, the definition of what constitutes Personally Identifiable Information, Protected Health Information, and Personal Data differs between regions as well. That difference makes the standardization of this data complicated. Heterogeneous semantic mapping between these, and the differences in how data is stored, highlight inconsistencies (e.g. Social Security Numbers in the USA and Social Insurance Number in Canada).

## **5. Right to be Forgotten and Right to Deletion**

In the EU, there exists a Right to be Forgotten (RTBF) and in the CCPA there is a “right to deletion” RTD which were both created for a consumer to be able to exert a level of control over any future use of their data. When RTBF or RTD are exercised, the personal information related to that consumer is to no longer be used.

## 6. Data Provenance

The requirement to be able to share with a consumer where a given article of data originated, how the possessing company came to own it, and potentially whether or not it was purchased, guessed, inferred, or otherwise arrived at is a part of this privacy landscape that this paper will show may drive some of the most significant changes. The granularity of data, the mapping to origin metadata, and the combination of technology and process required to deliver that information are going to drive costs in enterprise development.

## 7. Privacy and the Network Operator

The GDPR states that, as shown in Figure 2., a computer's IP address is part of the Personal Data protection. As a network operator who bills for connectivity services, this is a critical aspect of the data used throughout operational systems. While this paper is not offering legal advice, Article 49 of the GDPR does have derogations and special conditions where there is a contract in place or in service of public interest.

## 8. Costs of Non-Compliance

While the above technical implications and process flows for potentially protected data collection are potentially costly, so too are the potential penalties for non-compliance with the laws. GDPR has potential fines of up to €20 million or 4% of annual turnover, whichever is greater. The CCPA has a \$7500 cap per intentional violation (\$2500 cap for unintentional violations), but the definition of a violation is likely<sup>ii</sup> to be interpreted to mean per incident per consumer in line with the data breach class action section of the law.

## 9. Revenue Opportunities

While the costs are high, there are opportunities for companies who find the path to successfully navigate the privacy concerns, that can scale their solutions, and which can offer such solutions as shields to the misuse of protected data while simultaneously providing a similar protection for enterprises with legitimate needs for use of data.

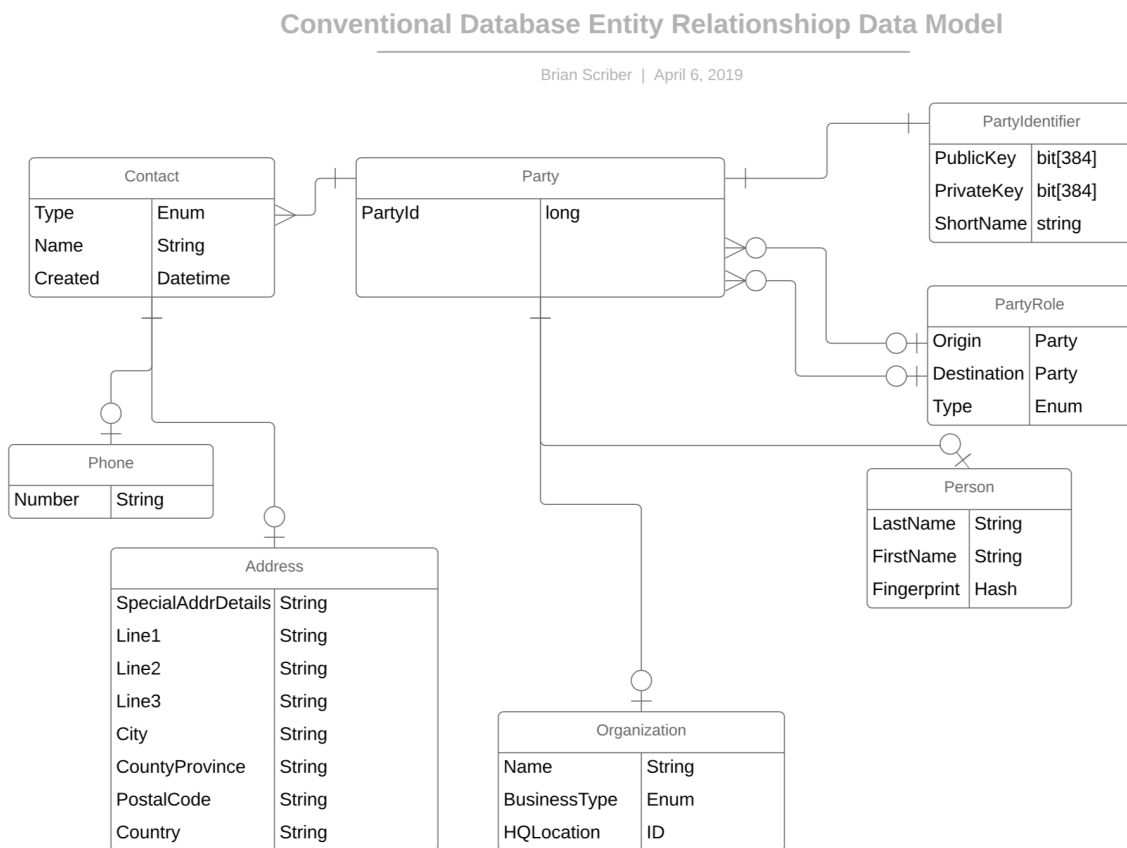
# Data Modeling

## 10. Databases

Since we are talking about requirements on data, the important place to look inside any enterprise is at the database level. Several databases will exist within a typical enterprise, including network operators. There are customer support databases, billing databases, service databases, advertising databases, operational databases, data warehouses, reports that utilize data, online systems, back-end systems – each with their own databases to support their missions. Data is pervasive, and it is important to protect it. Current data models follow a similar format to the ERD shown in Figure 4 Conventional Database Entity Relationship Data Model. This model shows the relationships between a person, their means of contact,



and any organizations they belong to, among a few other additional relationships and details.



**Figure 4 Conventional Database Entity Relationship Data Model**

## 11. Adding Data Protection

Adding the elements described by the confluence of data protection regulations to the already existing data models means exploring how to add at least the following eight elements of metadata:

- Source (what is the provenance of this data?)
- Date Entered (when did it make it to this database?)
- Expiration (when can this data no longer be used?)
- Allowed Uses (for what can this data be used?)
- Sensitivity (what protections must exist and what authorization is required for access to this data?)
- Explicit Consent (was this permission for use granted explicitly? When? By whom? In what mechanism?)
- Revocation (has the right to use this data been terminated?)
- Right to be Forgotten (can this data be completely removed? Remove all null references to these rows as well)

Reconciling that with existing data models can be problematic. To further reduce scope of the above already simple data model example, this analysis will focus exclusively on the Person table.

## 12. Data Model Impact of Adding Protection

The Person table from the prior diagram is now shown in the upper left hand corner of Figure 5 Extending the Data Model for Sensitive Data. The expansion of two of those three fields, LastName and FirstName make up the rest of the example. The “Fingerprint” aspect of the prior table is included to show that one of the greatest steps an enterprise can take is to ensure protection of sensitive data is to not collect or store sensitive data. The collection of biometric information about consumers is dangerous because an enterprise cannot issue a new fingerprint or offer fingerprint-monitoring-services free for a year to anyone impacted by a breach. It’s simply safer to not collect or store some information.

No longer can LastName simply be a string. The need now is to create a StringProtectedDataElement and reference it from the Person table. The new table needs to be able to have the Id being referenced and the actual data element, but since we want to reuse this table for multiple lengths of strings, we may have to be creative about how to not end up with a sparsely populated column/field where only a small portion of the data we set aside for storage is actually used by the data going into it (e.g. if most first names are only five characters long on average, and most last names are an average of 14 characters long, then using the same table to store only those two fields could result in an average of nine characters of unused storage capacity which can be costly in its own right).

Extending the Data Model for Sensitive Data

Brian Scriber | April 7, 2019

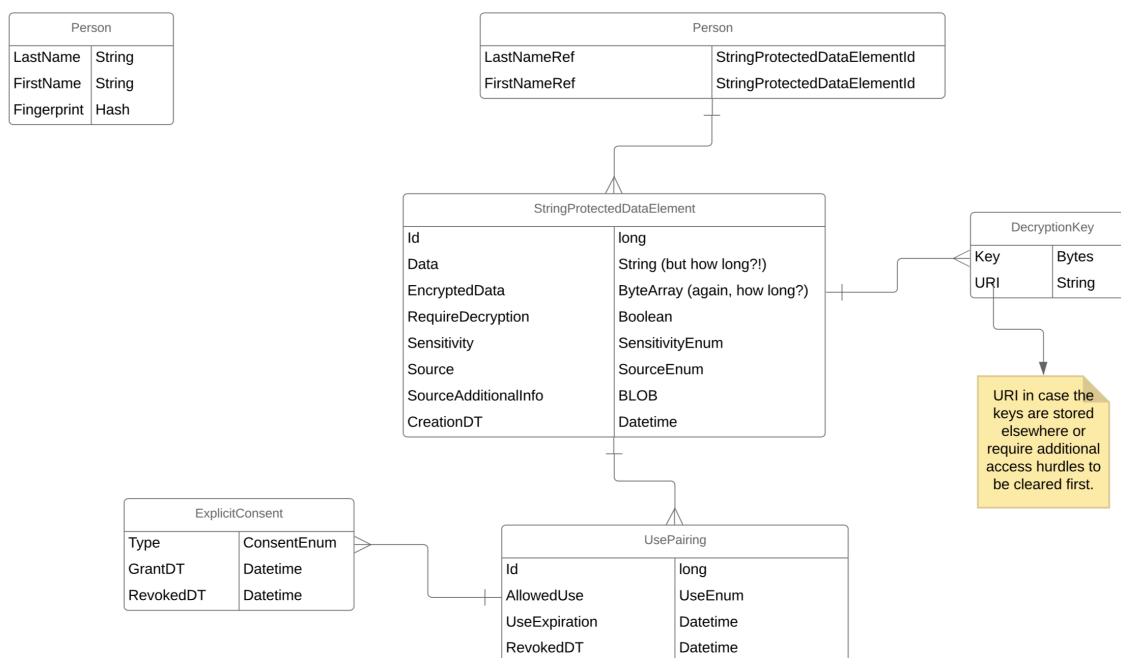


Figure 5 Extending the Data Model for Sensitive Data

This new table will also need to be able to store encrypted data if that data is designated as sensitive. Many databases have more efficient ways to deal with this than what is shown above, but the example here demonstrates that the requirement for encryption and the actual payload still need to be stored. Sensitivity may not be binary – the data may not be exclusively either sensitive or non-sensitive. There may be list of things used to qualify that sensitivity (e.g. “Top Secret”, “Confidential”, “Company-Only”,

“Partners and Company”, or even combinations like “Authorized Advertisers but Otherwise Confidential”, etc.). The source for this data needs a similar enumeration (e.g. “Company A”, “Employee Survey”, “Data Aggregator B”, etc.). It’s not always clear what additional metadata might be required to help differentiate between sources, so a binary object has been created for this additional info. Each field from this string data also needs to have the date it was entered into this database, and a relationship to the allowed uses of the data, if those uses have been expired or revoked, and if there was an explicit consent provided. If there is explicit consent, there is now the capacity to know when it was granted and when it was revoked for each use (e.g. you can use my name for two years related to the bicycle contest, but you can only use my name for 3 months related to the ice cream flavor game). This data model gets complicated when one realizes that this is only for a tiny portion of only one of the databases for only the first and last name of the user.

## Incurred Costs Related to Data Model Changes

### 13. Technical Transition

The technical transition will typically need to incorporate the aspects defined in Table 1 Technical Transition Costs. The right-most column in that table lists some of the typical roles that will be involved in the exercise of data management identified in the middle column. The three primary technical categories are the Data Model itself, the Software which directly accesses that data model, and then the Quality Assurance and Business Analysis required for each system. Project management expertise will likely be required at all phases of this work.

**Table 1 Technical Transition Costs**

Data Model		
	New Data Model	Data Modelers, Database Administrators, Application Engineers, Architects
	New Database Tables	Data Modelers, Database Administrators
	New Object-Relational Mappings	Database Administrators, Application Engineers, Architects
Software		
	New Allowed Use Checks	Software Design Engineers, Application Engineers, Architects
	Copying and Caching Validation	Application Engineers
	Integration APIs to be Updated	Systems Architects, Systems Engineers
Quality Assurance & Business Analysis For Each Group		

Review Requirements	Business Analyst, Quality Assurance (QA) Engineer, Architect, Systems Architect
Test/Confirm Privacy Requirements Met	Business Analyst, QA Engineer
New Bugs and Bug Fixes	Application Engineer, QA Engineer, Configuration Management Engineer

## 14. Operational/Procedural Transition

The technical transition is not the only part of a project like this, the impact to the operational, procedural and legal aspects of the enterprise should not be underestimated. Again, project management involvement as well as the office of the Chief Privacy Officer (a requirement in the EU, but optional elsewhere) are presumed for all of the following activities laid out in Table 2 Operational/Procedural Transition Costs.

**Table 2 Operational/Procedural Transition Costs**

Procedures		
	Notify Downstream Users of New Data Contracts	Legal, Marketing, Application Engineers, Business Analysts
	Data Warehousing and Reporting Data Curation Requirements	Data Warehousing Architects, Data Reporting Engineers
	Limiting Access to DBAs, Operational Teams, Strict Access Controls	Security Engineers, Business Analysts, Legal
Legal		
	Revocation (primary versus secondary/tertiary users of data)	Legal, Business Analysts, Application Engineers, QA Engineers
	Audit (data model, access control, even with decryption keys requires trust that regulators may still question)	Legal, Compliance/Auditor, Audit Committee, Business Analyst
New Customer-Facing Obligations		
	Exporting Data	Data Engineers, Legal, Standards Compliance Engineers
	Right to be Forgotten	Business Analyst, Data Engineers, Legal, Standards Compliance Engineers

Data Provenance Reporting	Systems Architects, Data Engineers, Legal, Standards Compliance Engineers
---------------------------	---

## 15. Privacy Impact Areas

Across the enterprise, the privacy and compliance changes will touch many departments if not all of them. The following list identifies several that will need to be kept at the front of mind while engaging in a holistic compliance effort:

- Online Account and Support Systems
- Production Applications
- Network/Abuse Monitoring
- Legal and Risk Management
- Billing and Accounting
- Purchasing (Contracts)
- Data Warehousing
- Marketing and Sales
- 3rd Party Data Clients & Partnerships

## Solution Options

There are several options available which can each address some of the concerns raised above. This paper presumes outright that a complete green-field implementation of all systems and integrations is cost-prohibitive. Given that, the following subsections each explore some of the pros and cons of different and progressive concepts. There are other options and other combinations that may make sense, but

### 16. Don't Collect Private Info in the First Place

This is likely the immediate direction many enterprises will wish to explore. How to minimize data collection and use/store only the bare minimum. In this model they will still need to address the privacy implications of the data they do collect.

### 17. Modify the Existing Data Models

The data model sections above go into the impact of what is likely to feel is the compromise solution of updating the data models of a few key systems, but the costs associated with this were explored in depth, above. This noted, there may not be easy shortcuts for some of the complications imposed by the privacy landscape.

### 18. Add Separate Data Model for Privacy

Looking at the data, it might appear that some of the data can be extracted into a privacy model, but even doing this would require the integration points that were highlighted in the “Modify the Existing Data Models” approach defined above.

## 19. Restrict Access to a Centralized Data Store

This approach puts all the privacy-related data in a single place, but the issue here is that now there would exist a single point of failure for the entire enterprise and the risk of shutting down all operations while any delay or issue with this core system was fixed is likely prohibitive.

## 20. Restrict Access but in a Decentralized Store with User Data Self-Sovereignty

User data self-sovereignty is a relative new concept in the Privacy Enhancing Technologies (PET) space. The idea is that the owner of the data controls who has access to their information. All data for all systems is encrypted and specific grants are given to different systems who can then access the same decentralized network to request access. The result of the requests are pointers to where the data resides, and the request also enlists aid from the network in decrypting that data pending the requestor authorization. Participant nodes are all independent, but the data stores may have single points of failure for any single element of data.

The data self-sovereignty options are built around the concepts of using some of the tools available to us from the security space and apply those to the privacy domain. These tools include:

- Encryption
- Databases
- Protected File Systems
- Audit Logging
- Indelible Ledgers
- Hashing
- Smart Contracts
- Trusted Execution Environments (TEE)
- Cryptographic Grants for Information Use

The TEEs and Smart Contracts open the door for:

- Mutual Trust Environments
- Audited Transactions \*
- Authenticated Access: Read/Write/Create/Delete/Update/Notify
- Grants for Information Use (time-bounded, use restricted, access controlled)

Distributed Ledgers Add:

- Indelible Transactions
- Distributed Access
- Potential to Store Critical Protected Data in a Manner Accessible by
- Byzantine Fault Tolerance

Self Sovereignty Provides:

- Putting Protected Information into the Hands of the Owners of that Data
- Enables Explicit Consent, Access Grants, and Updates

# Conclusion

The costs of compliance with global privacy initiatives can be far greater than some of the assumptions that are being made about how to achieve this compliance and about the penalties that are likely to be levied. The future is likely going to be complex during the transition period from our legacy data models to those which protect private information, but the long-term future has options like that of data self-sovereignty and individuals owning their own data and using cryptography to protect it and grant appropriate access.

# Abbreviations

CCPA	California Consumer Privacy Act (California, USA)
ERD	Entity Relationship Diagram
EU	European Union
GDPR	General Data Protection Regulation (European Union)
PIPEDA	Personal Information Protection and Electronic Documents Act (Canada)
QA	Quality Assurance
RTBF	Right to be Forgotten
RTD	Right to Deletion
TEE	Trusted Execution Environment

# Bibliography & References

*EU General Data Protection Regulation (GDPR)*: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ 2016 L 119/1.

*EU Directive 95/46*: Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data, OJ 1995 L 281/31.

*Papers of the Article 29 Working Party (example)*: Article 29 Working Party, ‘Opinion 2/2003 on the application of the data protection principles to the WHOIS directories’ (WP 76, 13 June 2003), at 4

*Personal Information Protection and Electronics Document Act*: Amended April 1, 2019 <https://laws-lois.justice.gc.ca/PDF/P-8.6.pdf>

*EU Charter of Fundamental Rights*: Charter of Fundamental Rights of the European Union, OJ 2010 C 83/389.

*European Convention on Human Rights*: European Convention for the Protection of Human Rights and Fundamental Freedoms, Sept. 3, 1953, ETS 5, 213 UNTS 221.

*Judgments of the Court of Justice of the EU (example)*: Bodil Lindqvist, Case C-101/01, [2003] ECR I-12971 (ECLI:EU:C:2003:596), at para. 74.

*Council of Europe Convention 108: Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data 1981, ETS 108*

*NIST Privacy Framework: An Enterprise Risk management Tool*: National Institute of Standards and Technology, US Department of Commerce, April 30, 2019,  
<https://www.nist.gov/sites/default/files/documents/2019/04/30/nist-privacy-framework-discussion-draft.pdf>

---

<sup>i</sup> <https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html>

<sup>ii</sup> <https://www.clarip.com/data-privacy/california-consumer-privacy-act-fines/>



# **Cable Edge Compute: Transforming Cable Hubs into Application-Centric Cloud**

A Technical Paper prepared for SCTE•ISBE by

**Rajiv Asati**

Distinguished Engineer  
Cisco Systems  
rajiva@cisco.com

**Alon Bernstein**

Distinguished Engineer  
Cisco Systems  
alonb@cisco.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
1 Overview.....	4
1.1 What – Definition, Background and MEC Relation ! .....	6
1.1.1 MEC Background! .....	7
1.1.2 Cable Cloud vs. IT Cloud .....	7
1.2 Why – Benefits? .....	8
1.3 Where is the Edge? Centralized vs. Distributed ? .....	9
2 Cable Edge Compute – Application Functions .....	12
2.1 Infrastructure Use-Cases .....	13
2.1.1 (Virtual) RAN Functions .....	13
2.1.2 CBRS .....	15
2.1.3 Subscriber Edge/User Plane Functions (e.g. CMTS, BNG, PGW) .....	16
2.2 B2C Service Use-Cases.....	16
2.2.1 Gaming .....	17
2.2.2 LiveTV.....	18
2.3 B2B Use-Cases .....	18
2.3.1 IoT & Public Cloud Hosting.....	18
2.3.2 3 <sup>rd</sup> Party CDN .....	19
2.3.3 Video Surveillance .....	19
2.3.4 Security .....	20
3 Cable Edge Compute – Architecture .....	20
3.1 Infrastructure – Common Hardware.....	21
3.2 Infrastructure – Software (NFVI).....	24
3.2.1 Common Cloud Orchestration Platform.....	24
3.2.2 Orchestration Control Nodes combined with Compute Nodes .....	25
3.2.3 Deterministic NFVI Performance.....	26
3.2.4 Remote Storage with Optimization & Security .....	26
3.2.5 Remote Management & Monitoring.....	27
3.3 Network Transport Fabric.....	28
3.4 Automation, Orchestration and Assurance.....	29
4 Summary .....	31
4.1 Opportunities .....	31
4.2 Challenges.....	31
4.3 Recommendations .....	32
Abbreviations.....	34
Bibliography & References .....	34

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 Applications – Cloud Enabled vs Cloud Native .....	5
Figure 2 IaaS, PaaS, BaaS .....	5
Figure 3 Customer Experience Driving Applications Enablement Closer and Closer .....	6
Figure 4 Public Cloud Edge Locations - Example .....	10
Figure 5 Centralize or Distribute the Edge Clouds - Pros & Cons .....	11
Figure 6 Upstream Bandwidth could be ~30% of downstream Bandwidth.....	11
Figure 7 Example Application Functions for Edge Cloud Locations .....	12
Figure 8 Mobile RAN Evolution .....	13
Figure 9 5G Bandwidth in Access, Aggregation.....	14
Figure 10 Radio Access Network showing 5G Cell Site with mmWave .....	15
Figure 11 Subscriber Edge Functions, CUPS.....	16
Figure 12 Gaming Service .....	17
Figure 13 Managed Video/CDN .....	18
Figure 14 IOT and Public Cloud Hosting .....	19
Figure 15 Video Monitoring adds significant Upstream bandwidth Consumption.....	19
Figure 16 E2E Architecture Blueprint Showing Cable Edge Compute (CEC) .....	20
Figure 17 Network Platform - Abstraction is KEY .....	21
Figure 18 Hardware Infrastructure - Edge PoD SKUs.....	22
Figure 20 Virtual Forwarding Options.....	23
Figure 21 Converged Cloud Platform .....	25
Figure 22 Collapsed VIM Controller and Compute Nodes.....	25
Figure 23 Deterministic NFVI Performance Logic.....	26
Figure 24 No more Local Storage for VIM .....	27
Figure 25 No more local Management Nodes .....	28
Figure 26 Network Transport Fabric.....	29
Figure 27 Automation, Orchestration and Assurance .....	30

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 Factors that can Influence Distributed Edge Clouds.....	11
Table 2 Infrastructure Hardware - POD SKUs .....	23

# Introduction

The mobile industry has popularized and has already started embracing the concept of Multi-Access Edge Computing (MEC). Clearly, the same concept can be applied to cable networks to benefit wide range of use-cases, especially the ones that are latency or bandwidth sensitive. This concept drives placing time sensitive applications e.g. IoT and/or bandwidth hungry applications e.g. CDN/cache at the network edge or Hub sites, closer to the customer. There is a tremendous opportunity for Cable MSOs in transforming their Hub sites into Next Gen Application-Centric Cloud sites.

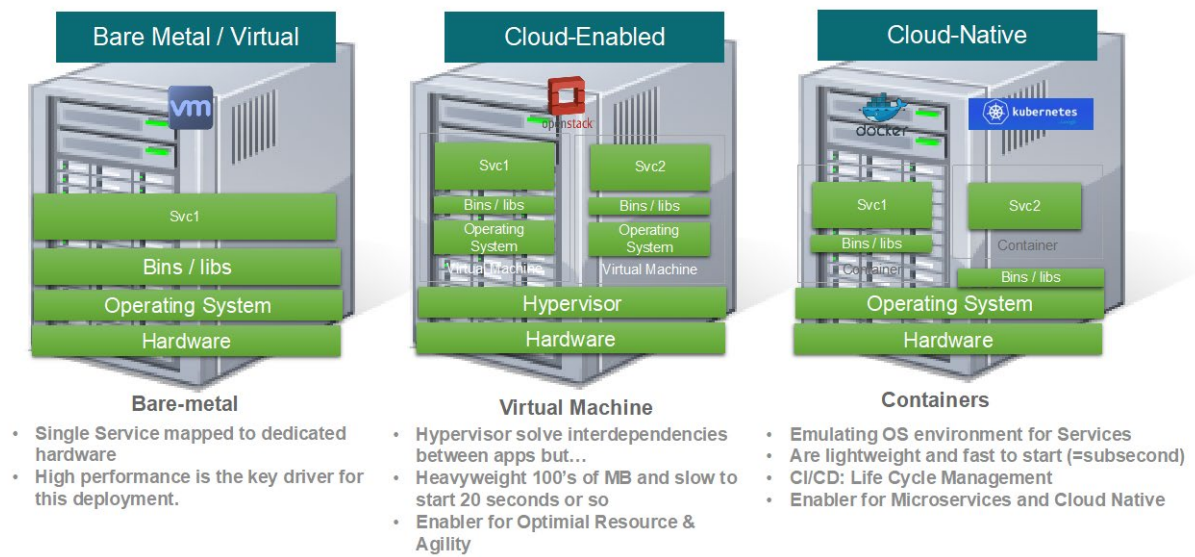
This paper defines the notion of Cable Edge Compute (CEC) along with ‘Edge POD’ and explores how it could be organized to serve a range of use-cases (e.g. what functions could be placed). The paper outlines the Architectural building blocks (suitable for CEC), their key attributes and captures unique opportunities, challenges and recommendations for Cable Operators.

## 1 Overview

Cable Multi System Operators/Service Providers are undergoing multi-pronged digital transformations.

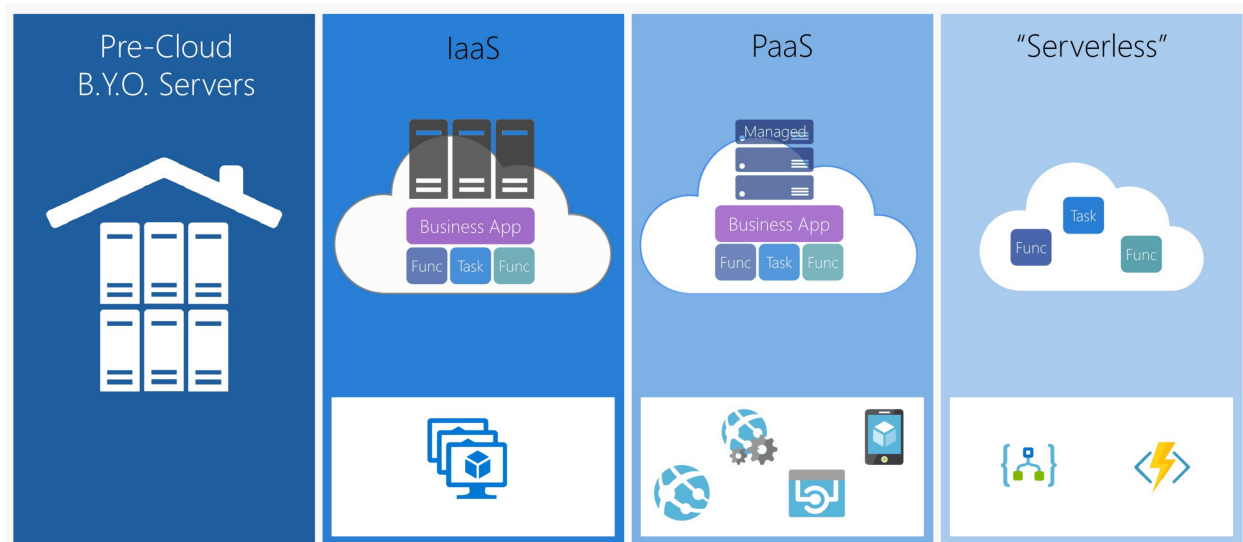
A few are attempting to change the game, not just play it better, by morphing services and solutions offerings (connectivity-centric and/or consumption-centric) that ultimately yield the best customer experience, while recognizing the fact that they have to serve not only the ones with eyeballs/eardrums, but also the ones without them i.e. machines such as Internet of Things (IoT) sensors. It is somewhat obvious that wireless and/or wired endpoints, whether deployed in few tens or millions, whether mobile or not, require not only the optimal seamless connectivity constructs, but also the most optimal consumption experience, given that these endpoints consume one or more services (where each service in turn comprise of one or more applications).

In fact, Applications have become foundational to growth and experience, no matter where they run – private cloud, public cloud, hybrid cloud, as long as the customer SLAs are satisfied. (note that two of the phenomena fueling applications evolution are micro-services and cloud native constructs that have enabled more abstraction than ever before, as illustrated in the figure below)



**Figure 1 Applications – Cloud Enabled vs Cloud Native**

These trends have been fueling the usage of as-a-Service such as infrastructure as-a-service (IaaS), function as-a-service (FaaS), serverless aka backend as-a-service (BaaS) etc. for executing the application functions, as illustrated in the figure below. One or more application functions could execute in one or more cloud locations in a scale-out manner for whatever time-period, independent of the location (centralized, partially distributed, fully distributed).



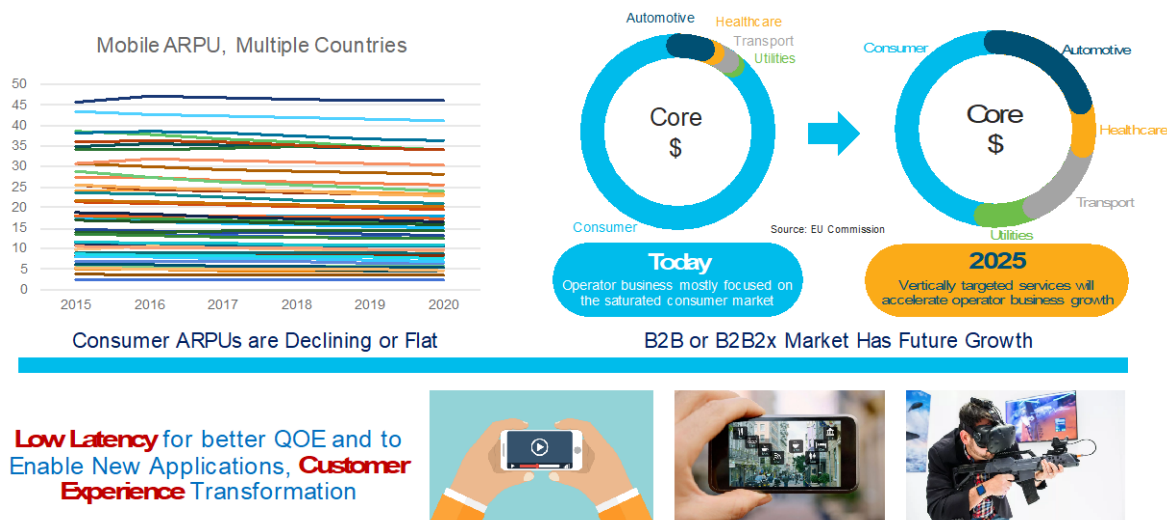
**Figure 2 IaaS, PaaS, BaaS**

Courtesy - <https://stackify.com/function-as-a-service-serverless-architecture/>

Application-Centric means that the network administrators manage a system for a set of applications rather than managing individual nodes like they did in the past.

The faster the Cable Operators leverage the **Application-Centric paradigm with flexible, distributed and intelligent network architecture**, the faster they get towards enabling superior customer experience. Applications could be related to internal usage (e.g. Infrastructure) or external usage (e.g. subscribers) or both.

The Operators could deploy Application-Centric Clouds in a centralized manner or distributed manner. It is important to point out that Cable Operators have had precious Hub Sites distributed across the footprints and they could be the ideal candidates for Application Centric Cloud. This is pertinent especially since the majority of revenue growth is now expected in B2B or B2B2x space, as illustrated in the figure below -



**Figure 3 Customer Experience Driving Applications Enablement Closer and Closer**

Minimizing the latency is now more important (than increasing bandwidth) for improving QoE, and by hosting applications closer to the customers can greatly help. Increasing/throwing more bandwidth data rate doesn't help much after a while. For example, browser application could load a webpage in around 3500msec with 200ms E2E latency, but in around 2000msec with 100ms E2E latency – 45% improvement. However, increasing the bandwidth bandwidth from 5Mbps to 10Mbps yields around 5% improvement in page loading experience. See Reference [9].

*Cloud Services Providers related to Residential and Enterprises e.g. CDN operators, Gaming etc. are pushing hard to get closer to the eyeballs and last-mile networks.*

## 1.1 What – Definition, Background and MEC Relation !

"Cable Edge Compute" is intended to represent "IT & Telco centric" cloud-computing capabilities at the edge of the MSO network and in close proximity to Cable subscribers.

Cable Edge Compute (CEC) is a form of Multi-Access Edge Computing (MEC) that is applied to Cable MSO environment.

Cable Edge Compute aims to improve subscribers' experience by cutting out the often long and imperfect network path between the subscriber's device and the location where the application they are accessing is hosted, as much as possible, in order to lower latency, increase reliability and improve overall network efficiency.

*The concept of placing computing power near the customer's devices with the primary goal of improving customer experience while reducing latency, backbone capacity, etc. and getting better scale and availability. Edge Computing nodes are on the outer region of the core network or its backbone. Almost any device with computational power that is near or at the customers' devices location can act as an edge computing device, as long as it's practical.*

### **1.1.1 MEC Background!**

Multi-access Edge Computing (MEC), has picked up a variety of names:

- edge computing,
- edge cloud,
- fog computing,
- mobile edge computing,
- Etc.

So, what is MEC? It is about hosting one or more applications on compute, network and storage resources that are placed closer to the subscribers (residential or enterprise). Per ETSI, MEC is “an evolution of cloud computing [that] brings application hosting from centralized data centers down to the network edge, closer to consumers and the data generated by applications.” In other words, MEC is a cloud-based IT service environment at the edge of the network.

MEC was originally coined to benefit applications and subscribers with mobile access, however, it has since evolved to cover multiple types of access, including wireline access.

MEC essentially brings cloud capability (not only compute, but also networking and storage) to the network's edge and helps unlock superior experience to the “things” including eyeballs, sensors etc. It enables real-time, high-bandwidth, low-latency access to applications and subscribers, allowing operators to open their networks to a new ecosystem and value chain.

**MEC in the context of Cable access can be referred to as Cable Edge Compute.**

### **1.1.2 Cable Cloud vs. IT Cloud**

Clouds are about hosting application functions. However, depending on the applications, the Cloud could be designated as Cable/Telco Cloud or IT Cloud, independent of whether deployed on-premise or not.

**Cable/Telco Cloud** is about hosting Telco related infrastructure services such as Subscriber Edge Functions (e.g. CCAP, BNG, SGW, PGW), Access Network Functions (e.g. RPHY, OLT, eNB) as well as end-user services such as IMS voice, SBCs, video, media content, etc. It is worth pointing out that Cable Applications such as CCAP could be deployed as a single network function (e.g. single container) or multiple functions (e.g. multiple containers).

Cable/Telco Cloud applications tend to require high throughput (10Gbps+) and low latency/jitter centric infrastructure.

**IT Cloud** is intended to host IT related services such as Operations Support System (OSS) applications, Billing Support System (BSS) applications, end-user portal, media storage, etc. IT Application such as Portal could be deployed as a single function (e.g. single container) or multiple functions (e.g. multiple containers).

IT Cloud applications tend to require high compute and storage centric infrastructure.

Of course, the data centers implementing Cable/Telco Cloud or IT Cloud exhibit high degree of resiliency, faster convergence, etc. It is likely that the line between Cable/Telco Cloud and IT Cloud would continue to diminish and soon, there won't be any meaningful difference between Cable/Telco Cloud and IT Cloud, though Security posture may mandate them to stay separated.

Reference [1] and [3]

## 1.2 Why – Benefits?

Many Operators have 1000s of Hub Sites already deployed/operational across the country. These Hub sites already have Network Edge functions for certain services such as Internet Data etc. and are best suited for transformation into Application Centric Cloud that can host additional B2B and B2C services such as RAN, IoT, Gaming, AR/VR etc. and offer superior customer experience.

The Hub sites are usually already quite fiber rich and are employing innovative technologies such as distributed CCAP, Remote PHY, Full Duplex DOCSIS etc. These are quite complementary to the notion of Application Centric Cloud and Edge Computing that essentially brings cloud capability (compute, storage, network) to the network's edge and helps unlock superior experience to the "things" including eyeballs, sensors etc.

Cable operators can now be enabled to be the cloud providers, taking a page from the success of companies such as Amazon, Google, etc., and leveraging the networks & Hub Sites assets in a new way, and really strive to yield win-win:

1. Better Utilization – If the hub sites are transformed into application centric cloud sites, then they could allow hosting not only walled-garden services, but also 3<sup>rd</sup> party services on-demand and take advantage of geographic closeness to the subscribers.

In particular, Edge Computing is seen as key to massive IoT deployments and as crucial for analyzing large amounts of data coming from increasingly connected things.



2. Increased Security – If data is processed closer to the customer site instead of far away (public cloud, for ex), then the risk of data theft or illegal access is significantly reduced. One can localize mission-critical data processing to help meet security requirements.
3. Decreased Latency – If data is processed closer to the customer edge of the network in near-real time, then propagation delay could be significantly reduced.
4. Increased Control – If one is able to dictate which data stays local and what goes external for processing with utmost granularity to the device level, then it increases the overall optimality.
5. Better Operations – If key constructs of the network are virtualized on a common x86 hardware platform, then it could improve the operations in terms of seamless service creation environment and efficiency.

### 1.3 Where is the Edge? Centralized vs. Distributed ?

Since 1990s, the “Edge” has referred to the point where a "customer connects to the provider."

The provider being the organization providing a service such as broadband, telephony, video, mobility etc. to the customer belonging to enterprise, residential, retail etc.

However, since the last decade or so, the “Edge” has increasingly referred to the point “where the service is located”, largely because of the emergence of cloud service providers (CSP), which are more concerned about where the cloud services can easily run at scale.

Two things have changed – (1) more focus on workload centric services (less focus on network centric connectivity), and (2) more focus on proximity of the workloads wrt its users

So, while it is debatable where the Edge exactly is (the answer may vary quite a lot depending on who we ask – Content Providers/Aggregators e.g. Netflix, vs. Public Cloud Providers e.g. AWS, vs. Online Gaming Provider e.g. TakeTwo etc., vs. Subscribers e.g. eyeballs/ears/sensors etc.), it is important to describe the Edge in the context of Cable Operators paradigm. Arguably, many regard the Edge where the Subscriber Session Control function (e.g. CMTS) is instantiated – usually the Hub Site.

Few of public cloud providers have 20+ Edge locations in the USA and building more. They continue to expand their Edge presence, either by placing the applications and/or compute capacity in peering points and into the MSO/SP networks. However, most of their Edge locations are far from the access & aggregation/regional networks, and not as close to the subscribers (yet) as Cable Hub sites are. Nonetheless, one of the public cloud providers has partnered with an incumbent Telco SP to expand their number of Edge locations, as illustrated in the figure below.

Because each Azure CDN product has a distinct way of building its CDN infrastructures, Microsoft recommends against using POP locations to decide which Azure CDN product to use. Instead, consider its features and end-user performance. Test the performance with each Azure CDN product to choose the right product for your users.

Region	Microsoft	Verizon	Akamai
North America	Toronto, Canada	Guadalajara, Mexico	Canada
	Vancouver, Canada	Mexico City, Mexico	Mexico
	Querétaro, Mexico	Puebla, Mexico	USA
	San Juan, Puerto Rico	Querétaro, Mexico	
	Ashburn, VA, USA	Atlanta, GA, USA	
	Atlanta, GA, USA	Boston, MA, USA	
	Boston, MA, USA	Chicago, IL, USA	
	Cheyenne, WY, USA	Dallas, TX, USA	
	Chicago, IL, USA	Denver, CO, USA	
	Dallas, TX, USA	Detroit, MI, USA	
	Denver, CO, USA	Los Angeles, CA, USA	
	Honolulu, HI, USA	USA	
	Houston, TX, USA	Miami, FL, USA	
	Las Vegas, NV, USA	New York, NY, USA	
	Los Angeles, CA, USA	Philadelphia, PA, USA	
	Miami, FL, USA	San Jose, CA, USA	
	New York, NY, USA	Seattle, WA, USA	
	Newark, NJ, USA	Washington, DC, USA	
	Phoenix, AZ, USA		

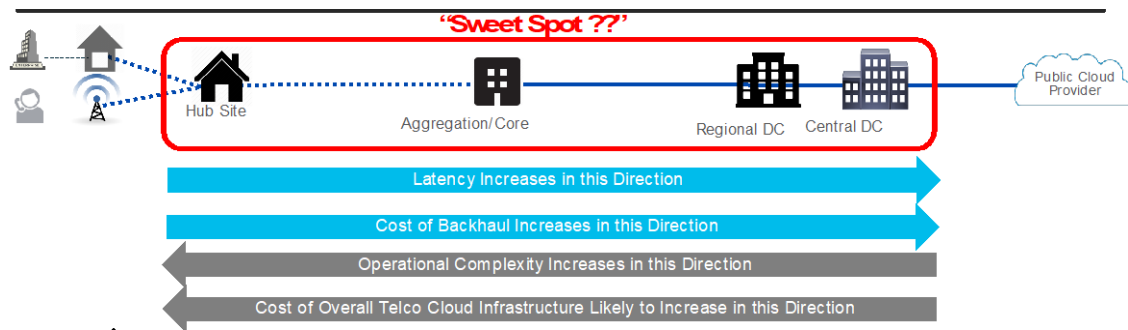
**Figure 4 Public Cloud Edge Locations - Example**

Source - <https://docs.microsoft.com/en-us/azure/cdn/cdn-pop-locations>

Subscriber closeness is a key advantage that Cable MSOs can utilize.


### **Centralized or Distributed or both**

It is important to highlight the pros & cons of centralized vs distributed in the context of what dictates the customer experience – latency, bandwidth etc. and what dictates the cost & complexity in MSO environment, as illustrated in the figure below -



**Figure 5 Centralize or Distribute the Edge Clouds - Pros & Cons**

It is somewhat clear that the applications that require SLAs comprising lower latency and higher bandwidth would demand Edge Clouds distributed in the network. For example, on-demand/online video consumption requires tons of downstream bandwidth, whereas physical security/monitoring (e.g. video surveillance) requires tons of upstream bandwidth. For example, the below figure illustrates a home network usage with growing upstream WAN traffic share.



**Figure 6 Upstream Bandwidth could be ~30% of downstream Bandwidth**

It is worth noting that almost all of ISP networks are built to optimize downstream traffic consumption, not upstream traffic consumption.

There are number of factors to consider in order to decide whether to distribute the Edge Clouds or not, as tabulated below –

**Table 1 Factors that can Influence Distributed Edge Clouds**

+ve Factors		-ve Factors
1	Reduction of Latency	Operational Complexity
2	Reduction of xHaul Bandwidth	Higher Infra Costs
3	Location Awareness	Location Availability
4	Regulatory / Compliance	Security Concerns
5	Localized Impact of Fault	Technology Maturity

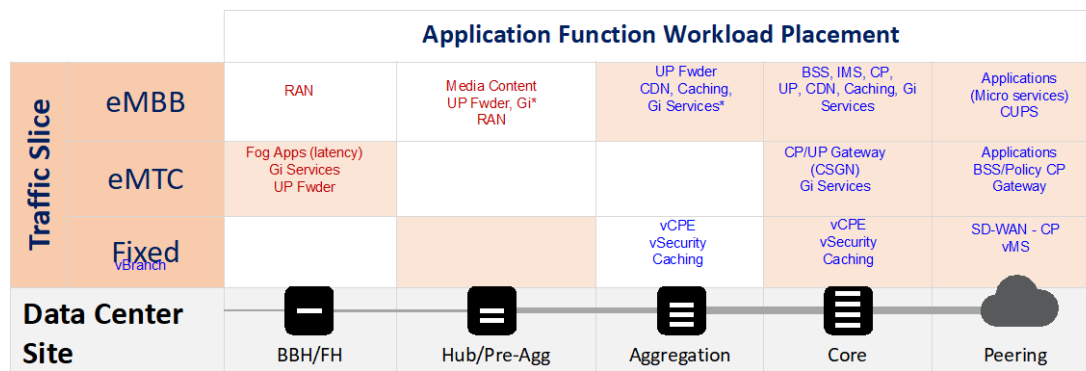
All in all, it is imperative to consider different applications – network centric, IT centric and customer centric and where all should the Edge Clouds be placed in the network in order those applications, considering the +ve/-ve factors from the above.

The below figure captures some of the applications (note that it doesn't cover all of applications) -

## Edge Cloud – Distributed vs Centralized

Bandwidth Intensive and Latency Sensitive Applications Demand Distributed

BBH = Baseband Hotel  
CP = Control Plane  
UP = User Plane  
CSGN = CIoT Serving Gateway Node



### Imperatives

- Offload mobile video traffic (78% by 2021) at edge
- Ultra low latency infra with large volume traffic
- Decomposition of RAN () virtualized
- Need to manage east-west traffic at edge

### Edge Cloud

- Distributed Micro datacenter for vRAN and User plane
- CUPS : deploy a user plane at edges and offload video traffic
- Edge CDN, Live TV, IOT, Online Gaming , AR/VR

**Figure 7 Example Application Functions for Edge Cloud Locations**

It is important to keep E2E network architecture and the relevance of Hub sites perspective. Please refer to the architecture section 3.

## 2 Cable Edge Compute – Application Functions

Edge Computing is one of the disruptive paradigms in the network architecture that enables a myriad of industry-specific use cases. By becoming edge cloud providers, Cable operators can leverage their nation-wide wired access networks (in addition to any wireless access) to shift their relationships with application developers as well as the customers who consume those services, and ideally, to become more like the agile, cloud and application-centric operators focused on innovation.

Of course, application functions would vary depending on the use-cases that are targeted. A few may be more relevant than the others. The applications fall within the below three categories

1. Infrastructure Use-Cases – RAN, BNG, CMTS, PGW/SGW etc.
2. Services B2C Use-Cases – CDN, LiveTV, IoT, Gaming, AR/VR, AI/ML etc.

### 3. Services B2B Use-cases – CDN Hosting, Online Gaming, Surveillance etc.

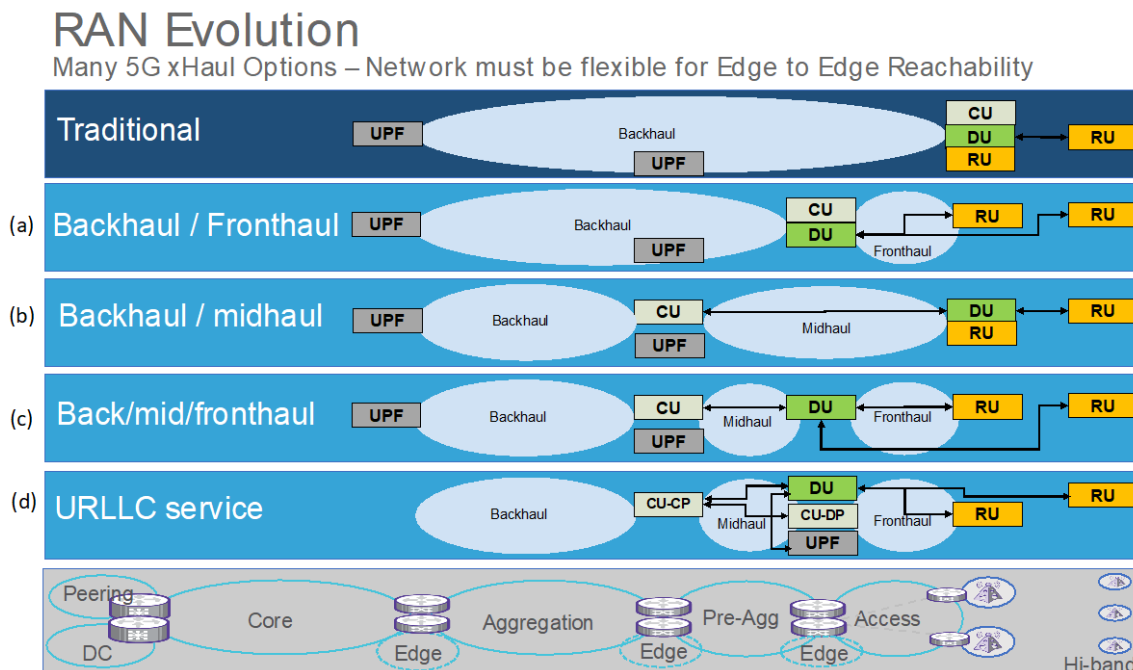
## 2.1 Infrastructure Use-Cases

### 2.1.1 (Virtual) RAN Functions

The way mobile networks have been built for decades are evolving and improving. Radio Access Network (RAN) with traditional cell sites, where the traditional monolithic eNodeB(s) have resided, are getting fast modernized with eNodeB getting decomposed into RU, DU and CU [4]. Operators have been developing RAN strategies around cost-effectiveness rooted in virtualization, cloud-native etc. In fact, they are the key tenets of the 5G RAN and made largely possible by “edge computing” with end-to-end automation for both infrastructure and services.

Modern RAN is centered around disaggregation and decomposition at multiple fronts e.g. two-layer split with cell site having only Remote Radio Units (RUs) and Antennas, virtualized Distributed Unit (vDU) functions processing lower layers of the radio stack and virtualized Central Unit (vCU) functions processing upper layers of the radio stack, will reside. RU – vDU connectivity being fronthaul and vDU – vCU connectivity being midhaul.

In essence, vRAN infrastructure service could comprise of 2 workload applications – Distribution Unit (DU) and Control Unit (CU). Each could represent a singular function or plural functions that could potentially be executed at the same site or different sites depending on the mobile operator’s design, as illustrated in the figure below:

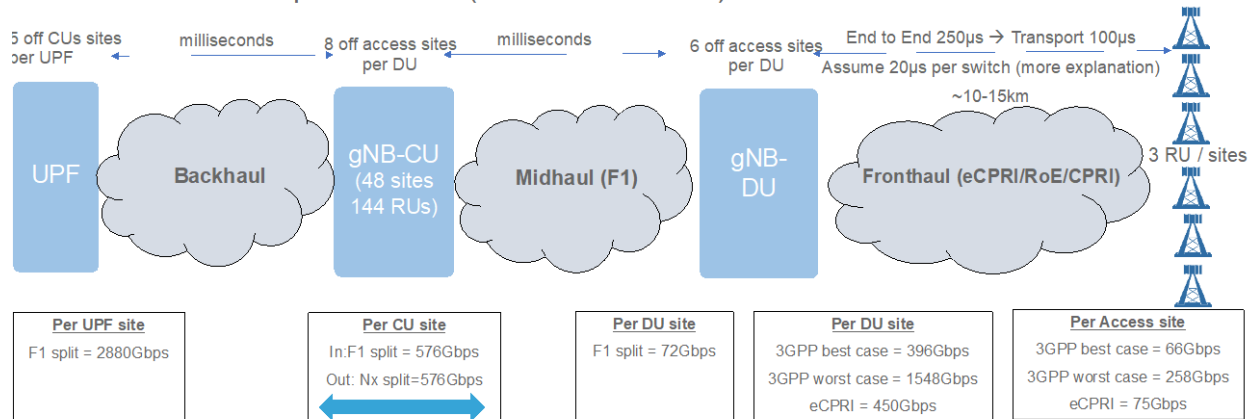


**Figure 8 Mobile RAN Evolution**

Note that the RAN fronthaul interface (RU-DU) deals with digitized RF signals, hence, it has a very high throughput (3Gbps+) and extremely low latency requirement (<250us) depending on the chosen split design, as illustrated in the figure below –

## RAN Evolution

5G bandwidth implications - (based on NGMN)



NOTE: No statistical gain assumed

DL	4 Gb/s*
UL	3 Gb/s*
Latency	1.5 – 10ms

(a*)	
DL	10 – 22 Gb/s
UL	16 – 22 Gb/s
Latency	250us

(b**)	
DL	38 – 86 Gb/s
UL	54 – 86 Gb/s
Latency	250us

(c)	
DL	10 – 22 Gb/s
UL	54 – 86 Gb/s
Latency	250us

Split numbers are 3GPP numbers based 5G 100MHz b/w (3GPP TR 28.801 v14); eCPRI figure approx 25 Gbps (UP: ~20Gbps CP: ~5Gbps) (e-CPRI specification v1)  
RRU: 100 MHz, 256 QAM, 8x8 MIMO IQ BW (7-16)bit, 32 Antenna Port) 3GPP TR 38.801 V14.0.0 (2017-03)

**Figure 9 5G Bandwidth in Access, Aggregation**

The point to take away here is that the usage of Cable Edge Hub sites could be suitable with fiber connectivity between the cell sites and edge sites.

Also worth noting that RAN focusing on the usage of higher frequency bands such as mmWave (24-86 GHz) (whereas the majority of current RAN deployments are around 2 GHz or below) want to benefit from increased radio efficiency/capacity etc, however, they also have to put up with corresponding limited coverage and strict line-of-sight consideration, which ultimately mean that RAN will likely have a lot more cell sites in a given area i.e. a lot more investment – in fact, according to a research report [<http://www.delloro.com/products-and-services/mobile-radio-access-network#5-year-forecast-report>], the 5G New Radio (NR) would propel the RAN market to around \$160B over the next 5 years.



**Figure 10 Radio Access Network showing 5G Cell Site with mmWave**

The point to take away here is that the usage of Cable Edge Hub sites would be suitable to provide fiber connectivity between these new cell sites with RUs and edge sites with DUs/CUs.

This is quite an opportunity for Cable Operators to position their distributed Hub Sites as the ideal places with CEC to host one or more virtual RAN functions for deeper and denser radio deployment.

## 2.1.2 CBRS

The US Government/FCC has provided 3.5 GHz (3550-3700 MHz) for Citizen Broadband Radio Service (CBRS). It has three tiers such that tier 1 “incumbents,” including ship-borne Navy radars, fixed satellite stations, and wireless providers, are protected from lower tier users at all times.

CBRS offers an economical path for Cable MSOs to enter the wireless industry via an MVNO strategy can now deploy LTE network and minimize network expenses by offloading the traffic to its owned CBRS LTE network (instead of sending it to the host MNO). This means that Cable MSOs can offer not only in-building & outdoor wireless coverage, but also capacity expansion for their own benefits or for their B2B partners’ benefits.

This is an upcoming opportunity for Cable Operators to leverage the Hub Sites as the suitable places with CEC to host CBRS functions.

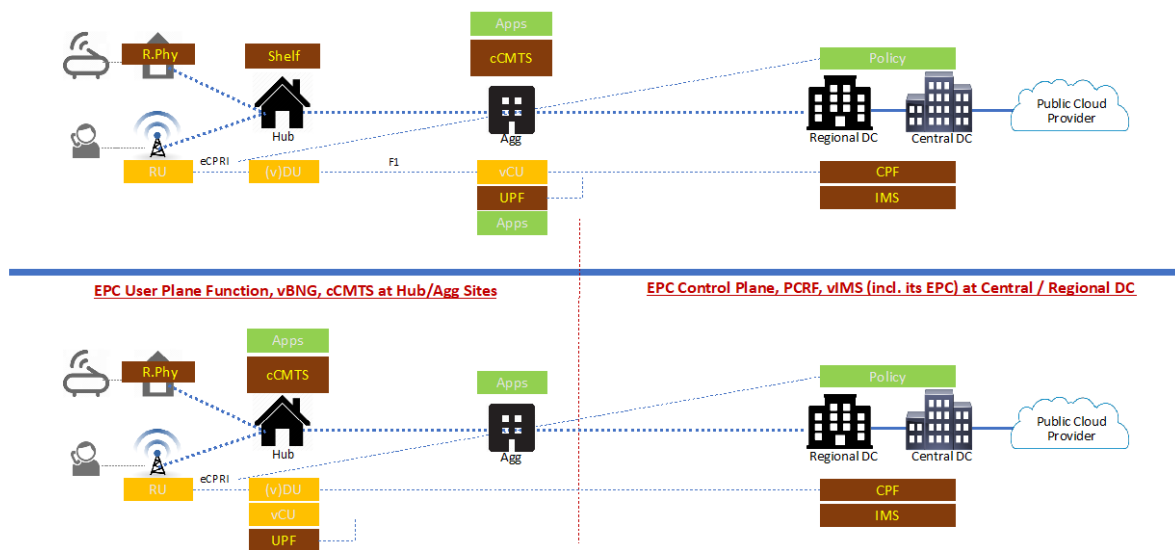
### 2.1.3 Subscriber Edge/User Plane Functions (e.g. CMTS, BNG, PGW)

Cable Operators have historically used CMTS / CCAP as the Subscriber Edge / Gateway function, which dictates the per-subscriber session and policy enforcement. Similarly, mobile Operators have used SGW/PGW (in 4G/LTE) as the Subscriber Edge / Gateway function. With the advent of Control Plane User Plane Separation (CUPS), Gateways could be decomposed, disaggregated and cloudified such that User Plane Function (UPF) could be deployed in a distributed manner, whereas CPF could be deployed in a centralized manner depending on the level of scale and aggregation needed.

Cable Operators can exploit the opportunity to embrace Fixed Mobile Convergence (FMC) by placing the Subscriber Edge Functions in virtualized manner at the transformed Hub sites. Furthermore, Control Plane User Plane Separation (CUPS) could enable a common converged UPF for cable access and mobile access distributed, while placing the CPF in a regional/centralized sites. This fits well with decomposed CCAP + RPHY approach many Operators are already pursuing.

The figure below illustrates the CMTS and SGW/PGW placement possibilities -

#### Virtualized Service Gateways



**Figure 11 Subscriber Edge Functions, CUPS**

Lastly, if certain traffic needs to be offloaded, then the CEC enabled Hub Sites could facilitate that, thanks to UPF.

## 2.2 B2C Service Use-Cases

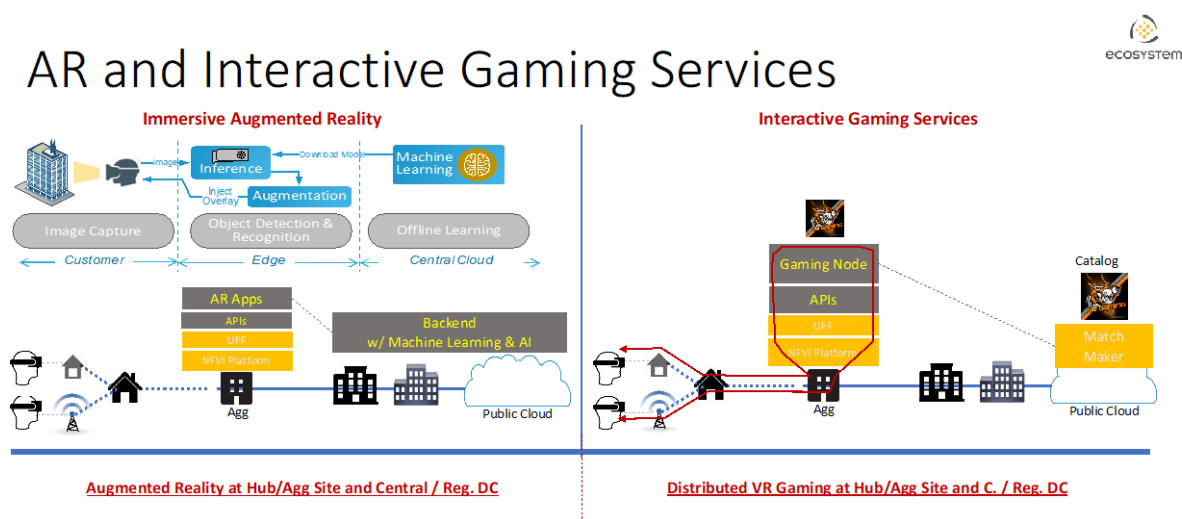
Operators have been offering plethora of services to the consumers. Some are highlighted here in the context of CEC.



## 2.2.1 Gaming

Online Video Gaming has attracted the attention of network engineers ever since the first multi-player online games appeared. Arguably, Online gaming has become one of the most profitable businesses on the Internet. The gamers do expect SLAs in both upstream and downstream direction notably in lowest latency and higher bandwidth with or without AR/VR. The game traffic is interactive and usually in Hub&Spoke pattern (as all client traffic is sent to the server(s) and sent back to the clients).

Google Stadia (recently announced) is making online gaming similar to that of content streaming that are hosted in the cloud allowing any type of endpoint devices. The below figure illustrates Gaming Node being hosted closer to the subscribers –



**Figure 12 Gaming Service**

Note that Gaming Services may require GPUs in addition to CPUs on the x86 server nodes.

Interestingly, even if the games are played offline, the games get downloaded online and mere game download consumes a significant network (downstream) bandwidth. For example, Call of Duty: Black Ops 3 is 101GB, Grand Theft Auto V is 65GB. In contrast, an hour of 4K video on Netflix is about 7GB per hour, making a Call of Duty download equivalent to watching over 14 hours of 4K video!

It is worth highlighting that game downloads can affect ISP network utilization far worse than windows/iOS/macOS downloads, or even 4k movie download. Imagine the network contention that may arise if 20 users in a domain are downloading the latest game edition, while 20 users in the same domain are playing the game. In other words, caching such games (similar to that of on-demand videos) closer to the end customers could help to minimize the severe downstream bandwidth stress on rest of the network.

## 2.2.2 LiveTV

This one is an obvious one in which the Operators can move the content caches (live and on-demand) or translators closer to the subscribers in a distributed fashion in order to reduce the network load exponentially -

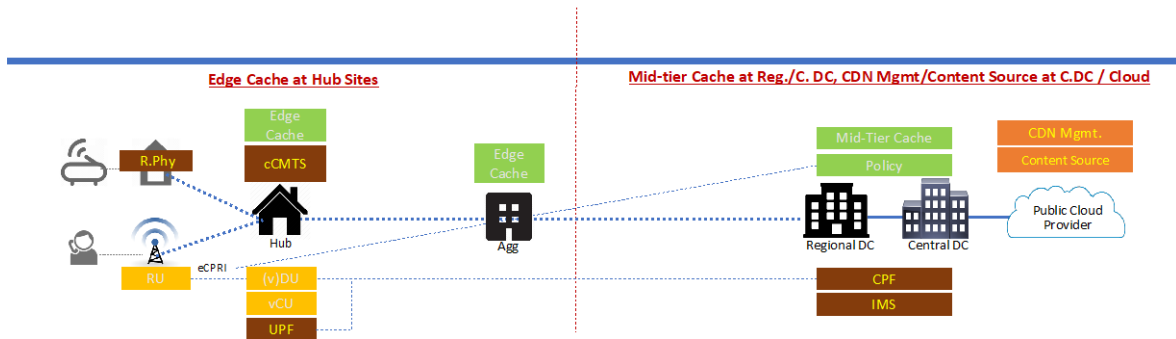


Figure 13 Managed Video/CDN

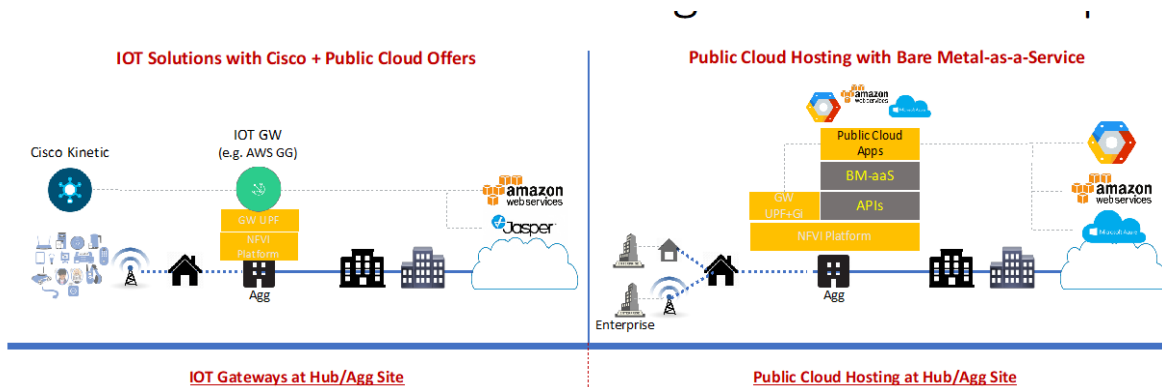
## 2.3 B2B Use-Cases

### 2.3.1 IoT & Public Cloud Hosting

One of the challenges of IoT is providing scalable connectivity to support a huge number of devices.

It is important to characterize IoT devices as – Heavy and Light. IoT Heavy devices are generally bandwidth intensive, require more power and sophisticated compute capabilities. Few examples are connected vehicles, autonomous control of large machinery, etc. IoT Light devices are highly constrained devices with minimal compute, memory, energy supply. Few examples are environmental sensors, under road monitors for smart parking, water meters etc. The IoT Light Devices require <200Kbps of data rate.

IoT Heavy devices such as the ones in factory automation etc. require lower latency treatment (say, few msec) and to do so, they would need to be hosted as close to the endpoints as possible for local processing.




**Figure 14 IOT and Public Cloud Hosting**

### 2.3.2 3<sup>rd</sup> Party CDN

Operators can offer the transformed Hub sites to the 3<sup>rd</sup> party CDN operators and reduce the network bandwidth consumption while improving their customer experience.

### 2.3.3 Video Surveillance

The Video Surveillance / Cloud Monitoring services continue to gain traction in the context of Home/Business Security. The HD/UHD cameras could stress the network in upstream direction quite a bit. One HD camera could consume 1Mbps+, while streaming. The figure below illustrates the monthly consumption, mostly in the upstream direction -



**Figure 15 Video Monitoring adds significant Upstream bandwidth Consumption**

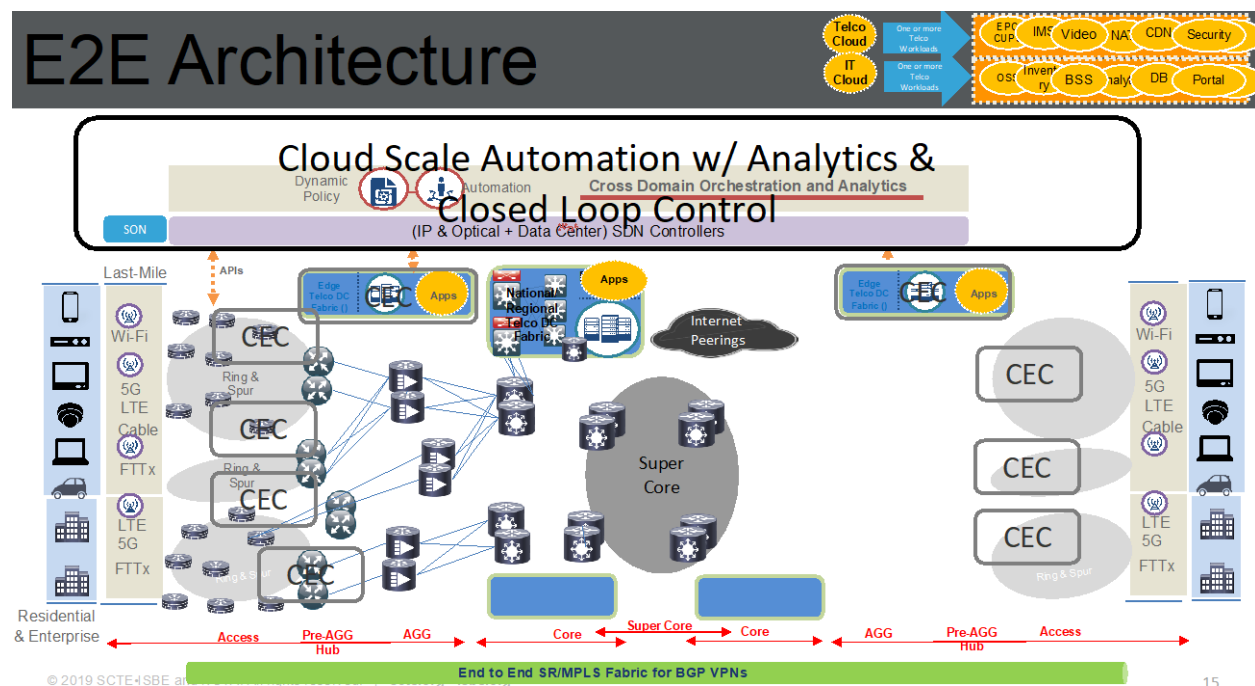
It is possible that such a service could be hosted in the Hub Sites, thereby significantly saving the network bandwidth in other domains.

### 2.3.4 Security

Security related workloads e.g. firewall, DPI, detection/prevention system etc. go hand in hand with the adjoining applications (of whichever category) in order to protect them, in addition to (cloud based) security related services that could be offered to the customers.

## 3 Cable Edge Compute – Architecture

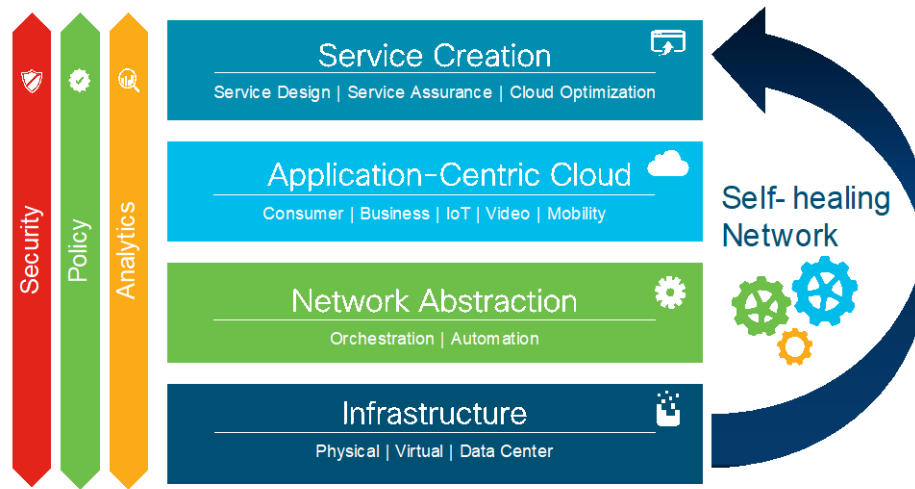
It is important to highlight the possible Cable Edge Compute (CEC) locations in the perspective of E2E network architecture, as illustrated below, even though the focus is more on the Hub sites –



**Figure 16 E2E Architecture Blueprint Showing Cable Edge Compute (CEC)**

Given the wide variety of possible use-cases/applications (few are discussed in section 2) that could be hosted at the designated “Edge” locations, they require the presence of the same consistent cloud Platform, on top of which the application functions (mostly virtualized/containerized) can be deployed in a seamless manner (e.g. IaaS, PaaS, FaaS). This is possible with sufficient abstraction. A blueprint of such a platform is illustrated below -

## Network Platform



**Figure 17 Network Platform - Abstraction is KEY**

There are 4 building blocks that constitute “Edge” architecture to facilitate the Application Centric Cloud construct, each with certain unique properties –

1. Edge Infrastructure – Hardware
2. Edge Infrastructure – Software Platform (NFVI)
3. Network Fabric (SR)
4. Automation (SDN), Orchestration and Assurance

### 3.1 Infrastructure – Common Hardware

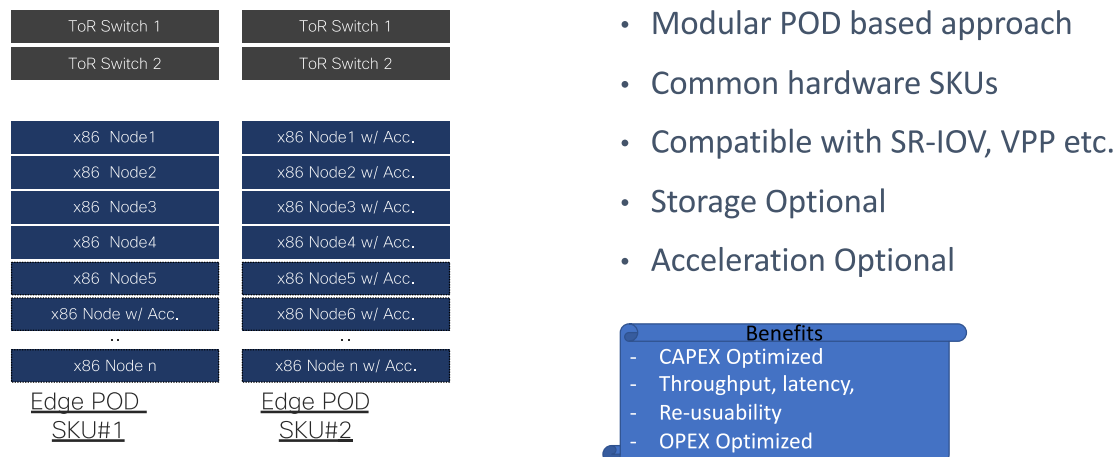
To morph MSO Hub Sites into the Application Centric Cloud that can accommodate wide variety of possible applications (mostly virtualized/containerized) for whatever time-period, the infrastructure should be flexible enough to not only provide higher throughput, lower latency, multi-tenancy, VM/Container workloads, security etc, but also enable scale-out in a policy-driven manner. Consistency is KEY for cost-effectiveness.

It is likely that many Hub sites may have smaller footprint (as compared to the traditional Data Centers) in terms of space, power, cooling, rack depth etc., whereas a few Hub Sites may have slightly larger footprint. Hence, the hardware infrastructure should be built while striking the balance among footprint optimization, performance and cost.

In other words, whether a Hub site has space for only 2 racks or 6 racks or more, the hardware infrastructure should be expandable *without requiring any/much architectural changes*.

This flexible expandability requires a **modular POD based approach** that can provide consistency using a set of common hardware configurations (SKUs) for x86 server, storage and network. We refer to them as Edge POD, as illustrated in the figure below –

## Building Block #1 Infrastructure – Hardware

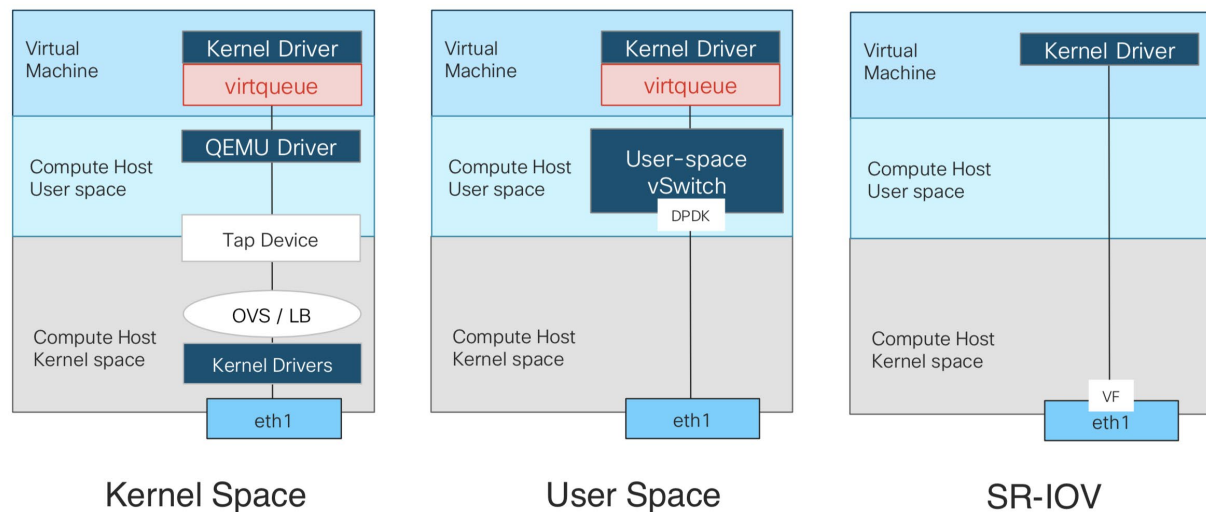


**Figure 18 Hardware Infrastructure - Edge PoD SKUs**

It is worth noting that recent innovations have helped to decouple storage and move the storage nodes to the locations such as conventional Data Centers that don't have similar constraints as Hub Sites do. This allows for simplicity and cost efficiencies, if/when application workloads don't require storage. This is further covered in the section 3.2.

The hardware infrastructure could also optionally include Accelerated Units (e.g. GPU, FPGA) on compute node(s), if required by any application workloads. For example, vRAN's vDU application workload may require eCPRI radio signal processing to be done on those x86 compute nodes housing specific Accelerated Units (e.g. FPGA), or Gaming's Application workload may require certain processing to be done on only those x86 compute nodes housing specific Accelerated Units (e.g. GPUs).

Additionally, It is very important for the hardware to be compatible with more than one virtual forwarding innovations whether user-space forwarder such as VPP/fd.io etc., or kernel-space forwarder such as OVS etc. or something that bypasses host kernel and user-space altogether such as SR-IOV etc. (as illustrated below) to ensure optimized forwarding behavior as expected by the application functions.



**Figure 19 Virtual Forwarding Options**

The Edge POD **consistency** ensures not only repurpose-ability, but also operational simplicity.

An example is shown for 4 NFV Infrastructure hardware SKUs that could be consistently deployed in the Hub sites depending on the requirements –

**Table 2 Infrastructure Hardware - POD SKUs**

	Hardware SKU1	Hardware SKU2	Hardware SKU3	Hardware SKU4
<b>CPU / GPU</b>	2 sockets [2.4GHz, 48 Cores, ...]	2 sockets [2.4GHz, 48 Cores, ...]	2 sockets [2.4GHz, 48 Cores, ...]	4 sockets [2.4GHz, 48 Cores, ...]
<b>Memory/RAM</b>	256GB	256GB	256GB	1TB
<b>Memory/Disk</b>	2x1TB SSD	2x1TB SSD	2x1TB SSD	2x1TB SSD
<b>HW RAID</b>	Yes	Yes	Yes	Yes
<b>NIC</b>	4 NICs: 2x10Gbps	4 NICs: 2x10Gbps	4 NICs: 2x10Gbps	4 NICs: 2x40Gbps
<b>Acceleration</b>	Yes/FPGA	- NA-	- NA -	- NA -

The smaller set of SKUs help to simplify operational and budgeting.

While the current Edge POD design assumes dedicated x86 server nodes, in the future, it may be possible to leverage the compute capacity of routers to be able to host containers or functions.

## 3.2 Infrastructure – Software (NFVI)

Software Infrastructure comprises of the Operating System (e.g. Linux/KVM) that facilitate reliable and deterministic “Virtualization” environment on top of the Hardware Infrastructure (i.e. Edge POD) that would be deployed during the Hub Site transformation, as well as Cloud Orchestration Platform that facilitate virtualized Infrastructure management.

Given the varying set of Hub sites constraints (mentioned in section 3.1), the Infrastructure Software stack must allow for maximizing the usage of Edge POD hardware resources for the designated Cable/Telco/IT applications workloads. This is KEY for superior cost efficiencies.

This means a typical NFVI software stack inc. Virtual Infra Manager (VIM) must take as least overhead as possible to keep most of the resources available for application workloads related to the use-cases.

Put it other way, VIM such as Openstack that may require at least 7 server hardware nodes (3 nodes for control, 3 nodes for storage and at least 1 node for management; see [6] and [7] for more details) either in the same rack or in adjacent racks in the site) for non-service purposes might NOT be acceptable. Consider a space/power constrained Hub site that can accommodate only 10 server nodes, then if the VIM software takes a large percentage (e.g., 70%) of the server nodes dedicated for cloud management, and not for hosting application workloads, then it would yield “Poor cost efficiencies”.

The Infrastructure Software Stack should have the following attributes in order to minimize the Edge POD hardware footprint suitable for CEC cloud in Hub Sites –

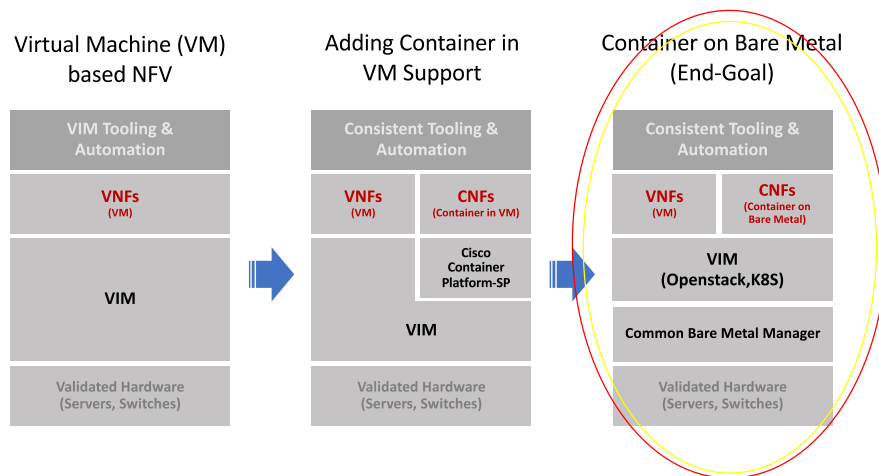
1. Common Cloud Orchestration for VNF or CNF
2. Cloud Orchestration Control Nodes combined with Compute Nodes
3. Deterministic NFVI performance
4. Remote Storage with Optimization and Security
5. Remote Management & Monitoring

Each of these are now further detailed below.

### 3.2.1 Common Cloud Orchestration Platform

The infrastructure software stack should allow for orchestrating both VM based Application Functions aka Virtual Network Functions (VNF), Containers based Application Functions aka CNF as well as Physical Functions (e.g. Bare Metals) on the chosen Infrastructure Hardware i.e. x86 by appropriately leveraging Virtualized Infrastructure Managers (VIM) independent of where the functions are instantiated (e.g. Serverless or not). This allows for future compatibility.





**Figure 20 Converged Cloud Platform**

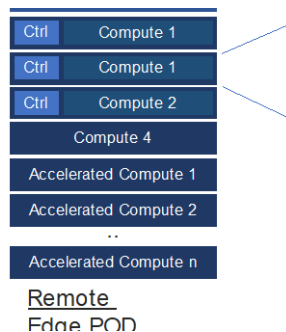
Containers may be hosted on Bare Metal or inside a VM. Few Examples of VIM are Openstack, K8S, vSphere etc.

VNF Manager (VNFM) and NFV Orchestrator (NFVO) would likely live in the centralized / regionalized data centers, though a lite version of VNFM could be hosted on each Edge PoD for VNF/CNF monitoring.

### 3.2.2 Orchestration Control Nodes combined with Compute Nodes

Instead of dedicating 3 Nodes for VIM/Cloud Controller purposes (e.g. Openstack Controller), they could be configured to run both VIM/OpenStack controller and compute functions. More specifically, the compute nodes run the host operating system (e.g. Linux/KVM) along with VIM, as well as the application workloads (IT/Cable/Telco Cloud).

The VIM/OpenStack controllers on the chosen 3 nodes continue to be in an active-active-active cluster configuration (with load sharing) for redundancy. Any additional nodes may be used as a pure compute node to scale the cloud on an as-needed basis. This is shown in the picture below –



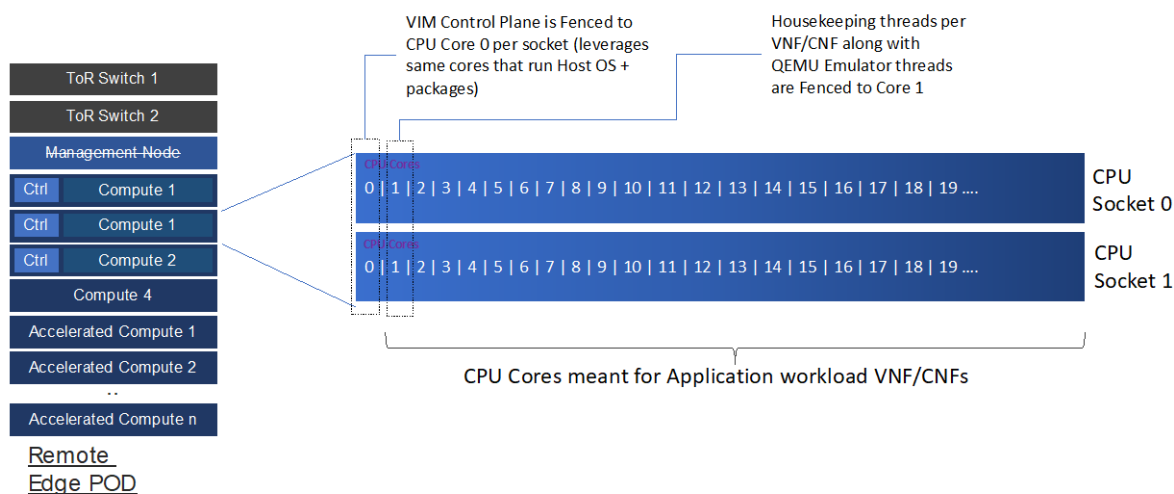
**Figure 21 Collapsed VIM Controller and Compute Nodes**

Additional compute nodes could potentially be reconfigured to have VIM controller function, if an existing compute+controller node failed.

### 3.2.3 Deterministic NFVI Performance

The collapsed control and compute functions are known to hamper the overall NFV performance. For example, latency may increase time to time.

To dramatically improve the overall NFV system performance in a deterministic manner, the software stack must be designed such that the VIM and VNF/CNF related task(s) run only on the specified CPU core(s), and the specified CPU core(s) are allowed to only run the chosen task(s), whether shared or not. This could be done by appropriately fencing the CPU cores from two distinct angles, as illustrated in the figure below -



**Figure 22 Deterministic NFVI Performance Logic**

At least one of the CPU cores per socket can be used to receive interrupts and cannot be used for guest workloads. The logic employed is that the application VNFs can share the CPU core(s) on a particular socket with other VNFs' to run non-real-time tasks. This logic is followed whether or not hyper-threading is enabled.

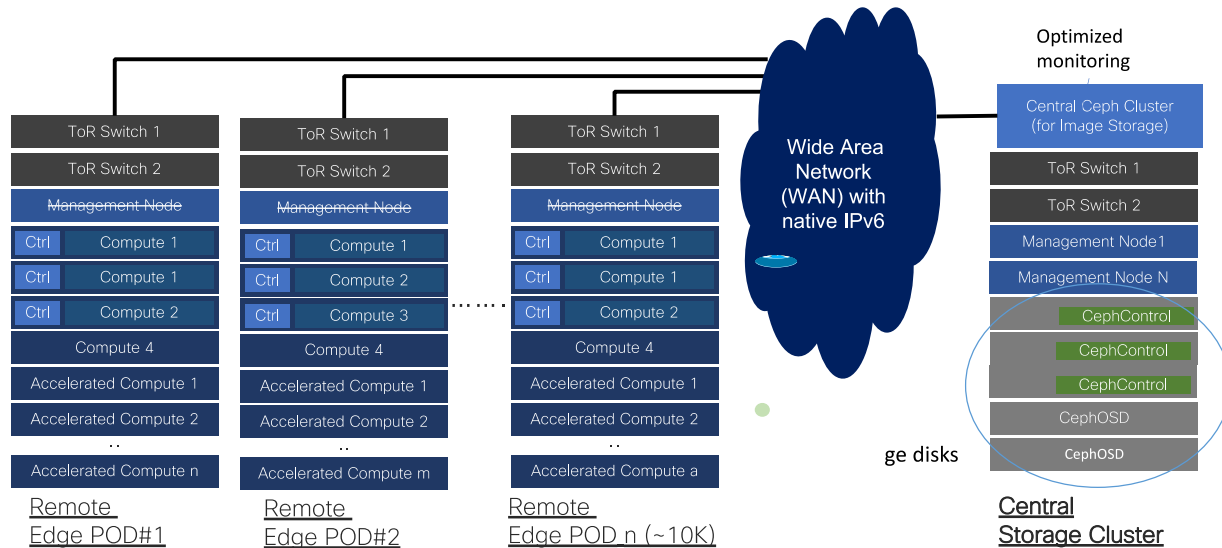
In addition to specific CPU reservations, the software stack should allow for leveraging the virtual forwarding innovations such as SR-IOV, VPP etc. for optimized forwarding behavior as expected by the application functions. For ex, VNF1 may require SR-IOV, whereas CNF2 may require VPP based forwarding on the same x86 host.

### 3.2.4 Remote Storage with Optimization & Security

To avoid having to dedicate any nodes for Storage purposes (e.g. persistent storage, object storage and similar, unaffected by latency variations), Storage can be moved to remote location (e.g. centralized or regional Data Centers that don't have similar constraints) and accessed by the CEC Hub sites over the network.

For example, the Ceph service for glance image services (including bulk transfer such as software image download) is no longer available locally inside the CEC Hub site. This is due to the assumption that Image-based (and object-based) storage is infrequently fetched and can be cached at the expense of latency (and would be costly if done at the edge). Thankfully, Latency isn't an issue, given the time it takes to download a VNF image is a lot more than WAN latency. Avoid using block-based storage because it may not be feasible to centralize it due to slow responsiveness (and would be costly at the edge).

This is illustrated in the figure below -

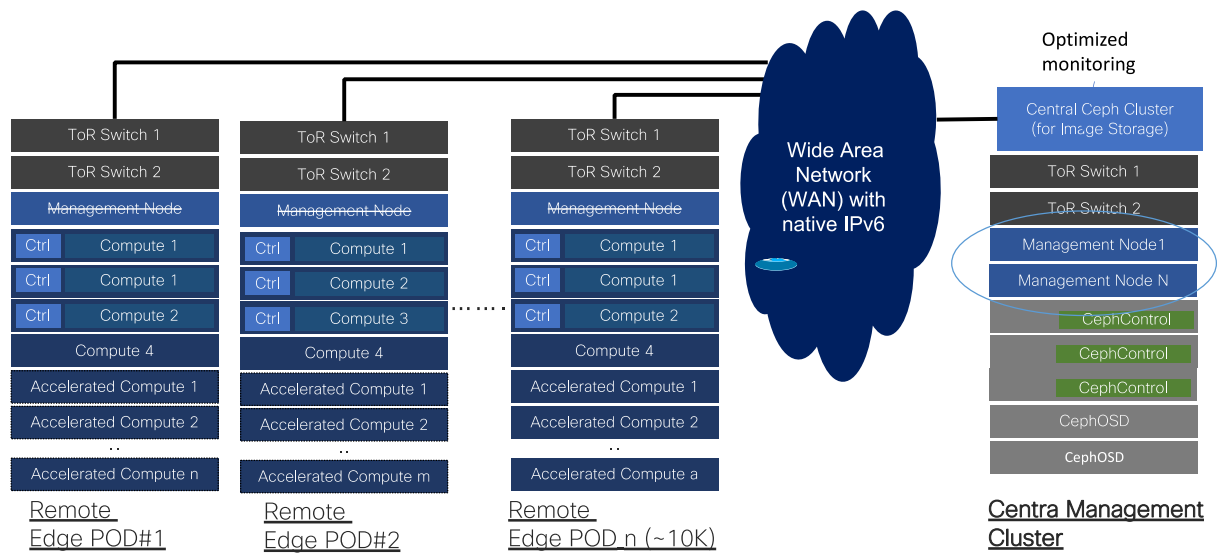


**Figure 23 No more Local Storage for VIM**

It is important to limit the potential bottleneck if multiple CEC edge PODs simultaneously interact with a single central Ceph cluster. Also, to ensure tight secured access between Edge PODs and Storage cluster, proper authentication and encryption (if necessary) of their communications (REST o TLS) is enforced.

### 3.2.5 Remote Management & Monitoring

Management Nodes perform number of important tasks – VIM deployment, monitoring, operations (e.g. node addition, replacement), version control, software version changes etc. To avoid having to dedicate any nodes for Management purposes in CEC Hub sites, management functions can be moved to remote location (e.g. centralized or regional Data Centers that don't have similar constraints) and accessed by the CEC Hub sites over the network. This could be quite challenging, but can be achieved, as illustrated in the figure below -

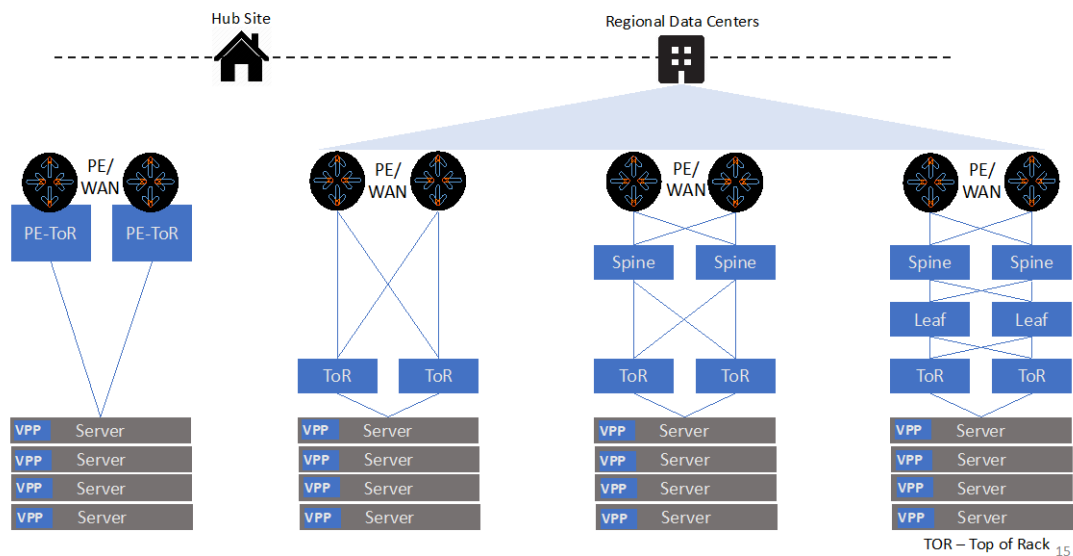


**Figure 24 No more local Management Nodes**

For example, the installation as well as monitoring of CEC nodes can be completely automated via a single intent file.

### 3.3 Network Transport Fabric

In order for the traffic to enter and exit the Hub Site, the Compute complex must appropriately connect to the network transport. There are 4 different options (4-tier, 3-tier, 2-tier, 1-tier) to design Network Fabric, as illustrated in the figure below –



**Figure 25 Network Transport Fabric**

For the Hub Sites, 1-tier or 2-tier design is more appropriate to keep the footprint minimal while ensuring SLAs. Network Transport Fabric should have the following attributes -

- 10/40/100Gbps+ Ethernet Transport
- Any-to-Any Reachability
- E2E IPv6 with Segment Routing
- BGP based VPN
- Programmable WAN and DC Fabric
- IP+Optical WAN

### 3.4 Automation, Orchestration and Assurance

The architectural framework must have two foundational elements to ensure the Application Centric Cloud transformation of Hub Sites becomes viable.

Firstly, the Hardware Infrastructure (e.g. x86 nodes, FPGAs etc.), Software Infrastructure (e.g. VIM) and Network transport (e.g. TORs) must be managed with stringent automation and orchestration for its entire life-cycle. Zero-touch instantiation of Infra management is also required. Consider FPGA installation or firmware changes should be handled in an automated manner.

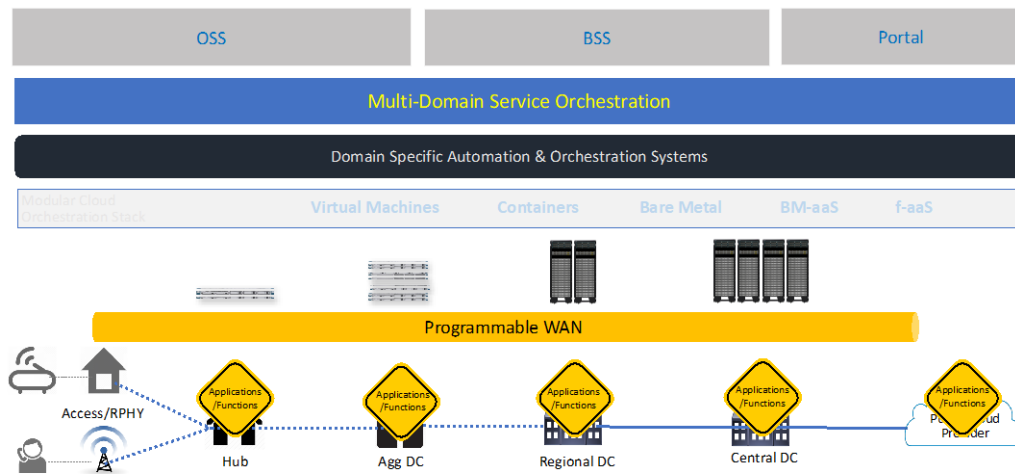
Secondly, each service instantiated at the Edge may comprise of one or more number of application functions that would need to be instantiated, provisioned and configured in the right sequence with the right service level agreements (SLAs) in a zero-touch manner. Assurance would be foundational to monitor the SLAs with closed-loop.

In order to appropriately enforce the policies (such as access control, inspection, prioritization, optimization etc.) on the traffic passing through any one or more of the (service) functions deployed at the edge, it is important to maintain the sequence in which the traffic must pass through those functions. This is referred to as service-chaining.

For example, Service chaining can help to regulate the traffic flow for service A to pass through application functions 1, 2 and 4 in a sequence, whereas traffic flow for service B to pass through functions 2 and 5.

Application Functions can be inserted or deleted into the flow by modifying the chain, as controlled by the orchestrator, as explained in [5].

The architectural framework is based on the concept of hierarchical management and orchestration, and consists of domain-level orchestration systems, where the domain would correspond to Network Transport, Edge POD, Data Centers etc. The domain-level orchestration systems would cater to domain specific management. For example, Edge POD in CEC Hub site would be managed via the MANO stack (VIM, VNFM, NFVO etc.). These orchestration systems are glued together in a modular architecture framework with an end-to-end service orchestration that would interact northbound with OSS and BSS systems, which offer a comprehensive set of service instantiation, service lifecycle management and operational workflows.



**Figure 26 Automation, Orchestration and Assurance**

The architectural framework should have the following attributes –

1. Infrastructure Software - lifecycle management with CI/CD pipeline
2. Automated Service Validation
3. Self-Service Portal
4. Closed Loop Assurance
5. Single Pane of Glass

## 4 Summary

Cable Operators can transform their Hub sites into Next Gen Application-Centric Cloud sites (similar to how Telecom Operators could transform their Central Offices (COs)) and take advantage of next swath of revenue generating services that can be hosted closer to the subscribers in order to ensure superior customer experience.

They should ensure that the Hub transformation approach adheres to key architectural building blocks – IH (Infrastructure Hardware), IS (Infrastructure Software Stack), NF (Network Fabric) and AOA (Assurance, Orchestration, Automation) along with key attributes.

### 4.1 Opportunities

Many Operators have 1000s of Hub Sites deployed across the country and they are not currently fully utilized to host applications pertaining to B2B or B2C services and offer superior customer experience.

If these Hub sites can be transformed into hosting qualified applications such as IoT, bandwidth hungry applications e.g. CDN/cache edge, lower latency applications e.g. gaming, then these applications (and resulting services) would get a lot closer to the customer.

### 4.2 Challenges

There are number of challenges in sufficiently utilizing Edge Computing in MSO environment. Few are captured below –

1. IT Cloud and Cable Cloud separation – Diverse application types (in IT and Cable/Telco space) would demand different SLAs (some may be more stringent than the others), hence, a common cloud could be deemed difficult. However, maintaining separate Clouds for Cable/Telco and IT applications / workloads could result in many clouds distributed on the network and less efficient usage of infrastructure.

In other words, one hand, the desire to push for high performant, efficient use of Edge resources have to be balanced with a cookie cutter way of deploying a common cloud.

2. People/Skills/Culture – People get less excited about making drastic changes or get entrenched with identifying tons of issues that will make things not work. Sometimes, it is because of lack of appropriate skillsets, sometimes, it is because of the organization culture. This can single-handedly make or break the Hub transformation. Also, Network Centric teams tend to downplay the IT Centric teams and vice versa.
3. Stringent Assurance Needs – While one of the most important, Assurance topic doesn't surface until later in the conversation. That could hamper the overall efficacy. Programmatic connectivity to ensure the “Service Assurance” would be difficult if Cloud and Network together are managed with intense Automation.
4. Hub Site Physical constraints – Power availability (AC vs DC) and Space availability (e.g. raised floor, rack size etc.) to install x86 platforms would become important. Hub sites conventionally

may not be built according to the cloud needs. Also, the failure rate with x86/storage nodes may be higher than the routing/switching equipments, so frequent swapping may require people (if not robots) with suitable skills.

5. Lawful Interception (LI) – It may not be feasible to comply with LI if applications are designed and deployed where there is no way to place taps.
6. Multi-vendor and 3rd Party/Partner workloads – B2C vs B2B business approach may drive application workloads from different vendors / partners with their prerequisites that may break the feasibility. For ex, vendor-proprietary virtualization manager co-located with the workload instance, or VNF/CNF not fully compliant with APIs etc. If they are not able to take advantage of the APIs for Easy Consumption of the Platform, then it could jeopardize the overall efficacy, and may even derail the Edge Computing insertion.
7. Challenging Software Lifecycle Management – In virtualization paradigm, the agility is key. So, the number of software changes is a lot more than typical operations expect. If lacking 100% automation, then it would require human intervention.
8. Cloud-Native Applications – Not all applications are fully containerized yet. This means they would not be able to take advantage of Edge Computing resources that tend to assume disaggregated applications for utmost efficiency.
9. Too many Standards and Industry Groups – This makes Cable MSOs choosing particular group(s) a bit difficult and also hampers the overall progress. For ex, 3GPP, MEC WG, CORD, TIP, OpenFOG, OPNFV, ONAP, etc.

## 4.3 Recommendations

This paper highlights several recommendations that Cable Operators could consider as they look to transform their Hub Sites with a focus on Application-Centric -

1. Develop a roadmap for what (IT and/or Telco) applications could run on the common CEC platform initially vs later on. Carefully Analyze and Select the Application type that would yield better bang for the buck (e.g. bandwidth savings, latency reduction etc.) on a common CEC platform.
2. Train and retrain the workforce and push them to drive the changes (not just accept the changes). Take risk, fail fast approach should be advocated.
3. Mandate “Assurance-first” approach while evaluating the CEC offerings, even if the Assurance systems are different for Applications, CEC infra and Network. They must be programmatically linked to ensure AI/ML powered Operations.
4. Choose Hub sites that have limited or no physical constraints early on. Ensure CEC PODs with highly optimized hardware and software stack in fewer variations (e.g. 8 server nodes, 12 server nodes, 24 server nodes) that can be well standardized for different Hub sites.
5. Applications that run on CEC must be first cleared for the LI compliance.
6. Specify a common set of requirements so as to accommodate application workloads from different vendors / partners
7. Develop 100% automated software lifecycle management with CI/CD pipeline and ensure that it works in as many failure scenarios as possible.



8. Push the applications owners to provide decomposed and/or disaggregated workloads that can run on Kubernetes etc.
9. Choose fewer standards groups that are not limited to mobility only. For ex, Edge Computing Consortium.

# Abbreviations

CEC	Cable Edge Computing
VNF	Virtual Network Functions
CNF	Cloud Native Functions
HFC	hybrid fiber-coax
MEC	Multi-Access Edge Computing
B2B or B2C	Business to Business, or Business to Consumer
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers
CUPS	Control Plane and User Plane Function
CPF	Control Plane Function
UPF	User Plane Function
BNG	Broadband Network Gateway
CCAP	Converged Cable Access Platform
cCMTS	Containerized Cable Modem Termination System
VIM	Virtual Infrastructure Manager
MANO	Management and Network Orchestration

# Bibliography & References

- 1 Managing the 5G Telco Cloud – <https://blogs.cisco.com/sp/managing-5g-at-the-data-center-level-first-steps>
- 2 Reimagining the End-to-End Mobile Network in the 5G Era - <https://www.cisco.com/c/en/us/solutions/service-provider/mobile-internet/reimagining-mobile-network.html>
- 3 Real-World 4G/5G Use Cases [https://www.cisco.com/c/dam/m/en\\_us/network-intelligence/service-provider/digital-transformation/knowledge-network-webinars/pdfs/0522-mobility-ckn.pdf](https://www.cisco.com/c/dam/m/en_us/network-intelligence/service-provider/digital-transformation/knowledge-network-webinars/pdfs/0522-mobility-ckn.pdf)
- 4 Open vRAN Ecosystem [https://www.cisco.com/c/dam/m/en\\_us/network-intelligence/service-provider/digital-transformation/knowledge-network-webinars/pdfs/0920-mobility-ckn.pdf](https://www.cisco.com/c/dam/m/en_us/network-intelligence/service-provider/digital-transformation/knowledge-network-webinars/pdfs/0920-mobility-ckn.pdf)
- 5 Using SDN Controller in Telco DC <https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-740717.pdf>
- 6 Openstack Example Architecture <https://docs.openstack.org/install-guide/overview.html#example-architecture>
- 7 <https://www.mirantis.com/blog/making-openstack-production-ready-kubernetes-openstack-salt-part-1/>
- 8 Towards MEC Edge Compute [https://www.cisco.com/c/dam/m/en\\_us/network-intelligence/service-provider/digital-transformation/knowledge-network-webinars/pdfs/0417-DC-CKN-PDF.pdf](https://www.cisco.com/c/dam/m/en_us/network-intelligence/service-provider/digital-transformation/knowledge-network-webinars/pdfs/0417-DC-CKN-PDF.pdf)
- 9 Latency matters more than Bandwidth <https://www.igvita.com/2012/07/19/latency-the-new-web-performance-bottleneck/>

# **Winning the Gaming War: Play for Cable Operator**

## **Assuring video game experience across multiple domains**

A Technical Paper prepared for SCTE•ISBE by

**Alon Bernstein**

Distinguished Engineer  
Cisco Systems  
alonb@cisco.com

**Rajiv Asati**

Distinguished Engineer  
Cisco Systems  
rajiva@cisco.com

**Sangeeta Ramakrishnan**

Distinguished Engineer  
Cisco Systems  
rsangeet@cisco.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Overview .....	5
1. Classification of Games.....	6
2. Game Downloads .....	7
3. The Roles of the Game Server.....	8
4. How Gamers Perceive “Lag” vs. How Network Engineering View It .....	9
5. Lag Compensation.....	10
6. Human Response Time.....	11
7. How is the Network Performance Profiled?.....	12
8. Packet Transport.....	13
9. Gaming Traffic Characteristics – Packet Sizes .....	14
Domains & Recommendations.....	15
10. Game Rendering.....	15
11. Home Network Domain .....	15
11.1. WiFi – WiFi5 vs WiFi6.....	16
12. SP Network Domain – First Mile.....	17
13. Achieving Low Latency on a DOCSIS Network.....	18
13.1. Sources of Delay in a DOCSIS Network.....	18
13.2. CableLabs LLD.....	19
13.3. DOCSIS Delay vs. PON delay .....	20
14. SP Network Domain – Second Mile.....	21
15. Internet Domain .....	21
16. Data Center Domain .....	22
Gaming as a Managed Service.....	23
17. Classification.....	23
18. Gaming as a Marketing Play .....	25
19. Assuring gaming performance across domains .....	25
B2B vs B2C monetization .....	25
Future topics .....	26
20. Cloud Game streaming .....	26
Conclusions.....	28
Abbreviations.....	29
Bibliography & References .....	29

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 Entertainment / Media Revenue and Gaming - 2017 View.....	5
Figure 2 Gaming Revenue Growth Y-o-Y and Device Distribution .....	5
Figure 3 Gamer statistics. source - <a href="https://www.theesa.com">https://www.theesa.com</a> .....	7
Figure 4 A 1000 pings jitter sample.....	9

Figure 5 Human Response Time .....	12
Figure 6 CS;GO scoreboard .....	13
Figure 7 Wifi Building blocks .....	17
Figure 8 Game Signaling .....	24
Figure 9 Stadia Bandwidth Usage from Google.....	27

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 Game Engine Examples.....	11

# Introduction

There are people who play games and there are people who watch games. While it is debatable whether Video Gaming is a sport, it is a fact that Video Gaming rivals traditional media as a form of entertainment, whether offline or online (i.e. available over the internet). As Netflix stated in its shareholder report that “We compete with (and lose to) Fortnite more than HBO” (see ref [1]).

As online gaming (mobile, PC, console, etc.) continues to exponentially increase (see ref[10]), the network performance (not just availability) is critical for the superior game experience and in particular to action games such as FPS (first-person shooters) and e-sports in multi-player mode. Arguably, network performance may be more critical for cloud gaming such as Google Stadia, since the gaming experience solely relies on the cloud (for almost all of the processing) over the network. The network performance (mainly, bandwidth) is also somewhat critical while broadcasting one’s game in real-time on any of popular social platforms e.g. twitch, youtube etc. for others to watch (note that online game watching is the second most popular viewing with ~100 million viewers, more popular than MLB, NBA or NHL per ref[12]).

This paper covers the challenges in getting an excellent online game experience for action games, not only from the game developer and players point of view, but also from the network point of view. The paper does not focus much on game streaming (e.g. Twitch) or Cloud Gaming (e.g. Stadia).


As we discuss in the paper, the online game experience wrt the network is multi-domain by nature – home network domain, service provider network domain, Internet domain and the data center network domain (where the game servers are located) are all important parts, each one with its own set of technical and non-technical issues. Of course, there are other challenges in ensuring superior game experience outside the network, such as the game engine, match making, rendering rate of a graphics card etc., but these are not the focus of this paper.

We see a trend of service providers marketing “low latency” as a differentiator that is essential for gameplay. They equate “Lag” to “Latency”. It intuitively makes sense. Isn’t multi-player gaming similar to a duel where the fastest to draw is likely to win? As we explore in this paper, the answer is more complex and depends on the type of lag compensation algorithms used by the game server, and what the gamer defines as “Lag” is not exactly what a network expert defines as “Latency”.

The above has been the trailer to our paper, now let the game begin!


# Overview

Online Video Gaming has attracted the attention of network engineers ever since the first multi-player online games appeared. Arguably, Online gaming has become one of the largest \$\$\$ businesses on the Internet since 2017 and growing with the high CAGR, as illustrated in the figure below (see ref[17]) -



**Figure 1 Entertainment / Media Revenue and Gaming - 2017 View**

According to NewZoo 2019 Global Games Report (see ref[18]), the Gaming revenue is expected to be around \$200B by 2022, as illustrated below –



**Figure 2 Gaming Revenue Growth Y-o-Y and Device Distribution**

When researching this area, it's important to filter out papers and presentations that are more than a couple of years old, because of the advances that have been made in game development. In other words, it is safe to ignore papers on "Quake III" performance that date back more than 5 years ago. They do have value, but they don't represent the state of the art.

# 1. Classification of Games

There are different types of video games, categorized by their characteristics or underlying objectives. Game developers and publishers categorize their game titles accordingly. The popular categories are listed below, please see the complete listing here (ref[11]):


- Strategy
- Action
- Adventure
- Sports
- Simulation
- Board

The categories or genres can also have subgenres, and games could fit into multiple genres! For example, Action genre has multiple subgenres – platform, shooter, fighting, stealth etc.

In many respect, MMORPGs (Massive Multiplayer Online Role-Playing Games) and Action/FPS (First-Person-Shooting) are among the most popular online games, and now attract millions of users who play in an evolving virtual world simultaneously over the Internet.

Suffice to say, not all genres require the same network performance. Action oriented games, for example, would require a lot more strict network performance (latency, bandwidth etc.) than Board games or strategy games. The latter categories do not require strict latency/jitter treatments from the network. Figure 1 shows the types of games and devices consumers use:





**Figure 3 Gamer statistics. source - <https://www.theesa.com>**

This paper focuses on multi-player shooter type games because those are:

- (a) Among the most popular.
- (b) They require the strictest network performance.

There are of course many other types of games, and even the multi-player games may have a standalone mode with AI (artificial intelligence) avatars instead of human competitors.

Another type of multiplayer game is the live-arena gamer, e.g. the live-arena built by Comcast (ref [1]) or other venues that host live games and are becoming popular. However, these are connected with a local network and are not the focus of this paper either.

## 2. Game Downloads

As a sidenote on network load, the game downloads can put stress on the network as well. Most games are sold via download now, and some games are huge. Call of Duty: Black Ops 3 is 101GB, and Grand Theft Auto V is 65GB. In contrast, an hour of 4K video on Netflix is about

7GB per hour, making a Call of Duty download equivalent to watching over 14 hours of 4K video!

For these game downloads, the network does play a role, but its limited to downstream speed (in bps), since new editions of the games could take a lot longer depending on the speed. While standard caching architectures can possibly handle these game updates well, and those are not the focus of the paper, it is important to highlight that game downloads can affect ISP network utilization far worse than anything else out on the internet – Windows downloads (~3.5GB in case of Win10), iOS downloads (~2GB in case of iOS 12), macOS downloads (~6GB for mohave), 4k movie download (7GB per hour) etc.

Also note that games (similar to mobile apps) get updated frequently and these updates could be multiple of GBs each (e.g. League of Legends minor update in Sept'2019 was ~2GB). In other words, just a minor update of a game could be more than the entire mobile/laptop OS download.

Imagine the network contention that may arise if 50 users in a domain are downloading the latest game edition, while 10 users in the same domain are playing the game. In a given network domain, this may even cause a congestion.

Network QoS should be considered to deal with any potential network contention when new games or new versions comes out.

### **3. The Roles of the Game Server**

Some assume that the game server acts as a directory used for the initial game setup and that following the initial setup the gamers play peer-to-peer. This is not entirely correct. As described in the section 5 (“lag compensation”), the game servers play a critical role in processing the “world view” messages that the clients send them and gamers connect to the server, not to each other (for security and privacy reasons as well).

The game server provide “match making” to pair up gamers correctly. For some game companies, the match making algorithm is a closely guarded secret, however generally speaking, these two criteria play a role in the matching:

- Skill level: game companies want to keep gamers engaged. If they pair a beginner with an expert, then the beginner is likely to be eliminated quickly, get frustrated and stop playing the game. Such a pairing is not enjoyable for the experienced player either.
- Latency: By pairing gamers with similar latency, it's easier to place everyone on the same timeline and reduce instance of “shooting behind the corner” (see more in section 5).


For effective match making, its best to have access to as many players as possible. This means that many massive multiplayer online (MMOs) gaming providers are unlikely to place their game servers close to the clients. There is a price to be paid for that flexibility, as we discuss in peering (section 15).

Some games can benefit from placing a server closer to the user, in particular, games that are geo-local such as “Pokémon GO”, however, these are not the focus of this paper.

#### 4. How Gamers Perceive “Lag” vs. How Network Engineering View It

Gamers often talk about “lag” when their avatar is not responding well to commands or does not move smoothly. Some assume lag is the same as latency, but it’s not. “Lag” is a measure of game experience, and all of the following network impairments may result in what a gamer perceives as lag:

- Packet drops
- Packet jitter
- Packet latency



**Figure 4 A 1000 pings jitter sample**

To get a sense of what “jitter” means let’s take a look at Figure 2 which a sample of 1000 pings over a cable network. It’s meant illustrate what “jitter” looks like and how it can impact games without getting too much into the details of the reasons this particular sample looks the way it does. In a perfect world the ping time would be constant, but here we see its hovering around 10ms-15ms. As long as the ping time remains in this range one would expect a smooth playing experience. However, each jump over 30ms may represent an instant where the gamer experienced “lag”.

As we explore in section 5, a stable latency is the least damaging for game play because lag compensation can deal with it well. It is packet drops and packet jitter that are the most damaging, therefor an obvious first step for improving gaming experience would be to make sure

that the gamer's upstream and downstream connections are clean before deploying more advanced techniques to improve gameplay.

## 5. Lag Compensation

Let's consider the simplest case of two gamers, each with a 10ms delay to the game server. For our discussion we can assume no packet drops and no jitter. We would obviously love to have no latency, but that's physically impossible: the images are rendered 30 or 60 times a second (33ms or 16ms delay), packet forwarding over the network is limited by the speed of light (in practice, speed of light is 1/3 slower in optical fiber (silica glass) based network links), routing hops and more. But for our example, we stick to our 10ms figure. Here is the basic issue:

- Player A sees an image of player B and shoots at it.
- In reality what player A sees is player B position 20ms in the past
- What if player B is not in the path of the bullet anymore?

In order to know whether or not to register a hit, the server uses a set of algorithms called “lag compensation”. The general idea is to establish a common timeline in order to decide who ended up hitting who. There are three timelines to consider: shooter timeline, target timeline and server timeline. In general, preference is given to the shooter's point of view, most likely because if one sees the target avatar in their cross-hairs when they pull the trigger, then a hit is expected. A good discussion of lag compensation is covered in a videoclip from Blizzard in ref [3].

In a real network with packet loss and jitter, some movements need to be predicted because the player position updates may be lost or jittered too much and in any case are not synchronized with the frame rate. In such a case the players position has to be assumed and later verified. If the prediction does not work well, it will result in an artifact called “rubber-banding” where an avatar appears to be teleporting back and forth because its assumed position is replaced by an updated one that is significantly different (if the prediction was good, then the update would not change much and gameplay would be smooth). Similar prediction algorithms are deployed on the client side as well to make movement appear smooth with the same risk of rubber-banding that may occur every once in a while.

The lag compensation algorithm highlights a problem that occurs with any distributed system and is described by the “*consistency, availability, and partition tolerance* (CAP) theorem” (see reference [6]). Simply put, the CAP theorem proves that it's not possible to have a distributed system that is fault tolerant, consistent and high performing all at the same time. However, it is possible if two of the three are chosen. In gaming, the system is desired to be high performance (meaning low latency) and fault tolerant (meaning surviving packet drops and jitter) and so have to resort to “eventual consistency” – and that means prediction and lag compensation.

The combination of prediction and lag compensation can result in an artifact called “shoot-behind-the-corner” that may actually give an **advantage to the player with the higher latency**. The way it works is that the player with the higher latency can hide behind the corner, shoot a target and hide again. The target player may never see the attacking player. There is even a

hacking tool called “premium lag” (<https://premiumlag.com>) that advertises the following: “A lag switch works by cutting off your outgoing data without cutting your incoming data. When you hold the button, you are essentially off the radar and your character appears frozen to other players.” Because of the way the lag compensation works, once the lag switch is on again, the hits would be registered with the server.

Having said that, it is important to note that lag compensation makes sure that the shot was accurate but does not change who shot first. Assuming both players shot correctly, the one that shot first still has an advantage.

Table 1 provides a small sample of game engines used in multi-player games, though there is a surprisingly large number of these game engines. Most game engines employ lag compensation algorithm(s) so as to normalize server-side state (of the game) for each player as that player's user commands are executed. See ref[19] and [20] for detailed insight. Interestingly, each one of them might react to network conditions differently, as well as deploy different optimizations to deal with Lag Compensation, however, the basic issues that we outlined are fundamental to any distributed system and the appropriate solution/s need to be implemented on any game engine.


**Table 1 Game Engine Examples**

Game Engine	Sample game
Unreal 4.0	Fornite
Unity	Rust
CS;GO	Source
AnvilNext	Rainbow 6
Riot Games Engine	League of legends
MT framework	Monster hunter

The key take away from all the above is that **latency is not the primary experience killer for games, rather jitter/drops are**, because lag compensation can deal with fixed latency quite well (up to a point), and that’s in addition to the “match making” that pairs players with similar latency. On the other hand, it is packet drops and jitter that are more directly related to what gamers experience as “lag” (avatar not responding or rubber-banding) because jitter/drops are the artifacts that cause lag compensation to miss on its prediction algorithms.

## 6. Human Response Time

When discussing latency, it’s good to account for the human response time in the overall equation. By response time, we are referring to how quickly users react to changes they see on the screen. A fun little exercise is to measure response time using the website listed in reference [8]. Note that the website runs the app directly on the browser so there is no network latency in this measurement. The website states that a slow PC/browser can add 10ms-50ms to the measurement:



**Figure 5 Human Response Time**

The website states that from the statistics collected across 80 million clicks, the median reaction time is 273 milliseconds. The average reaction time is 284 milliseconds.

This means that small latency differences, in the range of 10ms or so, are in the statistical noise of human response time and do not give a significant advantage. Having said that, two issues need to be taken into account:

- There is a mismatch between the times it takes us to detect a change vs. how fast we can respond to the change : while response time is in the 250ms range, the human eye does perceive movement at 30-60 frames per second (33ms/16ms). Therefore, smaller latency makes the lag compensation algorithm work better and as a result, provide a more consistent and smoother gameplay experience.
- Delay/jitter can be additive, so reducing the delay in small amounts each domain (CM, home gateway, WAN, etc) can help for the end-to-end delay budget.

## 7. How is the Network Performance Profiled?

Game developers use ping from the client to the game server in order to estimate RTT and report ping times for all participants of a game. This means that if one uses fancy classification techniques to divert gaming traffic to a dedicated service flow it might not be detected by the game engine if the pings go on a default traffic path.

How are ping times related to gameplay ? Generally speaking, gamers tend to align ping times around a 30ms quanta that corresponds to a 30 frames-per-second updated.

- Excellent latency: anything below 30ms
- Acceptable latency: anything between 30ms-60ms
- Playable latency: 60ms-90ms

- Bad: 90ms and above, though depending on the game and lag compensation algorithms as much as 150ms is deemed “acceptable”

Looking into source code of Unreal 4.0 (see ref [7], Unreal 4.0 is the open source game engine for games such as Fortnite), one can see that that particular game engine refers to ping times as “QoS” - which is quite different from what network engineers refer to as QoS. See (ref[11]) for the latter.

Given the impact that drop/jitter have on gameplay, the reliance on ping as a performance metric is painting an incomplete picture, though it’s understandable that game developers want to reduce the complex issue of network performance to a single well-understood number. One should note that to some extent, ping time is a “marketing” number – the same way that comparing CPUs based on MHz rating paints only a partial number of a CPU performance, but at the same time is an easy to understand figure.

Figure 4 is a capture of a CS;GO scoreboard (based on the “source” game engine). Note that the “ping time” is a column in the scoreboard (client to server) and using a special debug command, a user can view their overall ping time (roundtrip).




Figure 6 CS;GO scoreboard

## 8. Packet Transport

Gaming traffic is mostly UDP – whether video or audio.

Game engines use UDP for transporting “world view” coordinates between clients and servers. TCP does not make sense for games because:

- a. Games do not need bulk transfers so all of TCPs abilities to sense the network capacity and optimize transfer rates are not useful at all for gaming.
- b. Games don’t need a reliable transport. If a player coordinates are lost, then it makes more sense to transmit more recent update (and rely on prediction by the lag compensation algorithm) than to re-transmit the old coordinates.

In addition to the video part of the game, there is a separate channel for voice which is UDP as well. Transfers of game updates are of course bulk transfers and use TCP.

Note that Cloud based Game streaming (see section 20) does not use TCP either, but a specific UDP based streaming protocol (originally developed by Google, now getting standardized as HTTP3.0 by the IETF) and WebRTC extensions.

## **9. Gaming Traffic Characteristics – Packet Sizes**

Gaming traffic flows can have many of the following characteristics, depending on the game genre/sub-genre and depending on whether console or PC or mobile or Cloud delivered. For online games, the following characteristics are typical (yet to be confirmed for Cloud gaming such as Stadia):

- Long lived flows
- High packet rate
- Small and regular packet sizes
- Fairly regular packet inter-arrival times
- Bandwidth usage (Low to High depending on the game)

Note – League of Legends (one of popular games) comprises constant bitrate with small packet sizes (~55Bytes). However, other games such as the XBOX battlefield can have packet sizes around ~700Bytes.

It is interesting to note that downstream:upstream ratio is proportional to the number of players. For example, in case of 16 player game, the downstream:upstream traffic ratio could be close to 16:1. In case of 64 player game, the ratio could be close to 64:1. The reason is that the higher the number of players, a lot more info (changes) for server (since it has to aggregate all the clients’ actions) to send to each player’s device.



# Domains & Recommendations

End-to-end game performance is a multi-domain problem. This section will outline the domains involved in game performance. The delay/jitter/drops that we discuss in the following section is additive, and in some cases even if one domain still meets reasonable performance criteria, the combined effect of several domains can be significant.

## 10. Game Rendering

The compute resources needed to render a game are not the focus of this paper, but they are part of the overall game experience so for completeness we will overview them.

The tradeoff a gamer needs to make is cost of the hardware vs. the quality of the game animation. If a computer does not have a powerful enough CPU, GPU and memory then either the frame-rate, or resolution or both have to be reduced to keep gameplay responsive. If latency is caused by lack of computing horsepower on the client hardware, then there is nothing that lag compensation can do about it.

It is a little easier to manage the quality/responsiveness tradeoff in console games, as opposed to PC games, because the game can be optimized for a specific platform (the console).

Another issue with PC games is that other applications might be running in the background, and a periodic software update or system backup might start in the middle of a game. While easy to fix, these are some of problems that we have to add to the list of things to debug when trying to assure end-to-end game experience and are clearly outside the scope of the Service Provider world.

It's worth noting that most Gamers playing action type games are savvy consumers and can usually resolve the issues that are within their control on their own. However, other types of gamers would not be as savvy and may expect gaming provider or Service Provider to look after their needs and assurance.

## 11. Home Network Domain

The home network is one of the more hostile environments for gamers. WiFi can incur delays of 40ms and have a large percent of packet drops. It is true that serious gamers are recommended to use wired connectivity (i.e. connect the gaming device directly to their home router using an ethernet cable), but in reality, many don't.

In addition, if the home gamer shares the connection with other people in the same household they may run into congestion and buffer-bloat issues in the home (before even getting to the cable modems). Some companies offer home gateways that can mitigate these issues and those seem to be working to reduce ping times and improve game performance. These solutions have built in buffer-bloat management algorithms.

As a side note, one should not confuse the home gateway and a cable modem. The cable modem handles the DOCSIS protocol, which is basically the interface between the CMTS and the CM,

while the home gateway typically performs functions such as NAT and WiFi access. Even in cases where both the CM and the home gateway are packaged in the same enclosure, we should still treat them as functionally different entities.

All the above is to make a clear distinction between applying buffer-bloat mitigation at the home gateway vs. doing the same at the cable modem as we will explore in section 13.2.

#### RECOMMENDATIONS:

1. Prefer Wired connectivity to WiFi, if possible
2. Prefer 802.11ax (or 802.11ac) if using WiFi connectivity, given the scheduled mode support
3. Prefer 5Ghz or higher (more number of non-overlapping and wider channels, lower contention); However, resort to 2.4Ghz non-overlapping channels (1,6,11), if possible, in case of having brick walls between AP and gaming device.
4. Use routers that avoid bufferbloat issue
5. Disable unused WiFi modes (such as 802.11a,b,g,n).
6. Consider dual-channel WiFi that can separate out Tx and Rx on different channels. See more details on cablelabs' work in this area (reference [16]).
7. Avoid Mesh WiFi solutions, given the latency impact

### 11.1. WiFi – WiFi5 vs WiFi6

Home network commonly employs (802.11 standards based) wireless connectivity to access the network and consume services/content, given the sheer convenience factor. Interestingly, 802.11 standards have evolved quite a bit over 2 decades :

- 802.11b (WiFi1\*), released in 1999
- 802.11a (WiFi2\*), released in 1999
- 802.11g (WiFi3\*), released in 2003
- 802.11n (WiFi4), released in 2009
- 802.11ac (WiFi5), released in 2014
- 802.11ax (WiFi6), being released in 2019

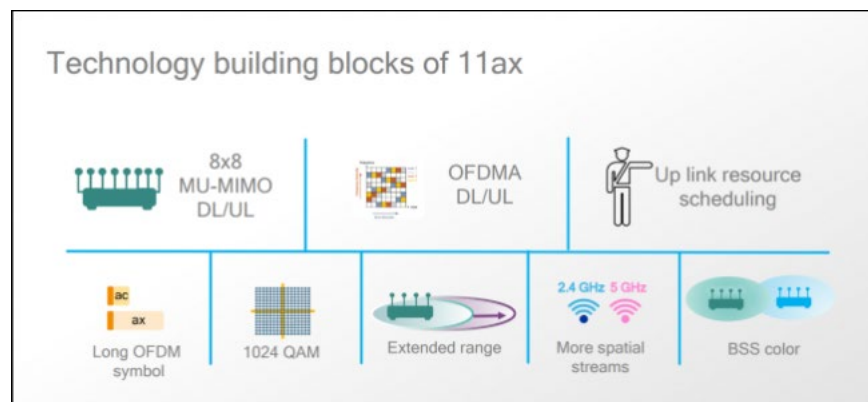
\*Not official naming,

However, despite the evolution, 802.11 wireless connectivity at home has remained somewhat problematic for superior gaming experience in certain cases, due to indeterministic latency and drops.

A major cause of latency in WiFi is that it's a carrier sense system and when many end stations compete for bandwidth, it can experience low utilization and high latency.

WiFi6 solves this latency problem by having a “scheduled” mode where it can predict bandwidth demands and schedule transmission time in advance and thereby avoiding issues with congestion and carrier sensing.

How does a WiFi6 detect an application need for bandwidth ? Well, the standard itself does not provide any means, but it can be as simple as setting fixed transmission time for all packets directed to a specific game server and by doing that assuring a dedicate channel/slot for game traffic. This can be one of the “advanced” options in the WiFi6 router settings and something that most serious gamers should be savvy enough to do.



**Figure 7 Wifi Building blocks**

Figure 5 depicts the base WiFi6 building blocks, see ref [13]

## 12. SP Network Domain – First Mile

First mile is the lag of the network that connects a home to the first active outside the home. Since this paper is written for a cable conference, we will focus on the DOCSIS drop/jitter/latency. As discussed in section 5, it is the drop and jitter that damage game experience more than the fixed latency.

Packet drops in cable networks are typically caused by RF impairments, and the first phase for improving a game experience would be to make sure the RF part of the plant performs well. Remember that lag compensation can mask packet drops up to a point and if it happens to predict well, the game experience might still be reasonable, else artifacts such as “rubber-banding” will creep up.

Jitter in cable networks could occur because of contention slots. In order to send a request, the CM has to first contend for a “DOCSIS contention slot” to send the request. The time it takes to acquire a contention slot is random, but it gets worse, as the utilization gets heavier. In bulk transfers, the impact of contention is secondary because the DOCSIS protocol allows additional bandwidth requests to be piggybacked, however in short transactional transfers, such as the ones used to transfer a game “world view”, the contention channel is the primary method for sending bandwidth requests.

In principle, buffer bloat and congestion on the access could cause jitter as well, but it’s a topic of further study whether the above manifests itself as a relatively stable delay or as jitter.

#### RECOMMENDATIONS:

1. Keep the RF part of the cable plant clean
2. A dedicated DOCSIS service flow would assure game traffic prioritization and easier debug.
3. Follow some or all of the Low Latency DOCSIS recommendations (at least reduced map times and proactive scheduling).
4. Reduce congestion, either by increasing upstream rates or node splits

### 13. Achieving Low Latency on a DOCSIS Network

When analyzing any system for delay the following contributing factors need to be considered:

- Propagation delay: typically, the speed of light, or in fiber 2/3 of the speed of light
- Transmission delay (serialization/encoding): the time it takes to send the bytes to the wire (e.g. serializing a 1500 bytes packet on a 1Gbps link will take 12 microseconds)
- Processing delay (media acquisition): any computation time it takes to process the packet, for example, calculating a CRC.
- Queueing delay: once the packet is queued for transmission how long it may end up waiting in the queue.

#### 13.1. Sources of Delay in a DOCSIS Network

Analyzing all the delay elements in DOCSIS could be an SCTE paper in its own right (and papers have been written on the topic, see ref [14]), so we will keep the discussion at a high-level, ignore the 2<sup>nd</sup> order artifacts and focus on the elements that are most relevant when comparing DOCSIS to other technologies (e.g. PON).

DOCSIS Media Acquisition Delay: Because DOCSIS is a request/grant/data system, the propagation delay is multiplied by 3:

- The time it takes the request to propagate in the upstream direction,

- Then the grant to propagate in the downstream direction
- Followed by the actual packet in the upstream direction.

DOCSIS 3.0 made a significant improvement because it provides the ability to “pipeline” the requests. So, for a large file-transfer, the request/grant delay becomes relatively negligible. For example, if it takes 200ms to load a web page, then an initial delay of 5ms is not that significant. To illustrate what pipelining means, imagine the following scenario: say a person goes grocery shopping in a store that is a 10-minute drive from their home. In the pre-DOCSIS 3.0 case every item purchased would require a separate trip. Buying 5 apples would require 5 trips and a total roundtrip of  $10 \times 2 \times 5 = 100$  minutes.

And here comes the interesting part of the analogy...one might expect that DOCSIS 3.0 is analogous to packing the 5 apples into the car and finishing the transfer in one roundtrip, but a better analogy for DOCSIS 3.0 would be sending 5 cars at the same time, each carrying a single apple. The result however is the same, all the groceries would arrive in 20 minutes (roundtrip time). The above holds true for longer bulk transfers, but for short transactions, the delay is additive since there is no “pipelining” and short transactions are relevant to gaming. Back to our example, if you were really hungry for an apple you would not care that after 20 minutes you could get 5 apples, you just want one as quickly as possible.

DOCSIS Propagation Delay: as transmission rates over cable get faster the transmission delay is in the range of microseconds and not a significant contributor. Having said that the transmission over the fiber part of the network can be in the range of milliseconds depending on the length of the fiber part of the cable plant.

DOCSIS Media Acquisition Delay: the most significant contributor to processing delay is the downstream interweaver, usually in the range on 2ms or so. It can be made shorter but at the risk of reducing the fidelity of the transmission.

DOCSIS Queuing Delay: Queuing is how long a packet waits in a queue before it is served. DOCSIS has tools to deal with queueing delay by sorting packets to flows so that the high-priority flows get services quickly before they form a long queue (analogous to the faster boarding service for business class in an airport). In principle, gaming packets could be classified to a dedicated high-priority service flow to assure that queuing delay is minimized. Another option is the Dual-Queue part of LLD (low latency DOCSIS) which we will explore in more detail in the following sections.

For more detail please see reference [14].

## **13.2. CableLabs LLD**

In order to reduce latency over the cable plant, CableLabs has initiated the development of “low latency DOCSIS” (LLD). LLD is described in detail in reference [14], for our paper we will give a high-level overview of the two features LLD enables:

- *Proactive scheduling*: The simplest way to reduce the request-grant delay is to eliminate the need to send a request. Proactive scheduler is sending a stream of grants even if the modem is not requesting. It is like UGS (unsolicited grant service) but for data. It is a way to trade off bandwidth (since some of the unsolicited grants may be unused) for lower latency. It's worth noting that the mobile 5G standards advocate a similar technique to facilitate low latency over mobile, and similar technique is used in WiFi6. A way to mitigate the bandwidth loss due to unused grant is called "grant sharing" where a grant for service flow A can still be used for service flow B if service flow A happened to not use it.
- *Dual queue*: TCP traffic causes buffer buildup because the way TCP works is by sending as much data as it can until it notices packet drops. This is a good method for bulk transfers because it keeps the network full, but it triggers a phenomenon called "buffer bloat" which can cause excessive delay. The dual queue in LLD is a collection of methods to help keep buffers shallow and at the same time assure that "well behaved" traffic is rewarded and experiences less queueing delay. For a detailed description of Dual Queue see reference [TBD]
- *More frequent MAPs*: The typical MAP interval for many CMTS implementations is 2ms. This means that worst case the CMTS might have processed a request but has to wait 2ms before sending this grant in a map message. LLD specifies a preference for 1ms MAPs. By doing that it reduces the MAP wait time to 1ms.

### 13.3. DOCSIS Delay vs. PON delay

Both DOCSIS and PON (EPON/GPON) are point to multi-point technologies and are similar on the need for a request/grant/data cycles (even though these are called differently in PON it's a similar concept). So, why are the PON vendors marketing their networks as "low latency" ? Two key differences are:

- *No Contention in PON*: think of PON as RTPS (real-time polling service) only. There is no option to contend for bandwidth and all slots to carry bandwidth requests are pre-scheduled and dedicated to the end station.
- *Shorter processing delay*: there is no "downstream interleaving" and other physical layer related delays in PON because currently its only 0's and 1's with no modulation (aka "baseband"). This already cuts about 2ms from the delay
- *Immediate grants*: There are no periodic MAPs in PON. Every request can be granted individually so there is no MAP wait time. This can cut another 1ms-2ms.
- *Lower number of subscribers and higher bandwidth*: This is not a protocol advantage, but because PON cannot be split to a large number of subscribers (being "passive" it loses half the power for every split) we have a smaller number of subscribers and more bandwidth for each Service Group than what is typical for cable.

As we explored in this paper, we have several tools to reduce the DOCSIS delay getting within a few milliseconds of what PON can achieve. Furthermore, as is evident from our discussions on lag compensation, frame-rates and human response time, it is clear that a difference of single digit milliseconds is a non-issue for gameplay.

## **14. SP Network Domain – Second Mile**

In most networks, the “2<sup>nd</sup> mile” is the network lag that connects the aggregation point to the internet exchange (or peering point). In most networks the 2<sup>nd</sup> mile is not over-subscribed and typically not a source of congestion, jitter or drops. However, it may be important to prioritize gaming UDP traffic in case of congestions resulting from network link/node failures etc.

### **RECOMMENDATION**

- Keep it simple and keep the 2nd mile rich with bandwidth, monitor performance and buffer depth in the routers/switches that make the 2nd mile.
- Mark UDP gaming packets with a higher DSCP code point.

## **15. Internet Domain**

This is an area that is an unexpected source of pain in terms of latency and jitter. Here is why:

- The game server usually are centralized, so that it has the most flexibility for matching players (with the exception of localized games such as “Pokémon GO”). Some game servers have been located in the center of a continent to have equal distance from most gamers.
- Centralizing a server is not necessarily a problem, but because of the way the Internet is built it becomes a problem. What we call “Internet” is not the simple network cloud that we draw, but rather a collection of smaller networks, transports, peering points and service provider networks. Routing IP packets between all these networks is based on minimal cost does not always equate to minimal latency. Reference [4] includes an excellent discussion showing how Riot Games built their own network in order to avoid routes that are minimal cost optimized but not latency optimized. Even at fiber speeds, these inefficient routes can add up to 10’s of milliseconds of delay.
- The interesting challenge is that there isn’t any big Gaming traffic Aggregators (similar to how content aggregators mushroomed a decade ago or so), nor any dedicated gaming exchange. Perhaps, an opportunity for Internet Exchange Points (IXPs) to become Gaming Exchange Points (GXPs). Else, ISPs and Gaming Providers will need to have direct or indirect peering that is latency optimized.

In general, it is beneficial for ISPs to have dedicated peering, if possible, to offload gaming traffic, whether via the gaming exchanges or via direct connectivity, to the networks having the optimal routing (latency, hops etc.) to the data centers hosting the gaming servers.

There are several strategies and several companies with products that can help with the above. They generally fall into one of three categories:

1. *VPN solutions*: by encapsulating game traffic into a VPN it's possible to steer it into the best peering point and may help divert traffic to a good path.
2. *Overlay solutions*: these solutions take advantage of the fact that the Internet is more than a simple "network cloud". There are many data centers in various places in the Internet and traffic can be routed from one data center to another thereby creating an overlay that effectively overrides packet routing decisions based on "minimal cost".
3. *Custom-built networks and custom-built edge solutions*: some game houses host the game servers on public cloud solution and these public cloud solutions can have their own network optimizations. In other cases the game provider will create its own peering network (see reference [5] presentation Blizzard networking solution as well as reference [4] already mentioned).

## RECOMMENDATIONS

- Use one of the solutions mentioned above to optimize the IP traffic paths between the game server and end devices.
- Seek and Implement direct and optimized peering relationships with gaming providers.

## 16. Data Center Domain

The game servers are hosted in a cloud, either public or private, and cloud applications are subject to delay/drop/jitter like any application.

The data center is a dynamic environment where compute loads can be moved at any point and each such move can cause delay/jitter/drops on the server side. A private data center may be easier to control, and even in a public data center it is possible to pay more to get a higher service assurance but clearly there is cost to both options.

## RECOMMENDATION

- Deploy an application monitoring system in the data center to track how well the application is running.
- Perform Cost analysis of server performances vs. data center options (private/public/higher service tier)



# Gaming as a Managed Service

Establishing a dedicated game connection can be viewed as similar to connectivity services sold to business subscribers. The basic tools that DOCSIS provides to build such a connection are the packet classification rules, DOCSIS service flows, and a policy framework. The following section will go into the detail of creating a “game connection”.

Note that if game packets are directed to a dedicated DOCSIS service flow, then there is little need for active queue management (such as dual queue which is part of LLD) because there will be no buffer bloat in a queue that serves only the gaming data. In addition to that, a dedicated service flow can use RTPS/nRTPS or dedicated predictive scheduling which will further reduce jitter because no contention slots will be used.

## 17. Classification

Game traffic needs to be classified in order to be directed to a DOCSIS service flow. As we discussed in section 11, the home gateway is a separate entity then DOCSIS and in the home gateway a user may directly program any packet with a destination of a game server to have “special treatment” but there is no such option for the cable modem, so how can it be done ?

PCMM (Packet Cable MultiMedia) is a solution where a user application can request QoS for the DOCSIS network. Having said that, PCMM was not a commercial success and we use PCMM as an example for a policy framework. In the solution outlined below, PCMM is meant as a reference that’s understood by cable operators. However, a wireless policy framework such as PCRF can be used just as well.

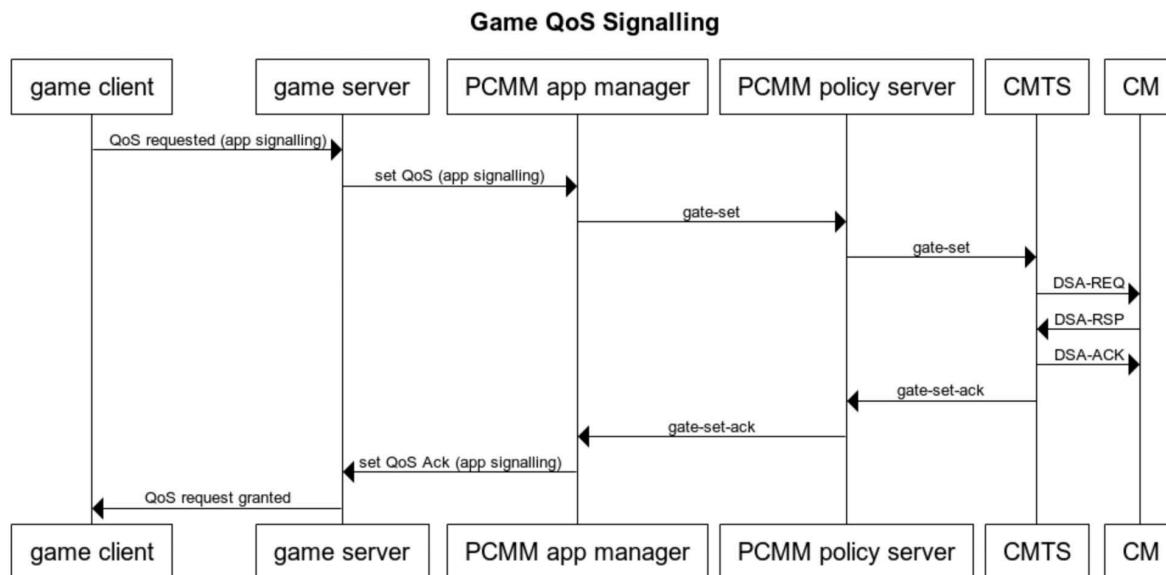
The proposal below puts emphasis on game server to cable network interaction as opposed to game client to cable network. It may be intuitive to have the client send a request for a dedicated service flow, but in this proposal, the game server is the one to signal the exact classifier for the following reasons:

- It’s easier to establish trust between the game server and a policy server than between the home client and the policy server. This is for two reasons : (a) the game company already establishes trust with clients, filtering out known abusers, so a request coming on behalf of a client is already sanitized, and (b) easier to build trust with a small number of business entities on the server side then with a huge number of home clients.
- The server sees UDP packets post-NAT. The client side can report packet pre-NAT. This way, we don’t have to deal with NAT issues since we have the full picture (when we view it from the server side)

The following ladder diagram shows how this proposal can work with PCMM; but it’s only a reference so we can have a concrete example. In this day and age of micro-services, it’s easy to build a function that just sends common open policy service (COPS) messages and has a different policy framework (possibly a PCF/PCRF) on top of it.

This proposal does require some changes to the game client. By way of inspecting the open source code for Unreal 4.0 (one for the most popular game engine, which happens to be open source) it seems possible.

Figure 6 depicts a proposed exchange between a policy framework (PCMM in this example, based on reference [9]):



**Figure 8 Game Signaling**

The initial signaling is between the game client and the game server. In fact, it can be part of the normal announcement that a gamer is joining the game and not a dedicated message.

The game server is in the best position to validate that the client joining is not a bot or was blacklisted for any reasons (such as cheating or flagged by other gamers). If the client is approved, the game server will signal the PCMM application manager a request for QoS, and the application manager will forward the request to a policy server for approval, this is all in accordance with the PCMM specifications.

From this point on we are dealing with the DOCSIS domain and a standard DOCSIS request for creating a service flow. We are focusing on the “success” case where QoS is allocated and a positive indication is propagated all the way back to the game client.

As outlined in Section 7, if the game is using pings to measure network performance it is desirable to classify pings to the same flow as the actual gaming traffic.

Since the flow created is designed to carry only game traffic, there should be no need for queue management because the game traffic itself, being UDP based and at fixed interval, will not cause buffer bloat.

## 18. Gaming as a Marketing Play

The marketing departments in any big operators are aware of the interest gamers have in “ping times” and in several cases use it as a sales tool, either to promote specific products such as “gaming package” or as a selling point against other operators. The latter is a point made by PON providers against cable providers. As we explored in this paper, the technical merits of PON against cable when it comes to latency are not significant. However, as a marketing tool, it can be effective to claim an advantage.

From a quick survey of “gaming packages” offered by cable operators, it seems like they fall into two categories in the moment:

- Higher speed services that are packaged as “better for gaming” because they reduce congestion at the home
- Collaboration with a game optimization vendor that is packaged with the cable provider offering.

## 19. Assuring gaming performance across domains

Once a packet is classified to a flow, it can be marked in a way that can make it recognizable as a gaming packet across domains. For example, the CMTS can place it in a particular VRF or VPN. It can also help with preventing attacks on the game because if an authentic game flow is marked, then it’s possible to treat non-identified flows as “suspicious”.

Defining an API between the game server and the network for assurance can help as well. The API can help isolate networking problems even if the caller of the API is not a networking expert. For example, if traffic is sorted to flows and these flows associated with a particular user, then packet counts can be compared across the path to isolate domains where packets are dropped.

Because there are many game developers and many service providers, it may call for a “middleman” between the two organizations to help define and operate these APIs.

## B2B vs B2C monetization

As outlined in reference [4], game developers would invest in improving network performance to the point of building their own network. This means that instead of a B2C (business to consumer) monetization strategy of “have a customer pay X to improve gaming experience”, it may make sense to have a B2B (business to business) monetization strategy with the game developer. This has the following advantages:

1. Instead of a customer paying twice, once to “improve gaming” and once to a specific game developer, there is only a single payment to a game developer. One has to think about it from the perspective of a 15-year-old asking a parent for permission to buy a game for X dollars alone as opposed to asking for X dollars for game and then an extra Y dollars to the Service Provider. Note that the monetization can be per-use of the game.

2. As outlined throughout this document, the game networking performance is a multi-domain problem and the access network is only one part of it, and in actuality might not even be the long pole in the end-to-end performance of the game. When a cable provider charges money for “game performance” it implicitly assumes responsibility for end-to-end game performance and if the game experience is bad because of issues in a different domain then the toxic reviews on Reddit will follow very quickly. It’s the game developer that is in a better position to assume responsibility of the end-to-end game performance, and for them to craft an agreement with each of the domains they cross to give them a competitive edge.
3. If the APIs are changed its easier to update the server side, than force an upgrade to the clients.

Issues pertaining to net-neutrality and specifically to the relative advantages or disadvantages of a B2B vs. B2C from a legal viewpoint are outside the scope of this paper.

## Future topics

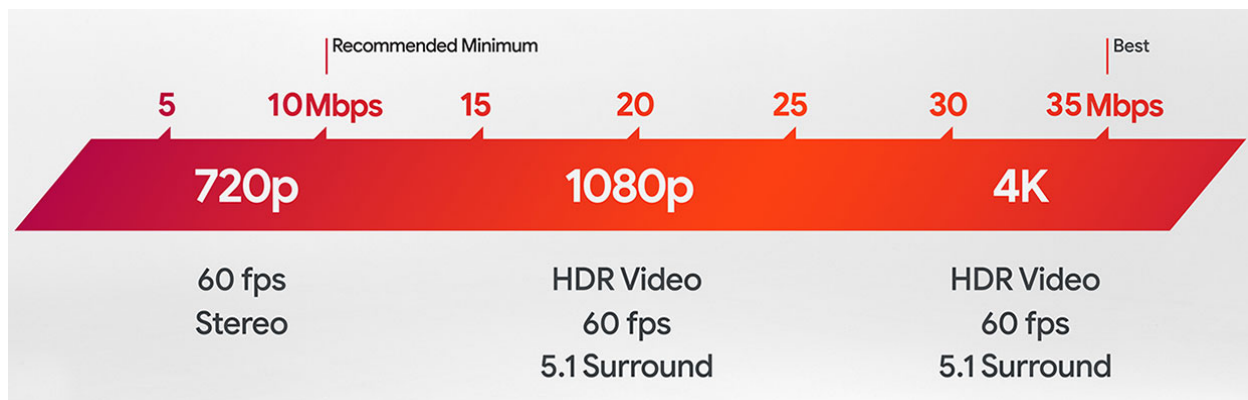
### 20. Cloud Game streaming

Recently, Google, Microsoft, etc. have announced their cloud game streaming platforms – Stadia, xCloud etc. respectively. Per Google, “*Our vision is to have Stadia available on all devices that stream YouTube—a truly platform-agnostic service.*” Stadia has the potential to become a popular alternative to game consoles since it could make playing games as simple as watching on-demand videos by clicking a button on any device.

Stadia promises that the game could start in less than five seconds after clicking on a link: no download, no patch, no install, no updates and in many cases, no hardware required. Just streaming, taking advantage of content distribution network technologies that most cable service providers have become well accustomed to.

Suffice to say, game streaming platforms such as Google Stadia demand specific network performance – around 15GB per hour of 4K game play (which is more than 2x than watching a 4K movie). Wrt connection speed:

- 1.5Mbps+ upload speed and 10 Mbps+ download speed for 720p resolution at 60 fps.
- 1.5Mbps+ upload speed and 20 Mbps+ download speed for 1080p resolution at 60 fps.
- 1.5Mbps+ upload speed and 35 Mbps+ download speed for 4K resolution at 60 fps.
- Stable/deterministic network latency.



**Figure 9 Stadia Bandwidth Usage from [Google](#)**

The following are very interesting points gleaned so far wrt Google Stadia:

- Stadia does “game streaming” which means that instead of rendering on a local console, it renders the images in the cloud and streams them as video to the consumer. Traditional console games do not take a lot of bandwidth (since they essentially communicate coordinates, not images). In contrast, game streaming can stream tens of megabytes. It’s definitely going to be a stress on the SP network if Stadia becomes successful. Some characterize Stadia as “Netflix for games”.
- Stadia may leverage AI based frame-by-frame latency prediction and be able to respond to gamers faster than their eyes could perceive the responses. This may indicate a sub-13ms latency, which is pretty astounding.
- At the time of this writing, Stadia is a new service and based on announcements made in E3 it targets single player games. Therefore, many of the issues discussed in this paper are secondary for this initial phase.
- There are people who play games and there are people who watch games. More watchers than players, obviously. Google’s vision is for Stadia platform to converge these two worlds together so that one can be watching a game, click and be playing a game or vice-versa.
- Stadia does require a Google controller, since the rendering is done in the cloud (around \$69, relative to \$300+ for an XBox). More details described at [15]
- Stadia can use the Google network, data center locations and thousands of network edges that Google manages around the globe to source the streams. Having said that, the last mile is naturally still in the SP domain.

# Conclusions

The service provider is in a unique position when supporting consumer gaming. The consumer sees the SP as the “one throat to choke” for any problem that is perceived to be network problem. But in the real world, the SP controls only network domains - the access, aggregation, possibly the core and possibly a part of the home domain. In some cases, these domains might not even be the most significant culprits during a degraded gaming experience. This basic catch is only going to be exacerbated if the SP charges fees directly from the consumer for an “advanced gaming service”.

How can we improve the experience for gamers and at the same time reduce the OPEX of debugging issues when the gaming experience is not good?

First of all, what can be improved in the access domain should be improved. Better RF will reduce packet drops. Dedicated gaming service flows or better queuing at the cable modem (such as LLD) will help reduce delay and jitter as well as make it easier to debug networking problems.

Secondly, collaboration with the game developers that will allow server-side probing for debugging of game issues can help reduce the OPEX required to debug experience issues. This can either be done directly or through a middleman between the service provider and the game developer. The key is the B2B relations between the game developer and service provider where it's easier to establish trust and business relations than with the B2C model on the client side.

# Abbreviations

API	Application Programmatic Interface
B2B	Business to Business
B2C	Business to Consumers
PPS	Packets Per Second
BPS	Bits Per Second
PCMM	Packet Cable Multi Media
QoS	Quality of Service
QoE	Quality of Experience
VRF	Virtual Routing and Forwarding
VPN	Virtual Private Network
CCAP	Converged Cable Access Platform
CMTS	Cable Modem Termination System
LLD	Low Latency DOCSIS
NAT	Network Address Translation
COPS	Common Open Policy Service
PCRF	Policy and Charging Rules Function
EPON	Ethernet Passive Optical Network
GPON	Gigabit Passive Optical Network
IXP	Internet Exchange Point

# Bibliography & References

- 1 Netflix shareholder report Q4 2018 (“competition” section, page 5):  
[https://s22.q4cdn.com/959853165/files/doc\\_financials/quarterly\\_reports/2018/q4/FINAL-Q418-Shareholder-Letter.pdf](https://s22.q4cdn.com/959853165/files/doc_financials/quarterly_reports/2018/q4/FINAL-Q418-Shareholder-Letter.pdf)
- 2 Comcast live arena : <https://www.nbcsports.com/philadelphia/fusion/fusion-arena-become-newest-state-art-gaming-facility-philadelphia-sports-complex>
- 3 Lag compensation explained by Blizzard:  
<https://www.youtube.com/watch?v=vTH2ZPgYujQ>
- 4 Riot games network: <https://qz.com/790208/how-the-company-behind-league-of-legends-rebuilt-its-own-internet-backbone-so-that-its-faster-for-gamers/>
- 5 Blizzard network @scale : <https://www.facebook.com/watch/?v=2090071161265977>
- 6 CAP theorem: <http://robertgreiner.com/2014/06/cap-theorem-explained/>
- 7 Github for QoS code in Unreal Engine 4.0: <https://github.com/soxueren/EpicGames-UnrealEngine/tree/59267dc158d4e919a579a98d472fbf21bb64508b/Engine/Plugins/Online/OnlineFramework/Source/Qos> - one needs a github account and an Epic games account to link to it.
- 8 Response time measurement : <https://www.humanbenchmark.com/tests/reactiontime>
- 9 Cablelabs PCMM <https://specification-search.cablelabs.com/packetcable-multimedia-specification>

- 10 Online Gaming Industry – Statistics & Facts <https://www.statista.com/topics/1551/online-gaming/>
- 11 Network QoS <https://www.cisco.com/c/en/us/products/ios-nx-os-software/quality-of-service-qos/index.html>
- 12 Online game viewing - <https://onlinebusiness.syr.edu/blog/esports-to-compete-with-traditional-sports/>
- 13 WiFi - <https://www.howtogeek.com/368332/wi-fi-6-what%E2%80%99s-different-and-why-it-matters/>
- 14 Low Latency DOCSIS, Greg White, SCTE EXPO 2019
- 15 Google Stadia [https://store.google.com/product/stadia\\_founders\\_edition](https://store.google.com/product/stadia_founders_edition)
- 16 Cablelabs Dual channel Wifi : <https://www.cablelabs.com/technologies/dual-channel-wi-fi>
- 17 Video Gaming vs Other Entertainment Revenue - <https://www.gamecrate.com/statistically-video-games-are-now-most-popular-and-profitable-form-entertainment/20087>
- 18 Global Gaming Revenue Growth - <https://newzoo.com/insights/articles/the-global-games-market-will-generate-152-1-billion-in-2019-as-the-u-s-overtakes-china-as-the-biggest-market>
- 19 Lag Compensation Algo - <https://enterprisecraftsmanship.com/posts/how-i-tried-to-get-into-game-development-and-failed/>
- 20 Lag Compensation Algo  
[https://developer.valvesoftware.com/wiki/Source\\_Multiplayer\\_Networking](https://developer.valvesoftware.com/wiki/Source_Multiplayer_Networking)



# Building a Cable-Friendly Internet of Things

A Technical Paper prepared for SCTE•ISBE by

**J. Clarke Stevens**

Principal Architect, Emerging Technologies  
Shaw Communications  
1401 Lawrence St, Suite 1550  
Denver, CO 80202  
720-723-2316  
clarke.stevens@sjrb.ca

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Executive Summary .....	4
1. Context .....	4
2. Operator Concerns.....	5
2.1. Security .....	5
2.2. Interoperability.....	6
2.3. Simple Development and Support.....	6
2.4. Other .....	7
3. Open Connectivity Foundation Features.....	8
3.1. Organization and Authorization.....	9
3.2. Security .....	10
3.2.1. Authentication .....	10
3.2.2. Authorization.....	10
3.2.3. Security Profiles .....	11
3.2.4. Onboarding .....	11
3.2.5. Secure Sharing .....	12
3.3. Interoperability.....	12
3.3.1. Data Models.....	12
3.3.2. Restful Architecture.....	14
3.3.3. Remote Access and Cloud Support.....	15
3.3.4. Bridging to Other Ecosystems .....	16
3.3.5. Common Management.....	16
3.3.6. Interop Events.....	16
3.4. Simple Development.....	17
3.4.1. Simple Device Description.....	17
3.4.2. Code Generation.....	18
3.4.3. Building the Application.....	18
3.4.4. Device Ownership.....	18
3.4.5. Introspection and Automatic User Interface Generation .....	18
3.4.6. Running, Testing and Debugging .....	19
4. OCF Development Process.....	19
4.1. Controlling the Server with DeviceSpy .....	20
4.2. Testing with the Compliance Test Tool .....	20
Conclusion .....	20
Abbreviations.....	21
Bibliography & References .....	21

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – A connectivity view of the Internet of Things .....	5
Figure 2 – The certification process. An automated test tool usually verifies compliance .....	8
Figure 3 – Features provided in the OCF Internet of Things .....	8

Figure 4 – OCF architecture featuring three pillar alignment (specification, open-source reference implementation and certification of devices) .....	9
Figure 5 – A read-only device (R) cannot be controlled unless an authenticated user obtains authorization .....	10
Figure 6 – The process for securely connecting an OCF device to the secure OCF network (which is implemented over a standard Internet Protocol (IP) network) .....	12
Figure 7 – The definition of a dimmable light bulb using component resources .....	13
Figure 8 – An asymmetric OCF bridge between an OCF client and non-OCF server devices on a non-OCF network.....	14
Figure 9 – A symmetric OCF bridge where client and server devices may be in either network (i.e. OCF clients can control non-OCF servers and non-OCF clients can control OCF servers).....	14
Figure 10 – A RESTful system in which any client can make requests to any server at any time. Notification is also supported, so servers can report events to be handled by clients .....	15
Figure 11 – The architecture and operation of the OCF cloud .....	15
Figure 12 – OCF developer event in China .....	16
Figure 13 – The Raspberry Pi board used in the tutorial (with Pimoroni Explorer Hat Pro daughter board).....	17
Figure 14 – Example input file for a simple binary switch.....	18
Figure 15 – Partial input file for the tutorial .....	19

# Introduction

The Internet of Things (IoT) is driving new products and technologies—as well as customer adoption rates. This increased integration presents an exciting opportunity for cable operators to reinforce their value, and support customers as they transition to an increasingly interconnected home. With this opportunity comes a handful of challenges around network security, device interoperability and increased customer support requirements for installation and troubleshooting. The *Open Connectivity Foundation* (a consortium of CableLabs, Comcast, Shaw, Cox, Midco, Mediacom and over 400 other companies like LG, Honeywell and Cisco) has been formed to overcome these challenges through a standardized approach to IoT.

The benefits of cross-organizational collaboration include industry alignment, cost-savings and an opportunity to develop mutually-beneficial products and platforms. For the cable industry, this partnership promises a seat at the table with leading device manufacturers, as well as economies of scale for troubleshooting and alignment on security, interoperability and scalable remote manageability. The partnership's work to develop an ecosystem which considers these impacts and allows for seamless integration of new devices will ultimately benefit customers and help usher in a new era of interconnectivity.

## Executive Summary

In this paper, we explore some of the major challenges faced by cable operators implementing an Internet of Things (IoT) offering for their customers. The IoT exponentially increases in number of connected devices on MSO networks. At the same time, these connected devices act more directly on the real world – controlling expensive equipment, unlocking doors and otherwise connecting the real world to the online world. This new type of network imposes an increased operational load on the operators while introducing new risks for both customers and the operator's network.

A cable-friendly IoT will make this workload manageable and reduce the risk. OCF is explored as a potentially important system that addresses these issues directly because cable operator requirements were part of the original design of the system. These requirements have been successfully implemented by OCF. OCF has also created a very complete development system that allows devices to be created automatically and very quickly. An IoT device based on OCF will be built live on stage during the presentation.

### 1. Context

When it comes to the Internet of Things, it is such a vast and varied topic that any particular discussion requires a bit of context. This paper will consider the IoT from a perspective of individual devices (lights, thermostats, etc.) and the applications that control them. The main topics to be discussed are the technical aspects of making these devices secure, getting them to communicate regardless of manufacturer or networking technology, and making things simple for both developers building the devices and the ordinary lay customers using the devices. In this context, operators are providing a valuable service to their customers in a way that can be profitable. This paper does not consider derivative services that can be based on the data associated with an IoT offering.

The following diagram provides an overview of the physical architecture of a comprehensive IoT network. The IoT service rides on top of existing operator data services, but should be considered as a separate platform for an unlimited number of specific IoT products to different customer segments.

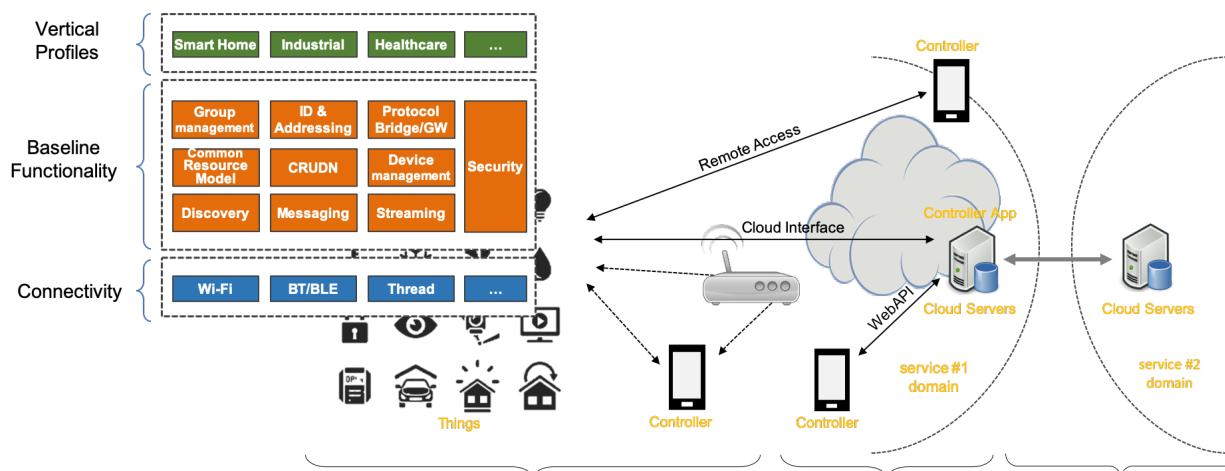


Figure 1 – A connectivity view of the Internet of Things

## 2. Operator Concerns

Operators find themselves in a unique position concerning the IoT. Cable operators are ideally qualified to be providers of the Internet of Things. They have fast networks. They already manage subscription services for customers. They have a strong 24/7 connected support system as well as fleets of technicians who can resolve problems both on premise in person. Additionally, they have relationships with equipment vendors and are accustomed to installing complex devices in customer homes and places of business.

The Internet of Things, however, comes with a host of challenges. The potential number of connected devices in a local network could be in the hundreds. The customer wants to buy the products that meets their needs regardless of manufacturer or what the operator wants to support. Each of these devices poses management challenges, enables a new attack surface for network criminals and could even threaten the integrity of the entire operator network. As with traditional cable video business, there is a great opportunity here for cable operators to be aggregators who simplify the customer experience and offer reliable, in-demand services at a great value. In order to do this effectively, the operator needs to provide the best network security available. At the same time, there must be a rich source of products that operators can provide – and support – to give customers the technology they want in a safe and easily-manageable environment.

### 2.1. Security

Security cannot be an afterthought. Security must be integrated into IoT systems from the planning phase to ensure that there are no obvious security gaps. Systems should be designed to incorporate the newest security principles, anticipate potential security breaches and have a plan for addressing them.

- **Devices** – Devices are on the front lines of potential attack so they must implement security best practices. This includes best-in-class security algorithms, the ability to update firmware on the device, and the ability to quarantine devices that can't be controlled.
- **Network** – The network can be defined as the collection of devices connected to it and the actual connection between those devices. Securing individual devices is a start, but the paths beyond the local network are also at risk. Malicious devices can be blocked to prevent them from creating havoc on the broader network. Tools that enable careful monitoring of network traffic can indicate sources of attack and isolate them before they cause too much trouble.

- **Privacy** – Privacy is related to general security but is concerned with protecting the identity of network customers. Most customers have little understanding of, and no feasible defenses to protect their privacy, so it is incumbent on network providers to do this for them.

## 2.2. Interoperability

Interoperability gives customers the freedom to install devices with features and services they want while allowing them to choose from a large pool of products from different vendors. In the current IoT environment, there are numerous smart phone applications that only support devices from the same company. This is untenable. Some vendors have seen this problem and have used the opportunity to provide interfaces to these products individually to enable a single controller. This is an expensive interoperability solution. A more efficient solution would provide a common framework compatible with many devices.

**Devices** – To provide scalable interoperability, devices need an open interface that complies with a standard in order to drive scalable interoperability. The benefits of standards are evident all around us. Web sites, mobile phones and electrical sockets are all good examples of the potential of standards. If there is only one standard, everything works together. Where there are several (as with electrical sockets worldwide), you need adapters. The fewer adapters you need, the better.

**Controller Applications** – Controller applications are the flip-side of devices. If devices are based on different standards, the controller must implement each of them, along with translators, in order to enable any level of interoperability. This problem grows exponentially with the increase of devices on the customer network.

**Groups, Scenes and Rules** – Groups allow devices to be controlled in tandem. For example, a group of lights can be turned on together or dimmed at the same rate. Scenes allow a favorite group of settings to be remembered and reproduced later. Rules monitor various conditions (time, weather, door sensors, etc.) to be monitored. When the right conditions exist, further actions can be taken (for example, sounding an alarm or turning on a light). Interoperability allows different devices to work as groups to implement scenes or launch other actions.

## 2.3. Simple Development and Support

Most operators have little interest in producing IoT products themselves. Instead, they are interested in developing relationships with a few key partners or enabling a few key ecosystems. Customers, however, do not want to be limited on which devices will work with their IoT network. The market will provide these devices, but it becomes prohibitively expensive to support multiple ecosystems on a device or write new code from scratch.

**Vendor Development** – The Internet of Things relies on connecting common devices to the Internet. These products become harder to sell if they are much more expensive than their non-connected equivalents. So the IoT part needs to be inexpensive. This can't be accomplished if you have to include software to support multiple ecosystems. Duplicating device data models, or supporting multiple ecosystems quickly becomes expensive. This is especially true for devices that have minimal memory or processing power.

**Operator Development** – Operators need to enable good customer experiences. The sheer number of IoT devices makes it impractical to deploy trucks and technicians to install every new device. This problem can be mitigated by making self-install extremely simple. If the customer does need help, customer service people need software tools to get an accurate view of the network and the ability to fix problem.

As with the IoT products themselves, the operator can scale support only if the interface to support devices is simple and consistent.

**Freedom to Select Devices** – One of the advantages of supporting an interoperable environment is that operators can choose which devices to support. Customers will ultimately decide which devices they want. If the operator already supports those devices, the customer is more likely to choose the operator-provided interface.

**Universal Backend Architecture** – In order to really make the IoT work operationally, devices, data and management need to rely on some common tools. It's hard to make a case for many IoT services independently. Devices run on different systems, connected through different networks. Still, they need common management, a uniform approach to data and the ability to control their security.

**Network Independence** – There are many different types of networks - Ethernet, Wi-Fi, Bluetooth, ZigBee and several mobile networks – and each has its individual strengths and weaknesses. Those strengths and weaknesses serve useful purposes, so it doesn't make sense to try to move everything to a common network. Instead, the control, management and data features that are independent of the particular network should be used to provide a common infrastructure.

Many IoT ecosystems are less about basic connectivity and more about the data produced by the individual devices. The principles behind data collection, however, are common. What is the format of the data? How often should it be sampled, or is it spontaneously generated? Where should it be stored? Should it be analyzed before it is collected or should it be retrieved in a raw state? These questions are independent of the particular ecosystem being used or how the device is connected. This fact makes data a good candidate for normalization in an IoT system.

**Common Support Infrastructure** – Regardless of the number of systems installed to minimize customer calls, customers are going to call. In order to efficiently dispatch these calls, the customer service representative must have tools that allow him to interact with the customer's network and IoT devices. This support will shorten customer calls and improve customer satisfaction. This type of tool requires a system where devices are recognized and controlled in a common way. They must also be managed from a security perspective in a way that is comfortable for the customer to allow and simple for the customer to enable and disable.


## 2.4. Other

There are a few other aspects of an IoT offering that might be considered by an operator in order to create a more universal and stable service.

**International Standard** – It helps if an IoT product can be based on international standards. A standard will include explicit instructions on how a device should be designed, constructed and how it should communicate. A standard has a greater chance of adoption as it is more likely to be supported by other companies in the ecosystem. Also, standards are likely to be maintained for some period of time by the authoring bodies by which they were produced. If the standard is open, it is easier for others to get access and to possibly influence how it works. This brings in more voices and more people who can review a standard to make sure it is consistent, efficient and implementable.

**Open Source Implementation** – While a standard goes a long way to enabling interoperability, there is still a chance that it will be interpreted differently by different readers. Despite the most valiant efforts of implementers, different interpretations are likely to lead to incompatible devices. This possibility can be

reduced by an open source implementation. Creating a valid working example helps to clarify the specification in the standard and is likely to be used by others even if the specification is not clear enough.

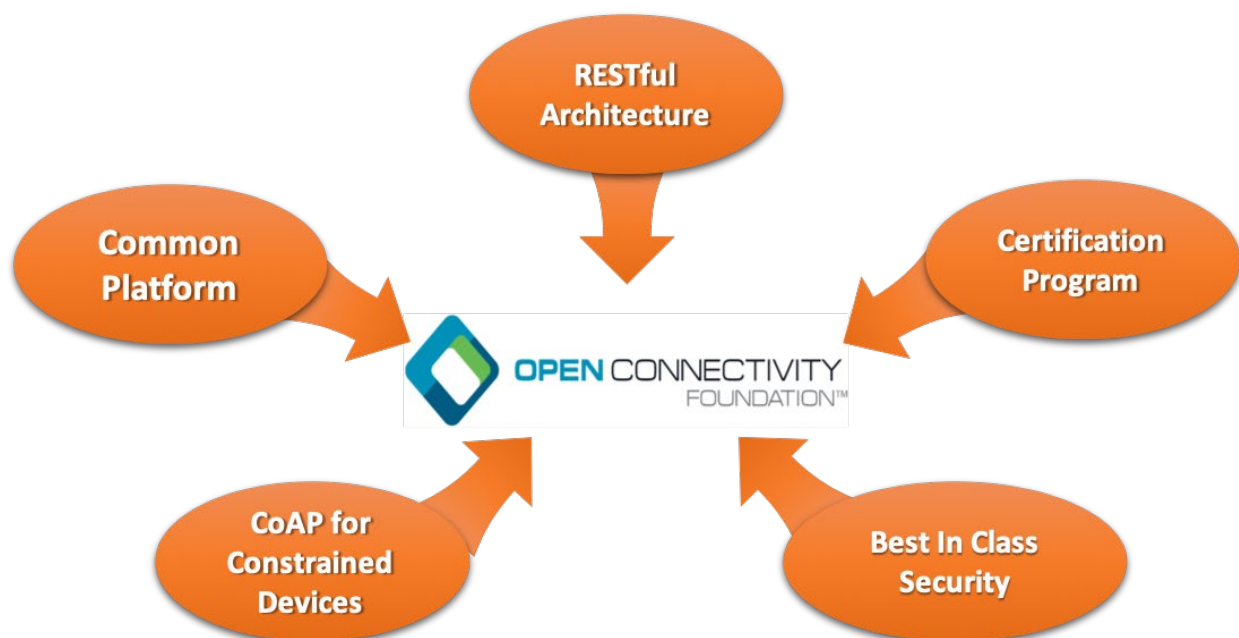


**Figure 2 – The certification process. An automated test tool usually verifies compliance**

**Certification** – The final check on an interoperable implementation is a certification process. A certification tool is really a reciprocal implementation with clear tests for each of the specified requirements. Of course the standard must be written with sufficient requirements to ensure interoperability and those requirements must be testable. Since some requirements are not testable by their very nature (negative requirements are sometimes an example of this), a certification tool can be augmented with attestation statements that an implementation is compliant with untestable requirements.

### 3. Open Connectivity Foundation Features

All the objectives stated above are not met by any single system at this point. Indeed the variability of these objectives between different operators indicates that there will always be a different combination of solutions for each operator. However, many of the objectives can be met by the Open Connectivity Foundation (OCF).



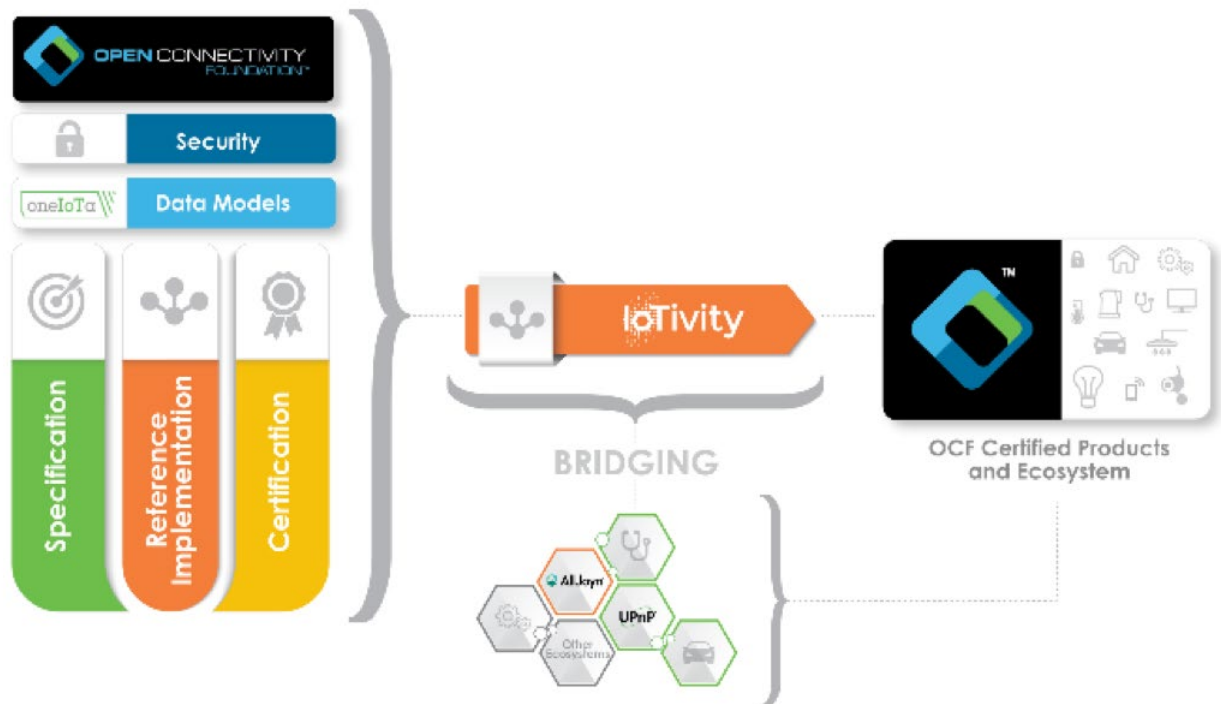
**Figure 3 – Features provided in the OCF Internet of Things**



### 3.1. Organization and Authorization

The Open Connectivity Foundation is an international standards development organization that writes standards for the IoT.

In OCF, there is something called three-pillar alignment. The standard, an open-source implementation (called IoTivity) and an automated tool are synchronized about every six months.



**Figure 4 – OCF architecture featuring three pillar alignment (specification, open-source reference implementation and certification of devices)**

The standard is frozen at a particular version number, then all contributors are notified in case they want to claim any intellectual property. This frozen version is the reference for the open-source implementation and the basis for the certification test tool (CTT).

The open-source implementation must implement all requirements specified in the standard at each release version. The open source implementation may also include optional features or features that are not even specified, but at a minimum, it must implement all mandatory features. Compliance is validated with the automated Certification Test Tool.

The automated Certification Test Tool is a software program that implements all the test cases noted by SHALL statements in the standard. Each test case will instruct the device under test (DUT) to set up the particular conditions of the test case, then verify that the DUT responds as specified. In three pillar alignment, the open source implementation must pass all the tests for a particular specification version number.

The success of any standard is determined by how broadly it is adopted. One way to encourage adoption is through membership. Those companies willing to join the organization and put in the work to develop the standard, the open source implementation or the test tool are quite likely to also implement the standard in their products.

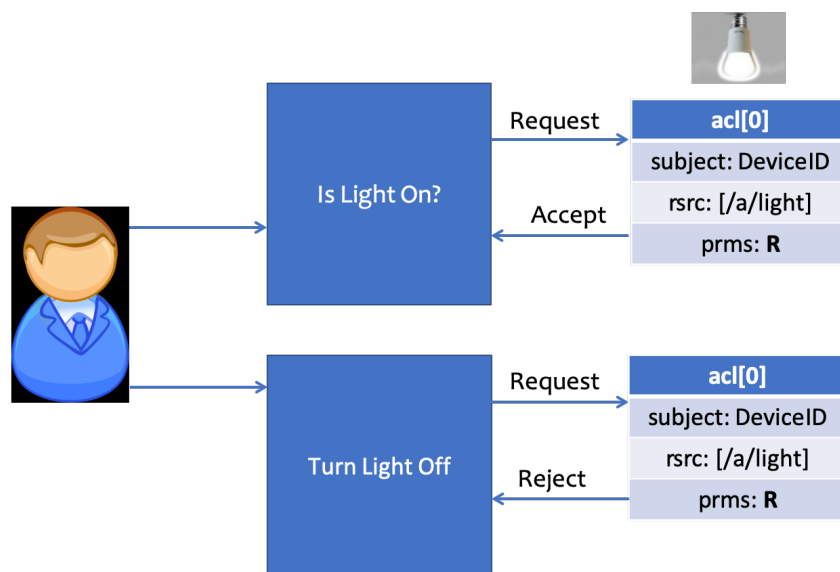
In addition to memberships, liaisons are a valuable means to encourage adoption or at least cooperation. Liaisons are particularly valuable for OCF. Since the primary objective of OCF is to create a secure interoperable system for IoT regardless of the underlying protocols, the success of OCF is partly dependent on bridging to other ecosystems. This is done on the physical level through hardware bridges with interoperable data translations, but the liaison relationship aligns the cooperation on a business level.

## 3.2. Security

Security is arguably the most important aspect of the IoT. Since IoT devices manipulate real-world objects under networked electronic control, there is a serious risk of doing something like opening a door lock without proper authorization. State-of-the-art security is the most practical solution.

### 3.2.1. Authentication

Authentication is the process of verifying that a network, a device, or a user is who or what they claim to be. This is done through digital signatures and an infrastructure that can be trusted. OCF has set up multiple registration authorities to generate and hold the root of trust for OCF certified devices. Certificate authorities then are then used to allow devices to follow the path to the root of trust to verify the authenticity of OCF devices.



**Figure 5 – A read-only device (R) cannot be controlled unless an authenticated user obtains authorization**

### 3.2.2. Authorization

Once a device is authenticated and its ownership status is verified, it can be trusted to be onboarded (joined to the secure OCF network) and controlled by authorized users.

- **Individuals** – Authorization in OCF is managed by the “owner” of the OCF secure network (a subset of the the user’s network). In a normal customer network, the owner is assumed to be the first user to onboard a device. The owner of the OCF network is authorized to authorize other users. The permissions on a particular device can be given to a particular authorized user.
- **Roles** – Roles allow for simpler management of authorization by assigning any user who has been assigned a particular role to share the authorizations of that role. Common roles include administrator, user and guest. Other roles can be created as needed.

### 3.2.3. *Security Profiles*

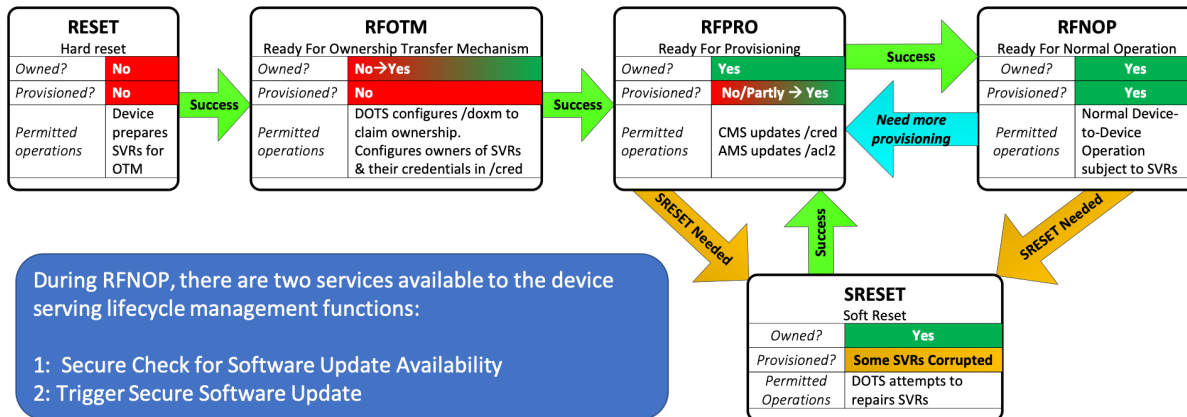
Security profiles define a minimum level of security for all OCF devices. OCF supports both symmetric and asymmetric security methods, but every certified OCF will have a minimum level of security.

- **Just Works** – Just Works security is the simplest form of security allowed in OCF. It is based on Diffie-Hellman symmetric keys. The public keys are distributed in advance and the secure connection is set up pairwise between a client and a server. This fact limits the scalability of Just Works security as it requires  $N \times N$  security relationships to be set up.
- **Random PIN** – Random PIN security also uses Diffie-Hellman, but relies on a random PIN generated by the server during onboarding. This strategy makes it harder for someone without physical access to the server screen to try to steal the connection during setup.
- **Certificates** – The other three security profiles are based on certificates issued by a Certificate Authority (CA). These certificates, in turn, are based on a root-of-trust issued by the Registration Authority under the management of a Management Authority and using policies provided by the Policy Authority. The OCF board of directors is the Policy Authority. The other authorities are contracted out. There are currently three certificate-based profiles. There is not a security hierarchy. Rather each profile has different features that can be useful in different situations
  - **Black** – The black profile uses public key infrastructure (PKI) defined by OCF and signed by designated OCF authorized certificate authorities. This guarantees that security meets the requirements defined in the OCF certificate policy.
  - **Blue** – With the blue profile, manufacturers are allowed to specify their own certificate authorities. They certify that they conform to the OCF-defined security criteria.
  - **Purple** – This profile supports a requirement to a piecewise boot process and secure online software update. This improves device integrity.

### 3.2.4. *Onboarding*

Onboarding is the process of connecting a device to the secure OCF network. The secure OCF network is really just a collection of devices that will only communicate with other similarly secure devices. In OCF, this secure “network” sits on top of the IP network. So a device must first be connected to an IP network (via Ethernet, WiFi, bridging, etc.) before the OCF onboarding process.

The onboarding process, then, is primarily using one of the security profiles described above to bring an authorized device on to the network and connecting it with authorized client devices.



Device can transition to **RESET** from any state (these transitions are not shown)

**Figure 6 – The process for securely connecting an OCF device to the secure OCF network (which is implemented over a standard Internet Protocol (IP) network)**

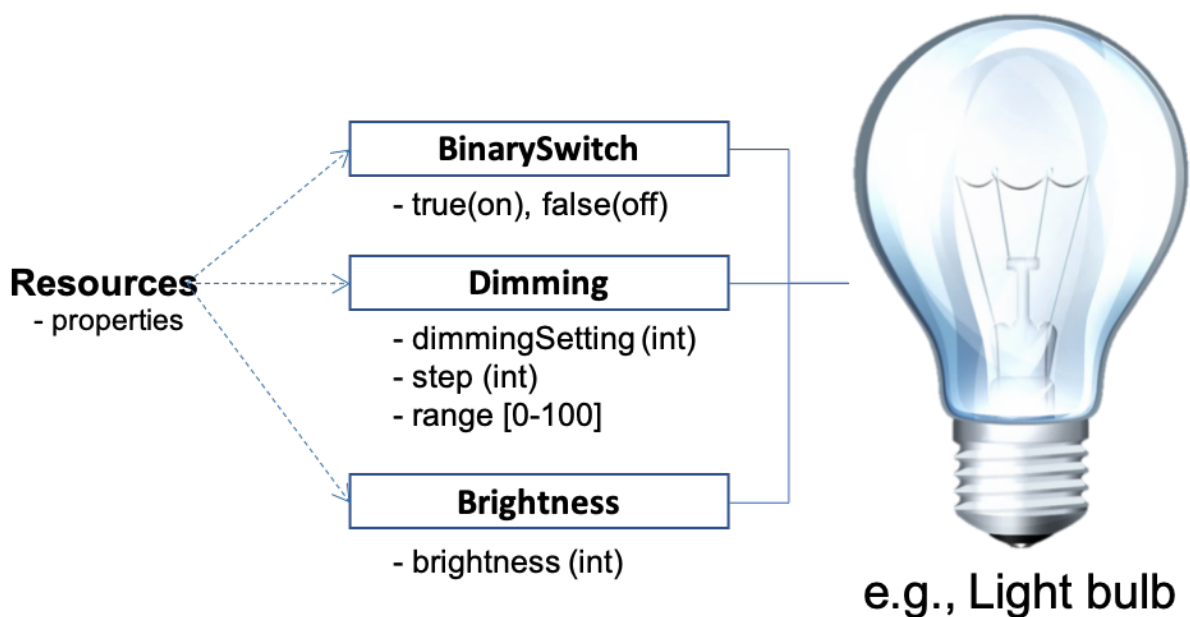
### 3.2.5. Secure Sharing

## 3.3. Interoperability

The second objective for OCF is interoperability. There are a number of features of the OCF architecture that make interoperability simple to implement. These features include a common data model, a RESTful architecture, cloud support and bridging.

### 3.3.1. Data Models

Data models are at the center of OCF. Data models describe resources that are composed into complete devices. These data models are used to generate some specifications, fully functional source code and test scripts.



**Figure 7 – The definition of a dimmable light bulb using component resources**

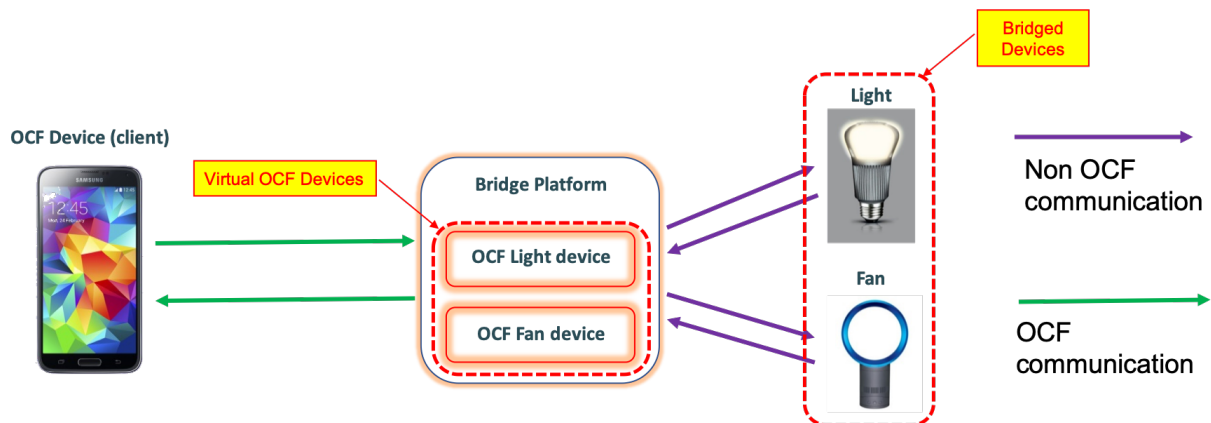
**Core Models** – The core data models are part of the core architecture in OCF. Since these models have aspects tied to the OCF architecture and particular protocols they are kept separate from the resource and device models.

**Resource Models** – Resource models are models that contain the minimum properties and definitions to describe a complete component. For example, a temperature. Interoperability in OCF is defined at the resource level. This allows for interoperability between any devices that use a common resource.

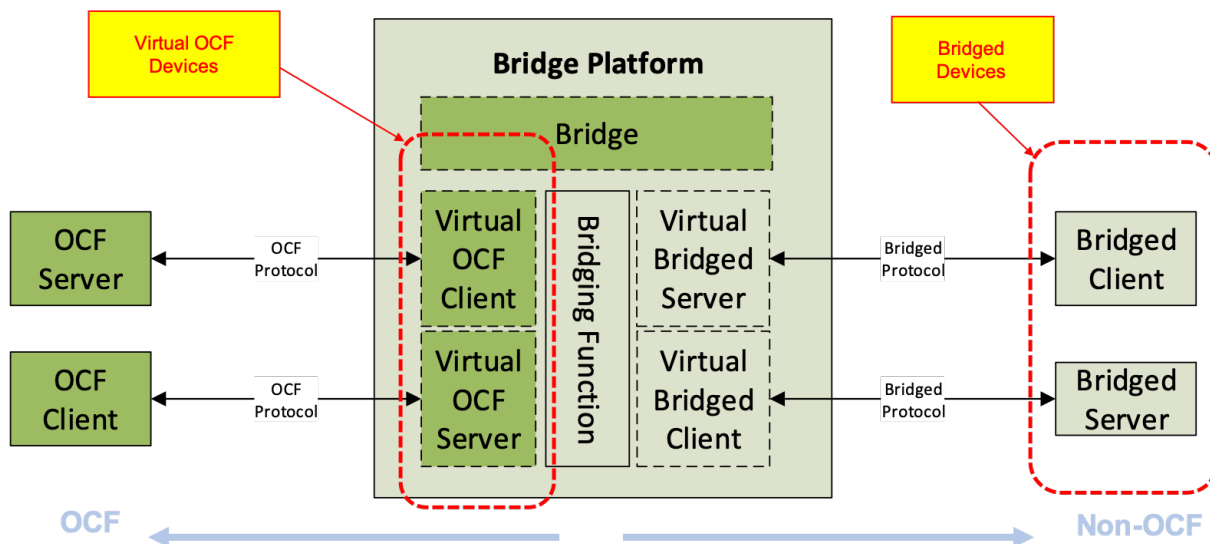
**Device Models** – Devices are fully functional items that a consumer might purchase. Devices are composed of a collection of resources. Device models are defined with the minimum number of required resources to be classified as a certain type of device (e.g. a thermostat). By defining devices in this way, a manufacturer is able to arbitrarily enhance their device to distinguish it from competitor devices while still be compatible with other such devices.

Resource models are managed within the oneIoTa online tool. This tool includes a simplified Integrated Development Environment (IDE), and process management to enable the review and approval of new resource data models. Once models are approved, they are pushed to a git repository where they can be obtained to build OCF devices.

**Other Organization Models** – The oneIoTa tool supports models contributed by other organizations. Any organization can set up an independent organization within oneIoTa to use both the IDE and the process management features. This allows other organizations to independently control their own models while making them accessible to OCF.



**Figure 8 – An asymmetric OCF bridge between an OCF client and non-OCF server devices on a non-OCF network**




**Figure 9 – A symmetric OCF bridge where client and server devices may be in either network (i.e. OCF clients can control non-OCF servers and non-OCF clients can control OCF servers)**

**Derived Models** – Derived models describe the mapping of data between native OCF models and the models from any other organization. The mapping describes mapping into and out of OCF. The mappings do not need to be one-to-one. In fact, any procedures that can be done in most programming language can be used to map between the models. This flexibility makes it possible to build bridges between OCF and most any other ecosystem.

### 3.3.2. Restful Architecture

OCF supports a RESTful architecture. Devices are normally controlled using Create, Retrieve, Update, Delete, and Notify with data models. However, OCF can also be used to remotely execute commands. This is useful for supporting streaming services and other services that are not well-modeled using a RESTful approach.

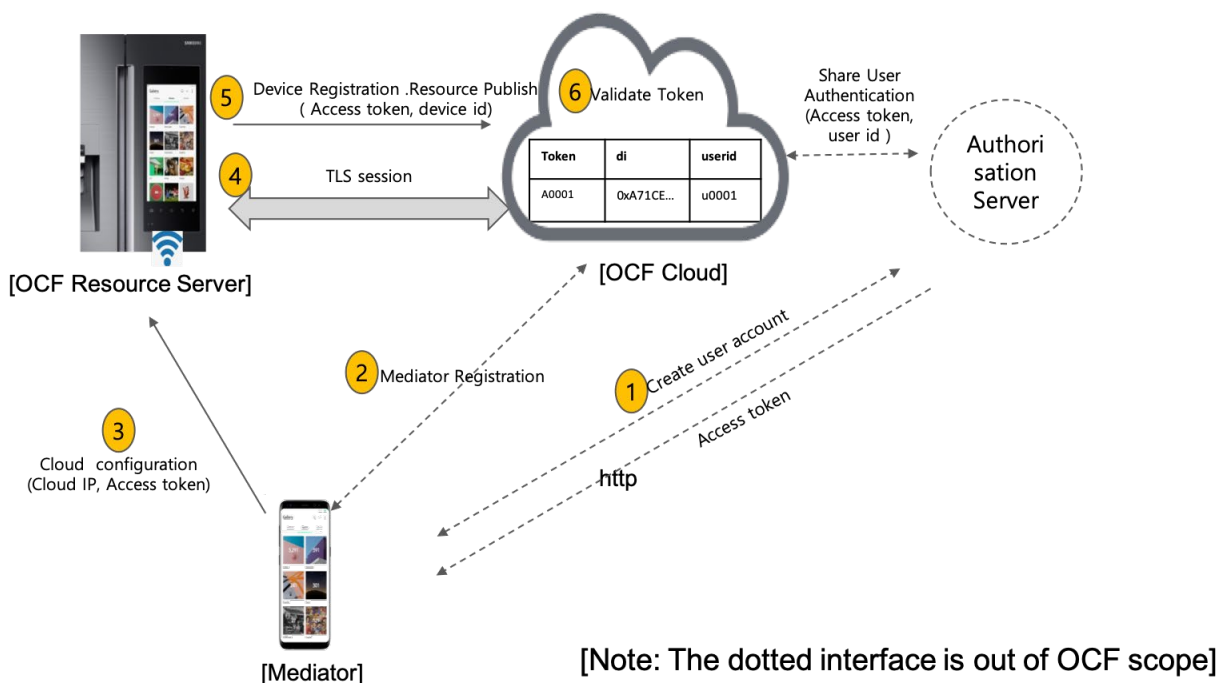


**Figure 10 – A RESTful system in which any client can make requests to any server at any time. Notification is also supported, so servers can report events to be handled by clients**

### 3.3.3. Remote Access and Cloud Support

The IoT is not truly an Internet experience if it is only on a Local Area Network (LAN). For this reason, OCF supports remote access to LAN IoT devices (e.g. using a smart phone) as well as cloud services and cloud-to-cloud connections between vendors.

**Device-to-Cloud** – Device-to-cloud supports connecting servers or clients to devices in the cloud. This means not only that devices can be controlled from the network but that virtual devices can be implemented in the cloud.



**Figure 11 – The architecture and operation of the OCF cloud**



**Cloud-to-Cloud** – Cloud-to-cloud allows clouds from two different companies or ecosystems to make connections between devices and clients anywhere. This means companies can define their own cloud infrastructures and still use OCF and bilateral deals to control OCF devices regardless of particular network connectivity.

#### **3.3.4. Bridging to Other Ecosystems**

Bridges in OCF allow connections between OCF and other ecosystems. Bridges connect between protocols at layers beneath OCF, then use OCF derived models to define the mapping between OCF and other ecosystems. This approach has the added benefit of providing compatibility with devices that have already been deployed as long as an OCF bridge is added to the system.

#### **3.3.5. Common Management**

OCF enables connectivity between devices from different ecosystems through OCF bridges or cloud connections. This strategy enables OCF management tools to extend to these other ecosystems. Tools that allow for devices to be onboarded, controlled and updated can all be used with only an OCF interface. This means there are lower costs for deployment and ongoing management of mixed IoT ecosystems.




**Figure 12 – OCF developer event in China**

#### **3.3.6. Interop Events**

Interop events hosted by OCF give vendors the opportunity to test their servers and clients in a safe setting with servers and clients from other vendors. This demonstrates the OCF promise of security and interoperability. It also gives manufacturers confidence that their devices will be interoperable when deployed to customers.





**Figure 13 – The Raspberry Pi board used in the tutorial (with Pimoroni Explorer Hat Pro daughter board)**

### **3.4. Simple Development**

The final objective of OCF is to make it easy for vendors to create, deploy and manage devices. OCF has several tools that make this easy.

#### **3.4.1. Simple Device Description**

A device is simple to define in OCF using JSON schema. Each device is described as a collection of the individual resources it implements. A thermostat, for example, would have a setting for the desired temperature, a thermometer for measuring the current temperature, and a switch that can be used to turn on the heating or air conditioning. In OCF, these descriptions are written in JSON (a data modeling language). An example is shown below:

```
[
  {
    "path" : "/binaryswitch",
    "rt" : [ "oic.r.switch.binary" ],
    "if" : [ "oic.if.a", "oic.if.baseline" ],
    "remove_properties" : [ "range", "step", "id", "precision" ]
  },
  {
    "path" : "/oic/p",
    "rt" : [ "oic.wk.p" ],
    "if" : [ "oic.if.baseline", "oic.if.r" ],
    "remove_properties" : [ "n", "range", "value", "step", "precision", "vid" ]
  }
]
```

**Figure 14 – Example input file for a simple binary switch**

### **3.4.2. Code Generation**

OCF has a tool called DeviceBuilder that will generate working OCF code using the device description file described above, a common code template (C and C++ developed so far), and the resource descriptions stored in oneIoTa. The generated code will compile without changes to create a working OCF IoT device. The code includes stub routines that can be populated with interface code to the hardware of a specific real-world device. Usually this is only a few lines of code that is linked to a library.

### **3.4.3. Building the Application**

Building the application is simply a matter of compiling and linking the code that was generated by DeviceBuilder. A makefile is used to describe all the dependencies and include the appropriate hardware support libraries.

### **3.4.4. Device Ownership**

The lowest level device security model supported is called “just-works” and is based on Diffie-Helman symmetric security. A variation can also be implemented with a random PIN being generated by the server device and entered in the client device. Asymmetric certificate-based security is also supported in OCF.

### **3.4.5. Introspection and Automatic User Interface Generation**

In addition to code generation, the DeviceBuilder script creates an “introspection” file that is used to create the user interface for the newly created device. The client applications (OTGC) reads this file and creates a generic user interfaces (generally buttons and text) that can be used to control any OCF server device. OTGC is available for Android, iOS, Windows and Linux and the source code is available, so vendors can modify it to add their own widgets for a nicer looking interface.

### 3.4.6. Running, Testing and Debugging

The device can be run by resetting the security of the device to Ready For Onboarding Transfer Method (RFOTM), the using the run script. Besides using OTGC, the developer can also run Device Spy to have explicit control of the exact message payloads that are sent.

## 4. OCF Development Process

Here is an example of the entire build process:

1. Start with the following device description file:

```
{
  {
    "path" : "/touch1",
    "rt" : [ "oic.r.sensor.touch"],
    "if" : ["oic.if.baseline", "oic.if.a"],
    "remove_properties": [ "range", "step", "id", "precision"],
    "remove_methods": ["post"]
  },
  .
  .
  .
  {
    "path" : "/output4",
    "rt" : [ "oic.r.switch.binary"],
    "if" : ["oic.if.baseline", "oic.if.a"],
    "remove_properties": [ "range", "step", "id", "precision"]
  },
  {
    "path" : "/oic/p",
    "rt" : [ "oic.wk.p"],
    "if" : ["oic.if.baseline", "oic.if.r"],
    "remove_properties": [ "n", "range", "value", "step", "precision", "vid" ]
  }
}
```

**Figure 15 – Partial input file for the tutorial**

2. Now generate the source code from the device description file by using the DeviceBuilder script:

```
gen.sh
```

3. Now that the source code has been generated, compile and link it:

```
build.sh
```

4. The server is now created. Next, set the security mode to RFOTM:

```
reset.sh
```

5. Finally, run the server code:

run.sh

6. With the server running, run the OTGC (using Android) client.
7. Press the “discovery” button to find the server.
8. Now select the discovered server and press the onboarding (+) button.
9. With the server onboarded, launch the automatically-generated user interface.
10. Test the user interface by moving the binary switch. The effect will be reflected in the server window.

#### **4.1. Controlling the Server with DeviceSpy**

DeviceSpy is a lower-level client that gives developers explicit control of the OCF payload. It implements the same functionality as OTGC, but is available only as an executable application on Windows.

#### **4.2. Testing with the Compliance Test Tool**

The final step of the development process is to test the device for compliance to the OCF specifications. This is done with a tool called Compliance Test Tool (CTT). This automated tool will test almost every aspect of an OCF implementation. There are a handful of requirements that cannot be tested by CTT. For certification, an implementation must pass the CTT and also attest to compliance with any of the OCF requirements are not testable.

## **Conclusion**

As mentioned at the beginning of this paper, the cable industry is well-positioned to become the premium IoT provider. While the industry has several natural advantages in the IoT space, there are serious security challenges it will need to address. Additionally, it has been difficult to find any single applications or services that justify the significant investment in infrastructure required to support an IoT service. That means the investment must support several different business cases on the same platform. Also, the Internet of Things exponentially increases the number and complexity of Internet-connected devices on the network. These devices must be easy to build, manage and support.

Open Connectivity Foundation makes up a piece of this puzzle. While its scope does not cover all cable industry IoT use cases, it does provide best-in-class security, a comprehensive infrastructure for interoperability and tools to support quick development.

Cable operators still have many problems to solve in this area. There are many industry sectors that are not adequately addressed by OCF. Even the sectors that are well-covered are just beginning to be deployed. The ability to be managed is not the same as having a management infrastructure. Still, it is a start. OCF can be used to as a first example of an open, secure, interoperable platform that can be used in key cable IoT markets.

## Abbreviations

OCF	Open Connectivity Foundation – an international standards development organization for the Internet of Things
IoT	Internet of Things – the collection of objects that interact with the real world and are enhanced by connection to the Internet

## Bibliography & References

openconnectivity.org, Open Connectivity Foundation, <https://openconnectivity.org>, 2019.

iotivity.org, IoTivity: Getting Started, <https://iotivity.org/getting-started>, 2019.

iso.org, ISO/IEC 30118-1:2018, INFORMATION TECHNOLOGY — OPEN CONNECTIVITY FOUNDATION (OCF) SPECIFICATION, <https://www.iso.org/standard/53238.html>, 2018.

# **Layer 1 Considerations for Extended Spectrum Utilization in Hybrid Fiber Coax & Distributed Access Architecture Networks**

A Technical Paper prepared for SCTE•ISBE by

**Ron Wolfe**

Sr. Director, Engineering  
Charter Communications, Inc.  
14810 Grasslands Drive  
Englewood, CO 80112  
+1 (720) 518-2307  
Ron.Wolfe@Charter.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Scope of the Extended Spectrum Challenge .....	3
Coaxial Cable .....	4
Analog Optical Transport .....	5
Digital Optical Transport .....	5
RF Amplifiers .....	6
1. Total Composite Power Requirements .....	6
2. Forward and Return Transition .....	9
3. Echo Cancellation .....	9
4. Power Consumption and Heat Dissipation .....	10
Taps and Passives .....	10
Customer Premise Equipment .....	10
1. Modems as Gateways .....	11
Conclusion and Evolutionary Outlook .....	11
Abbreviations .....	11
Bibliography & References .....	12

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Continuous Slope Extension 258MHz to 1,794 MHz .....	6
Figure 2 - Total Composite Power with Increasing Channel Load .....	7
Figure 3 - "Zig-Zag" Sloped Spectrum .....	8
Figure 4 - Total Composite Power with Zig-Zag Sloped Spectrum .....	8
Figure 5 - Guard Band Width vs Center Frequency .....	9

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Loss Table for Hardline Coaxial Cables (dB/100ft) .....	4
Table 2 - Loss Table for Subscriber Drop Coaxial Cables (dB/100ft) .....	4

# Introduction

Traffic needs in service provider networks continues to grow rapidly. Although the cumulative average growth rate for per subscriber data consumption has somewhat abated, the competitive pressure has not, and there is a growing consensus among service providers that capacity planning for the future needs to begin now.

There are numerous drivers to current and future growth and among them are the rapid growth of IP delivered video services, backhaul of wireless services and the emergence of cloud connected devices such as surveillance cameras, machine to machine communication, health and wellness monitors and inter-autonomous vehicle communications. The exponential growth of these internet of things (IoT) types of devices, though each is a relatively small consumer of resources, will begin to emerge from the “noise floor” of network traffic consumers. As service providers look at their networks today, there is little argument that further increases in capacity are necessary in the near future.

Fortunately, there are pathways to additional capacity. Fiber to the premise solutions remain a viable and highly effective option, particularly so in green field environments, but our industry’s strategic asset has always been its existing coaxial infrastructure, and that remains our most readily accessible and cost-effective path to higher capacity. Our hybrid fiber coaxial (HFC) networks have grown in capacity over the years through an ongoing succession of bandwidth expansion, digital compression, serving group splits and multicast solutions. The roadmap to capacity expansion for HFC networks beyond 1,218 MHz is a multi-faceted challenge, but there are numerous companies in both the supplier and service provider communities who are collaborating to identify the challenges and develop solutions to address those challenges. This paper will explore the challenges without necessarily offering solutions for each. It is likely that SCTE-ISBE’s EXPO 2020 will have numerous presentations proposing solutions that will lead to the realization of the Data Over Cable Service Interface Specifications (DOCSIS®) 4.0 technology and applications that take advantage of spectrum beyond the proposed DOCSIS specification extension to 1,794 MHz.

## Scope of the Extended Spectrum Challenge

This paper is limited in scope to Layer 1 considerations, as stated in the title. Layer 1 of the outside plant (OSP) network consists of optical transmitters and receivers, coaxial cable, radio frequency (RF) amplifiers and passive devices such as directional couplers and taps, splitters and power inserters. For DOCSIS 4.0 technology or other extended spectrum services to be fully leveraged, all of these devices must exhibit satisfactory performance over the entire frequency range of 5 MHz to 1,794 MHz, with a future upgrade path to 3 GHz. It is premature to make exact predictions of performance and requirements for this area of the spectrum as the DOCSIS 4.0 specifications are only now beginning development. In the following sections we will explore each of these key elements of the layer 1 network.



# Coaxial Cable

Coaxial cables have been in production for many years, and the materials and technologies used to manufacture coaxial cables has evolved over that time. Coaxial cables were not specified or measured to 1 GHz until the early 1990's, and cables older than this may have degraded due to environmental conditions. This doesn't necessarily mean the cable is not useable at frequencies beyond 1 GHz, regardless of age, but it does make it necessary to sample test older coaxial cables to determine the performance characteristics. It is a safe assumption that some cable spans will require replacement, and history tells us that an expectation of 10-15% of either replacement cable or newly added cable is realistic.

Attenuation in coaxial cables increases with frequency, so it will be necessary to account for this in the design process. As an approximate measure, the loss of coax increases by a factor of 1.4 each time the frequency doubles. Stated mathematically, the attenuation of coaxial cable increases proportionately to the square root of the frequency ratio. Table 1 provides loss data in dB for hardline cable and Table 2 provides loss data in dB for drop cable.

As can be seen from the rightmost columns, there is an increase in attenuation between 1,002 MHz and 1,794 MHz of roughly 40%. This will mean that long continuous runs of untapped coaxial cables such as trunk runs or "express" feeder runs will require careful attention as this means the gain requirement for the span goes up by 40% as well. Tapped runs tend to be less problematic as the increases in through loss for taps is significantly less than 40% over the same frequency span.

It is important to note that while coaxial cable performance is predictable in the majority of cases there are some spectral areas where environmental and process variations could cause performance to differ from that predicted. It is also possible that there will be areas of the spectrum that will not be able to support the same modulation profile as those areas that are more true to predicted behavior. Sample testing is the most reliable means to determine whether this is the case.

**Table 1 - Loss Table for Hardline Coaxial Cables (dB/100ft)**

	5MHz	54MHz	204MHz	750MHz	860MHz	1,002MHz	1,218MHz	1,794MHz	3,000MHz
.500	0.16	0.54	1.09	2.16	2.34	2.52	2.81	3.54	4.77
.625	0.13	0.46	0.92	1.78	1.93	2.07	2.30	2.89	3.87
.750	0.11	0.37	0.74	1.48	1.61	1.74	1.94	2.45	3.34
.875	.09	0.33	0.66	1.29	1.41	1.53	1.68	2.13	2.89

**Table 2 - Loss Table for Subscriber Drop Coaxial Cables (dB/100ft)**

	5MHz	54MHz	204MHz	750MHz	860MHz	1,002MHz	1,218MHz	1,794MHz	3,000MHz
Series 6	0.58	1.60	3.05	5.65	6.10	6.55	7.04	8.49	10.68
RG11	0.38	0.96	1.90	3.65	3.98	4.35	4.80	6.11	8.37

# Analog Optical Transport

Since the early 1990's analog optical transport has been the predominant means for signal distribution from hubs to neighborhoods served by optical nodes. Return path transport has been a combination of analog and digital transport. Analog transmitters are subject to clipping when the power of the input signal exceeds the dynamic range of the transmitter. Transmitters are typically designed to operate with a flat input spectrum. Increasing the channel load without making per channel drive level adjustments can cause clipping. The typical response to this is to lower the drive level such that the total composite power used to modulate the laser stays the same. That means that the optical modulation index (OMI) of the transmitter decreases on a per channel basis. This could have an adverse effect on the performance of the link.

In the case of digital transmitters used in the return path changes will also be required. Most extended spectrum planners anticipate a minimum extension of the return spectrum to 204 MHz, with frequencies as high as 492 MHz or even 684 MHz being discussed. Just as is true in the downstream transport, this higher power loading of the return transmitters could cause clipping without a reduction in per channel power resulting in an adverse impact on the performance of the upstream link and/or an increased risk of clipping.

Multi-wavelength analog optical links are complex designs with many dependencies and interactions in both the electrical and optical domains. While there will be design challenges, it is likely that these can be overcome, keeping analog optical transport a viable option even at these extended frequencies.

# Digital Optical Transport

DOCSIS 3.1 technology precipitated the development of the Distributed Access Architecture (DAA), wherein analog optical transport is replaced by a digital optical transport infrastructure and the RF spectrum for each serving group is generated by a device in the node. These devices fall under the descriptions of remote PHY devices (RPDs), remote MAC-PHY devices (RMDs) flexible MAC architecture (FMA) devices and the architecture utilizing them is referred to as DAA.

At Layer 1, each of these DAA devices performs essentially the same function, with the differences between them existing in higher layers of the architecture. The primary advantage of these devices is that the signal quality we would normally expect to see in the hub site is now able to be generated in the field with equal quality and without the variations inherent to a traditional headend combiner network feeding an analog optical transmitter. That quality then translates into the ability to support increasingly complex modulation schemes in the field. For purposes of this paper, the use of these devices and the digital optical transport improve our signal quality, at the node by about 6dB, and this headroom is useful in the design process for the network beyond the node.

While DAA devices are in active deployment today, many are designed to operate at frequencies up to 1,218 MHz and provide a single forward and either one or two return segments. It is anticipated that generational improvements in the digital signal processors (DSPs) and field programmable gate arrays (FPGAs) that are the core of RPDs and RMDs will track with the development of the DOCSIS 4.0 specifications. Reduced power consumption and thermal management will need to be key design considerations in this process as we approach heat dissipation requirements that exceed the physical limitations of housing sizes for strand mounting. A migration toward application specific integrated

circuits (ASICs) would help reduce power consumption, but this is not practical while the technology is evolving.

## RF Amplifiers

### 1. Total Composite Power Requirements

RF amplifiers represent one of the more significant challenges in designing an extended spectrum network. The expansion from 1,002 MHz to 1,794 MHz represents an increase in channel loading of 105%, and the RF power levels of the individual channels increase with frequency in order to match the existing design and provide for a flat spectrum at the customer premise. For example, a typical 1,002 MHz design might call for RF output levels at the node of +50 dBmV at 1,002 MHz and +38 dBmV at 54 MHz. To drop into this design and continue the linear slope to 1,794 MHz would require RF levels of +60 dBmV at 1,794 MHz and +40.5 dBmV at 258 MHz. Figure 1 illustrates this approach. This is typical of most upgrades to our HFC networks over the past decades. Its simplicity lies in providing a relatively flat spectrum at the customer premise.

By graphing the total composite power in watts as the load is increased from 1 to 253 channels from 258 MHz to 1,794 MHz we can see how the extension of the spectrum affects the TCP at the output of the amplifier. In this example the TCP required to generate this spectrum is 77.5 dBmV, a level that is unattainable across such a wide spectrum using today's Gallium nitride (GaN) gain blocks. Even if this were possible the power consumption of the gain block would be unacceptably high and current amplifier housings would be incapable of dissipating the heat they would generate. There may be many cases where current levels output from nodes and amplifiers allow a "continuous slope" design approach such as this, but it is unlikely that this will be a universally viable approach. Networks originally designed to 550 MHz are unlikely to be candidates for a continuous slope drop-in design.






Figure 1 - Continuous Slope Extension 258MHz to 1,794 MHz




**Figure 2 - Total Composite Power with Increasing Channel Load**

In most cases, engineers are considering a slope pattern that maintains current tap output levels up to 1,002 MHz then restarting a parallel slope line at 6 to 8 dB lower than the legacy plant levels. There are two advantages to this approach. First, it preserves the legacy design levels, making this an effective approach to brown field upgrades, where existing services and customer premise equipment (CPE) is designed to operate effectively. Second, the carriers at higher frequencies are the largest contributors to the TCP curve, and reducing their levels has a higher beneficial impact on the TCP than reducing lower frequencies. The resulting RF spectrum takes on a staggered look when viewed on a spectrum analyzer, as illustrated in Figure 3.

This reduction of 8dB in the carrier levels at frequencies will obviously have to be accounted for in the design of the customer premise network and some possible solutions are discussed briefly later in this document. The resulting TCP of this “Zig-Zag” sloped spectrum is 71.2 dB, which is within the reach with the newest generation of GaN amplifiers. Figure 4 is the resulting TCP graph of the carriers.




**Figure 3 - “Zig-Zag” Sloped Spectrum**



**Figure 4 - Total Composite Power with Zig-Zag Sloped Spectrum**

## 2. Forward and Return Transition

Since our early use of bidirectional amplifier stations that contain both forward and return amplifiers, we have relied on diplexers to keep signals in one direction from interfering with signals travelling in the opposite direction. Diplexers exist on the input ports of the amplifier. Although high quality components and manufacturing techniques have improved over the years, there is still a need for a “guard band” between the forward and return frequencies, as illustrated in Figure 5. The frequency at which the network is provisioned is often referred to as the “split” frequency and indicated by a combination of the highest return frequency and the lowest forward frequency. As the center frequency of the diplexer increases, more bandwidth is allocated to the return network and less to the forward, thereby enabling a higher degree of symmetry between the downstream and upstream capacity. However, the width of the guard band for the diplexer also increases, in this case from 12 MHz in a traditional diplexer to 54 MHz for a 5-204 MHz / 258-1,218 MHz, consistent with today’s DOCSIS 3.1 specifications. Traditional diplexers have a guard band that is roughly equal to 23% to 25% of the center frequency. It is probably reasonable to assume that higher quality filters can reduce this to 20% or a bit less,



**Figure 5 - Guard Band Width vs Center Frequency**

Although the specifications have not yet been drafted, there is a desire to incorporate even higher split frequencies, for example 492 MHz or 684 MHz as the highest return frequency. These splits would incorporate diplexers that would require guard bands of approximately 110 MHz and 160 MHz respectively, even after accounting for better quality filters, the amount of usable spectrum lost to the guard bands is a high cost to pay for the increase in return spectrum, and for this reason service providers are exploring alternatives.

## 3. Echo Cancellation

One approach that has received a good deal of attention in recent months is to utilize echo cancellation technology to provide isolation between the forward and return paths in place of diplexers. Echo cancellation allows the forward and return path to transition with almost no loss of spectrum between the two. Echo cancellation will be a requirement for a full duplex (FDX) DOCSIS solution, and leveraging

this for statically duplexed applications benefits both solutions through scale. While designs for echo cancellers can be complex, particularly for multi-output amplifiers, they do certainly deserve further research and engineering to determine performance and price points. Ultimately, if the cost of echo cancellation can be comparable to the value of the spectrum reclaimed it will likely be successful.

## **4. Power Consumption and Heat Dissipation**

As always, the contribution to the power consumption of the amplifier is a matter of concern. Next generation GaN amplifiers have a significantly higher TCP capability, but that comes with a higher power consumption. These devices also dissipate more heat than current generation RF amplifiers, and some compact amplifiers may be challenged to sufficiently dissipate the heat load of a high wattage RF Gain block and it's supporting circuitry, particularly amplifiers with multiple output ports.

## **Taps and Passives**

There are very few, if any taps installed in HFC networks today that are capable of satisfactory performance to 1,794 MHz, meaning there will need to be at a minimum a faceplate change required. This may prove to be a satisfactory option for some service providers, but there are considerations in the housing, sometimes referred to as the “back box” and the seizing mechanism that affect performance at frequencies higher than about 1,200 MHz. Often these are due to transmission characteristics and construction of the seizing mechanism and power bypass devices designed to maintain power when faceplates are removed.

Initial product samples of taps specified for performance to 3 GHz has been promising, with early samples showing satisfactory performance and an ability to improve on the loss performance at frequencies up to 1,002 MHz. This makes it feasible that a same value extended spectrum tap may be inserted into the same design location. The higher levels at amplifier and node outputs will likely make conditioning plug-ins a requirement, with cable simulators for higher value taps and equalizers for lower value taps.

Taps with amplification have also been discussed as a possible solution to placing full amplifier stations into spans that are prohibitively long to maintain the existing spacing. These “gain taps” could be spaced in optimal locations to minimize their impact on cascade performance.

## **Customer Premise Equipment**

While customer premise is beyond the scope of this paper, the development of CPE specifications will represent one of the most significant challenges in the extended spectrum development and implementation. There are no currently deployed set-top boxes or modems that support reception or transmission of signals above 1,218 MHz. This means that a new modem will be required to utilize this spectrum. The transmit frequency of most deployed modems ends at 85 MHz and that too will need to change to support at a minimum 204 MHz, and possibly as high as 684 MHz. With the increased loss of the coaxial cable and taps between the modem and the nearest amplifier, return transmitter power will need to be increased as well. All of this presents an opportunity to rethink how we deploy CPE in our networks.

## 1. Modems as Gateways

Today's customer premise installation typically consists of a drop feeding a 4-way splitter, one port of which serves the DOCSIS modem and the other three serving set-top boxes. If we were to transition the customer premise to an all IP network we could realize the elimination of the loss of that 4-way splitter. The drop could terminate on the modem and the set-top boxes could be replaced with IP set-top boxes connected to the network either by wireless or using the existing coax in a coaxial Ethernet network. By eliminating this loss, we can offset the loss of the drop at the higher frequencies and potentially open the door to using higher value taps that have less through loss. This is a scenario worth exploring with the condition that it must maintain compatibility with legacy customers who may be served by the same tap on a different port.

## Conclusion and Evolutionary Outlook

In the early 1990's cable service providers deployed HFC networks with an eye toward the future, and with the understanding that coaxial cable had enormous potential for capacity expansion beyond the limits of the technical limitations of electro-optical and RF components that comprise the active portion of the network. Extended spectrum operation to 1,794 MHz is a significant evolutionary step for HFC networks, but there is no reason to believe it is the last such step. Every HFC network is different in some way from others, so there is no single design approach that will work in every case. The most challenging requirements lie in the RF amplifier designs to support the additional spectral loading, the need for higher transmit power from CPE to overcome the greater coaxial losses at higher frequencies, addressing the guard band loss of capacity due to higher split frequency diplexers and management of power consumption and heat dissipation. From my experience in the industry, I am confident that these and other challenges will be overcome and that the incremental bandwidth and capacity made available through extended spectrum upgrades will position the industry to be competitive for years to come.

While it is true that most current efforts are focused on an evolutionary step that supports the needs of DOCSIS 4.0 technology, the spectrum beyond 1,794 MHz should also be considered for transport of other technologies. There are solutions for example that could utilize this upper end of the spectrum to transport Ethernet signals over the coax. This could be used to feed new DAA nodes for future segmentation, or to provide other Ethernet based solutions without the need to extend fiber. The same technology could be used to distribute IP traffic to IP set-top boxes within the customer premise. While this paper is based on operation of the network to 1,794 MHz, service providers are working with vendors to provide an infrastructure of cables, connectors and housings that support drop-in upgrades to 3 GHz.

There are significant challenges to overcome in the development of network specifications for performance to 1.8 GHz and beyond, but none of these challenges appear insurmountable. By next year's Cable-Tec EXPO it is likely that several manufacturers will be introducing products designed for operation to 1.8 GHz and even 3 GHz.

## Abbreviations

ASIC	application specific integrated circuit
CPE	customer premise equipment
DAA	distributed access architecture
dBmV	decibels relative to 1 millivolt



DOCSIS	data over cable service interface specifications
DSP	digital signal processor
FDX	full duplex
FMA	flexible MAC architecture
FPGA	field programmable gate array
GaN	gallium nitride
GHz	gigahertz (1 billion cycles per second)
HFC	hybrid fiber-coax
IP	internet protocol
ISBE	International Society of Broadband Experts
MAC	media access control
MHz	megahertz (1 million cycles per second)
OMI	optical modulation index
OSP	outside plant
PHY	physical layer
RF	radio frequency
RMD	remote MAC/PHY device
RPD	remote PHY device
SCTE	Society of Cable Telecommunications Engineers
TCP	total composite power

## Bibliography & References

Broadband Library: *Total Power and Power Spectral Density*; May 25, 2019 by Ron Hranac

With the exception of the aforementioned, specific documents were not used in the production of this document. The document is the product of numerous discussions on the topic with subject matter experts from other service providers and from industry suppliers.

The author wishes to acknowledge and thank the following companies for their advice and insights.

Amphenol Broadband Solutions

ATX Networks

Cisco Systems

CommScope/Arris

Cox Communications

GenXComm

Qorvo

Shaw Communications

Technetix

# 5G Backhaul/Fronthaul Opportunities and Challenges

A Technical Paper prepared for SCTE•ISBE by

**Joe Mocerino**  
Global Solution Architect  
Fujitsu Network Communications

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Radio Type Drives Optimum Transport Option.....	4
Optimum Transport for mmWave .....	4
TSN or Full Active WDM for Sub 6 GHz Radio Fronthaul .....	5
Virtual Networks via Network Slicing .....	6
Control in the Cloud .....	7
Scalable Slices .....	8

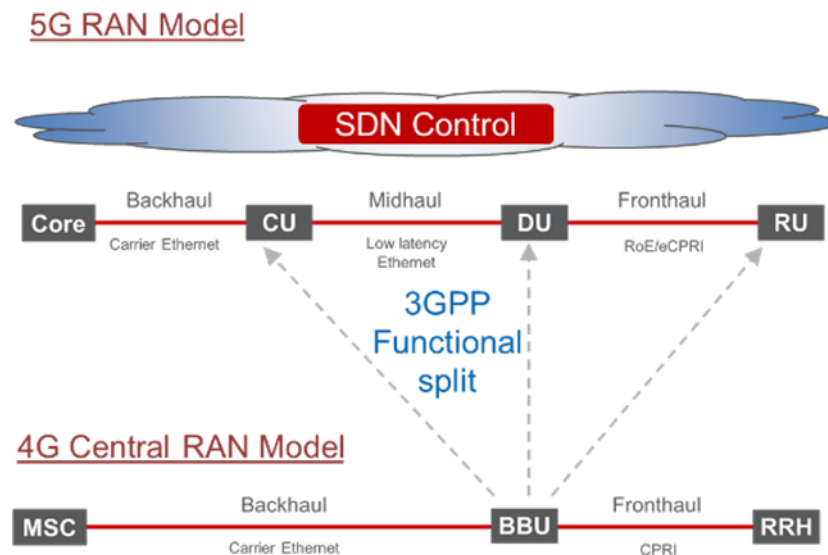
## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - 4G RAN evolution to 5G .....	3
Figure 2 - Comparison of Distributed to Centralized RAN Functional Splits .....	4
Figure 3 - Add 5G to 4G Distributed RAN.....	5
Figure 4 - Traditional vs. Cloud Control Plane Router.....	6
Figure 5 - Network Slicing Using Cloud Control Plane Router.....	7

# Introduction

Planning and deploying a mobile network to support a myriad of 5G applications will be no easy feat, considering the complexities of these new architectures and the interdependencies between the RAN and transport network.

RAN transport rates for 5G will be over 15 times greater than those available in 4G LTE and its variants. However, using the same operation of the 4G RAN for 5G would yield transport rates for optics and platforms price prohibitive. To mitigate this situation, the 3GPP standardized a new RAN model splitting the processing functionality of the 5G BBU into several blocks, thereby reducing the transport rate requirements.



**Figure 1 - 4G RAN evolution to 5G**


The key building blocks of the Next Generation RAN (NG-RAN) architecture are the centralized unit (CU), distributed unit (DU) and remote radio unit (RU). Fronthaul transport between the RU and DU will use the more efficient eCPRI protocol which provides higher performance at a lower cost per bit than CPRI used for 4G services.

The IEEE has standardized the latency budgets for the new 5G RAN. These budgets are similar to the original 4G fronthaul and backhaul segments with the exception of the new ultra-reliable-low-latency connectivity (uRLLC) use cases. In the fronthaul, these result in a 50 microsecond latency budget. Backhaul remains at 10 milliseconds and fronthaul at 100 microseconds for all but uRLLC applications. The new area of transport is the “midhaul” or next generation fronthaul II (NGFH II) section, which will vary from one to three milliseconds in latency budget as per the IEEE 1914.3.

With 5G services, a new form of RAN topology is emerging. The predominate topology in the RAN today is the distributed RAN. The distributed RAN consists of all the 4G elements- remote radio head (RRH) and baseband unit (BBU) at the cell site. This topology has the lowest latency. Next is the centralized RAN where the BBU is centralized at a location within 20 kilometers of the cell site. The centralized RAN configuration enables the BBUs at the central location to pool resources to address the demands of the cell sites. This eliminates the risk of over or under engineering the individual cell site

with a specific capacity of BBU. Cell site aggregation also enables two or more cell sites to address demands of an individual mobile user.

The “virtualized” RAN is the new model for 5G. The processing elements of the RAN, i.e. the DU and CU- will ultimately be virtualized because vertical network slicing will initiate in the DU. This is the most flexible topology as it can be dynamically repurposed. The Next Generation Mobile Networks (NGMN) consortium of service providers has developed several RAN topologies. These models vary from a distributed RAN with site cost and complexity balanced with less demanding transport to a centralized RAN, which provides a coordination gain and yields a high-performance transport layer.



**Figure 2 - Comparison of Distributed to Centralized RAN Functional Splits**

This flexibility is accomplished by splitting the functionality of the 5G RAN elements to deliver the performance requirements needed for the upcoming 5G use cases. The virtual RAN and transport topology will work very closely together. Service providers will be able to develop a single infrastructure to address the upcoming use cases and multi-tenant operation without having to dedicate assets to one topology type or having multiple network element overlays.

## Radio Type Drives Optimum Transport Option

The radio types used for 5G are millimeter wave (mmWave) and sub-6 gigahertz. Each has its pros and cons. The mmWave radio is using the higher frequencies offering high capacity service but coverage is limited to about 100 meters because its high frequencies are challenged with walls and obstacles. Therefore, separate indoor and outdoor RAN networks will be used. In the outdoor environment a densification strategy is needed resulting in deployments on streetlight and utility poles, sides and tops of buildings much like small cell installations. These non-traditional sites will require much more fiber facilities and have very limited power and footprint available.

Sub-6 gigahertz radios have less capacity than mmWave but have better coverage of about one kilometer. They do not have the penetration issues of mmWave and are installed on traditional cell towers.

## Optimum Transport for mmWave

Mobile Network Operators are deploying both radio types for 5G applications. Today these radios are used for fixed wireless access (FWA) to offer high speed Internet to residential and small to medium businesses. Transport options include dedicated dark fiber (DDF) as the first choice, but lacks integrated remote visibility. When DDF is in short supply, a multiplexing capability is needed to extend the capacity of the fiber. Two such multiplexers are the traditional wave division mux (WDM) and the new time sensitive networking (TSN) otherwise referenced as an Ethernet mux. Both technologies transport 4G CPRI (10 Gbps), 5G eCPRI (10 and 25 Gbps) and gigabit Ethernet up to 25G.


The WDM has several variants ranging from an all-passive technology, to semi-passive to a full-active technology. The all-passive technology is the lowest cost but lacks any remote visibility for performance monitoring. In semi-passive, the cell site end is full passive while the hub end uses active transponder technology. Smart optics are used in the semi-passive to establish self-tuning automation and a communications channel for remote visibility. Intelligence at the hub end provides integrated optical DDM, OTDR, performance monitoring and latency measurement capabilities working with the smart optics communications channel. Full-active WDM is the most expensive means of WDM transport but offers the most capability including full remote visibility and topology options for self-healing operation.

For mmWave deployments, given the challenges at the cell site in power and footprint, the semi-passive system is optimum in terms of providing a cost effective solution with a level of remote visibility. The cell site end is fully passive and does not require any power. The footprint for the outside plant WDM and enclosure is just a little bigger than a tall coffee container saving on footprint where it is at a premium. Linear or point to multi-point topology is available on the semi passive system for ease of deployment.

## TSN or Full Active WDM for Sub 6 GHz Radio Fronthaul

When deploying sub 6 GHz radios at traditional cell towers based on the coverage capability, there will be multiple sectors resulting in many channels for transport to a central location. Service providers will need to transport in the fronthaul the new 5G services along with the legacy 4G channels, which are highly inefficient. This presents a major challenge in the total number of channels to be transported.

Using the traditional WDM approach would require many expensive 10 Gbps and 25 Gbps optics to drive services over the fronthaul. This results in high cost and large footprint for the transport layer. However, up to a mix of 40 channels of: 4G CPRI, 5G eCPRI and Ethernet channels can be transported over a single fiber strand, justifying a WDM approach when channel counts are high and fiber assets are near depleted. Alternatively, if the capacity of traffic for transport is such that minimal fiber is required, the time sensitive networking (TSN) is a more cost effective option. The TSN approach will utilize a packet to multiplex channels in the fronthaul and can also translate the inefficient CPRI to eCPRI protocols reducing the total bandwidth capacity. The ORAN Alliance has specified a functional split, 7.2x, to translate highly inefficient CPRI traffic to eCPRI using the low order physical layer 1 processing (Low PHY). This function would reduce CPRI bandwidth capacity up to 5:1. Another translation approach for CPRI to eCPRI is the CPRI Cooperation's eCPRI v2.0. The eCPRI v2.0 does not do a full translation, instead it reformats the IQ data in the CPRI frame to that of eCPRI. This is useful in evolving distributed 4G cell sites to 5G while maintaining the legacy 4G service. When evolving the 4G cell site to 5G the backhaul capacity will increase from 1 GE to 10 GE.



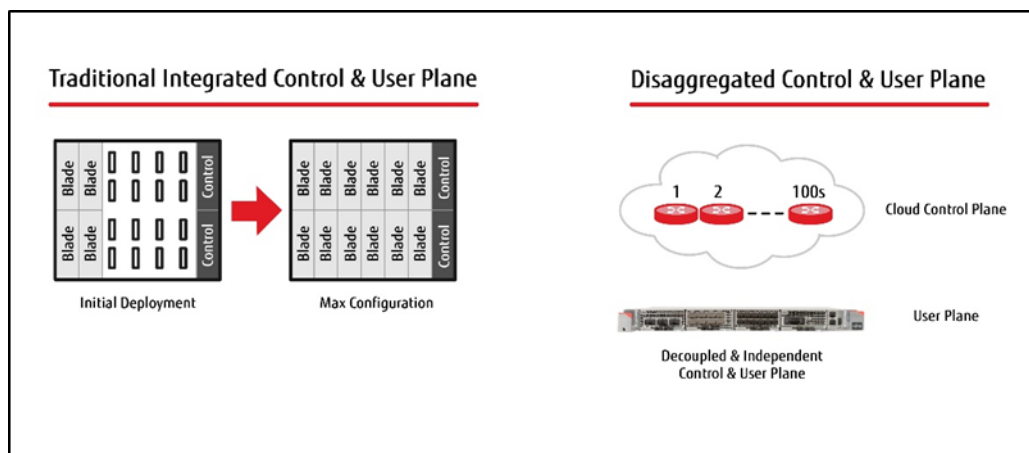
**Figure 3 - Add 5G to 4G Distributed RAN**

The compact TSN mux offers high density transport for 4G, 5G and Ethernet services, translation from CPRI to eCPRI, at costs lower and in a smaller footprint than full active WDM using one or more 100 GE network connections.

## Virtual Networks via Network Slicing

The virtual RAN will offer the service provider the greatest flexibility via a single transport infrastructure with multiple micro-services. This is accomplished using a mix of software-defined networking (SDN), network functions virtualization (NFV), and end-to-end network slicing for the RAN, edge transport and core networks. This RAN virtualization sets the groundwork for a single physical network infrastructure representing multiple virtual network configurations each representing a network slice, hence the term “network slicing.” Each network slice is a complete virtual network within the infrastructure.

The edge transport network establishes a common infrastructure using programmable and disaggregated network elements. Edge transport routers are used from the DU, where the network slice point begins, to the core offering dynamic multipoint connectivity. To assist in maintaining a predictive low-latency operation, MPLS segment routing (MPLS-SR) is the most common infrastructure technology used to facilitate network slicing.



**Figure 4 - Traditional vs. Cloud Control Plane Router**

Traditional router architectures are vertically integrated, self-contained network elements. They consist of a chassis with line cards deployed in predefined slots along with switch fabric and control cards in other slots. Connectivity between line cards and switch cards is enabled via electrical backplane traces commonly referred to as serializer/deserializer (SerDes). The number of traces between slots and the speed with which the traces are clocked determines the maximum inter-slot communication capacity. This architecture requires the alignment of three hardware components: the line card, the switch fabric cards, and the backplane. Service providers are challenged in three areas when specifying a router platform for their 5G network:

- Determining the right capacity and performance for the site demands
- Minimizing the physical and environmental allocations, and
- Scaling platform capacity and performance for the long term.

# Control in the Cloud

Router vendors typically offer a mixture of low-, medium-, and high-capacity performance units. Sizing the integrated router capacity is challenging because the control plane, backplane speed, and chassis capacity limit the performance and scaling of the user plane blades. Under-allocating the router performance can risk loss of opportunity, whereas over-allocated router performance results in capex inefficiency.

During initial installation only 20% to 30% of the router capacity is utilized but the chassis footprint, power, and thermal reserve all have to be fully allocated, resulting in cost-inefficiencies. Anytime the capacity of the slots is increased, all three elements must move in lockstep.

Because service providers loathe the idea of forklifting the chassis/backplane, vendors try to future-proof their node designs to support capacity expansions, including cooling, power, and backplane traces. However, because any chassis design utilizes the most cost-effective, commercially technology available at the time, there are limits to how far vendors can future-proof the network element. Once these measures are exhausted, additional capacity enhancements require replacement with a newer chassis.

To resolve the limitations of a traditional router, the next-generation router will employ a programmable disaggregated control and user plane architecture. The control plane is completely independent of the user plane, and in advanced models it is hosted and executed in the cloud. Incorporating cloud native technology and routing protocol isolation into the disaggregated router via a cloud control plane resulting in a single 1RU blade element capable of dynamically producing hundreds of router instances for RAN services and customer isolation. The virtual routing segments, quality of service (QoS), and resiliency requirements are provisioned in the cloud using automation for the virtualized service.

Figure 5: Network Slicing Using Cloud Control Plane Router

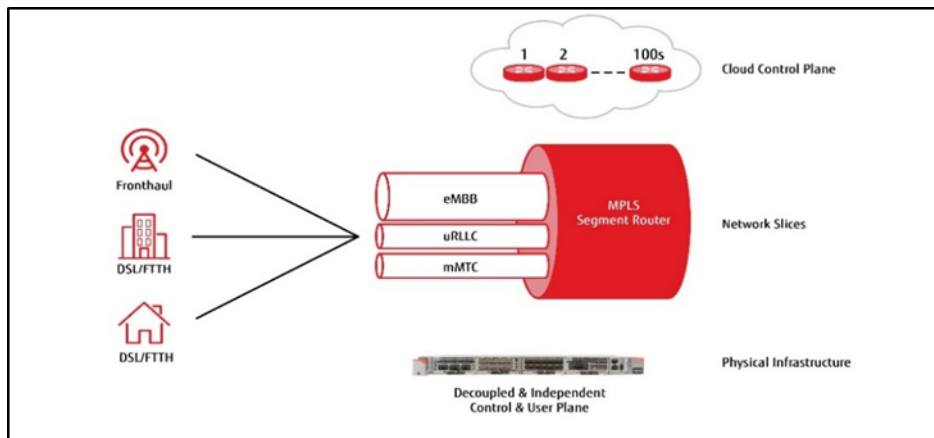


Figure 5 - Network Slicing Using Cloud Control Plane Router

Once the cloud control plane calculates mapping for each service, the control information is then pushed down to the router user plane infrastructure. If a physical site has a catastrophic failure, its virtual routing profile can be moved in the cloud control plane to another physical site, simplifying resiliency operations. Applying this architecture to the router optimizes physical/environmental cost-efficiencies, simplifies network engineering, reduces infrastructure capacity risks, and offers superior performance scaling.



# Scalable Slices

The network orchestrator coordinates this ecosystem between the core, edge transport, and RAN elements. As network slices are established via DU asset allocations and multiple CU-UP terminations, the transport network establishes router instances to support these individual services and customers providing the transport QoS guarantees.

A traditional router architecture with integrated control and user plane is initially cost-inefficient, has risks of over or under performance based on chassis size, and has limited scaling functionality over the long term. On the other hand, the disaggregated cloud control and user plane router approach establishes a single transport infrastructure with the ability to dynamically virtualize multiple networks cost-effectively in a highly scalable fashion. As today's networks continue to evolve, this dynamic flexibility will be key to allowing tomorrow's architecture to meet diverse needs for capacity, latency, and performance, fulfilling the 5G promise.

# **Blueprint for 3 GHz, 25 Gbps DOCSIS®**

## **Getting 25 Gbps PON-Like Performance Out of HFC**

A Technical Paper prepared for SCTE•ISBE by

**John T Chapman**

CTO Cable Access and Fellow  
Cisco Systems  
3700 Cisco Way, San Jose, CA 95134  
408-526-7651  
jchapman@cisco.com

**Hang Jin**

Distinguished Engineer  
Cisco Systems  
3700 Cisco Way, San Jose, CA 95134

**Thushara Hewavithana**

Senior Architect, Connected Home Division  
Intel Corporation  
5000 West Chandler Blvd, Chandler, AZ 85226  
602-245-1468  
thushara.hewavithana@intel.com

**Rainer Hillermeier**

General Manager Design and Operations Qorvo Germany  
Qorvo  
Loeffelholzstrasse 20, 90441 Nuremberg, Germany  
+49-911-9411-167  
Rainer.Hillermeier@qorvo.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	5
Concept.....	5
1. The Laws of Physics vs the Needs of the Market.....	5
2. Theorem 1: Three technical levers - power, bit loading and distance .....	5
3. Theorem 2: Coax lengths .....	6
4. 3 GHz ESD with Distributed Power Amplification .....	7
Spectrum Plans for 3 GHz ESD .....	8
5. Methodology .....	8
5.1. Transition Bands .....	8
5.2. Coexistence with MoCA.....	8
5.3. Coexistence with Other RF signals .....	10
5.4. Simplified Bandwidth Accounting .....	11
5.5. Rigorous Bandwidth Accounting .....	12
6. 3 GHz Premium Plan with a 1218 MHz Cross-over with 10 Gbps US.....	14
7. 3 GHz MoCA Plan with a 1100 MHz Cross-over .....	16
8. 3 GHz Legacy Update Plan with a 1002/862/750 MHz Cross-over.....	17
9. 3 GHz Low Power Plan with a 684 MHz Cross-over .....	18
10. Comparison with DOCSIS 4.0 with FDX.....	19
11. Comparison with DOCSIS 4.0 with 1.8 GHz Extended Spectrum.....	20
12. Summary .....	21
3 GHz Passive Taps.....	21
13. 3 GHz Tap specifications .....	21
14. 3 GHz Tap Prototype Test Results .....	23
Power Plans .....	27
15. Distributed Power Amplification .....	27
16. Optimization of Transmit Power for Capacity .....	30
16.1. Theoretical Framework .....	30
16.2. Backwards Compatible and Optimal Power allocation .....	32
16.3. Results and Conclusions .....	33
DOCSIS PHY Optimizations .....	34
17. Introduction.....	34
18. A Systematic Approach to Selecting OFDM Parameters.....	35
18.1. Target MER.....	35
18.2. Cable Propagation Channel.....	36
18.3. Cyclic Prefix for OFDM .....	37
18.4. OFDM Symbol Length .....	38
18.5. OFDM Time variations.....	39
19. Conclusion .....	42
Power Amplifier Circuit Design Development .....	43
20. Output Stage Gain Blocks for HFC Amplifiers and Nodes .....	43
21. 3 GHz Output Stage Gain Block Simulation Model .....	44
22. Conclusion .....	47
Additional Deployment Considerations.....	48
Summary.....	49

Acknowledgements .....	49
Abbreviations.....	50
Bibliography & References .....	51

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1: Three Technical Levers.....	5
Figure 2: Reference HFC Architecture .....	6
Figure 3: 3 GHz ESD with Distributed Amplification.....	7
Figure 4: MoCA Spectrum for a DOCSIS System.....	9
Figure 5: 3 GHz Premium Plan .....	14
Figure 6: 3 GHz MoCA Plan.....	16
Figure 7: 3 GHz Legacy Plan.....	17
Figure 8: 3 GHz Low Power Plan .....	18
Figure 9: Frequency Plan with Extended FDX.....	19
Figure 10: 1.8 GHz Static FDX Frequency Baseline .....	20
Figure 11: 3 GHz tap prototype.....	23
Figure 12: Test data for Tap14 insertion loss .....	24
Figure 13: Test data for Tap14 tap loss.....	24
Figure 14: Test data for Tap20 insertion loss .....	25
Figure 15: Test data for Tap20 tap loss.....	25
Figure 16: Test data for Tap26 insertion loss .....	26
Figure 17: Test data for Tap26 tap loss.....	26
Figure 18: The spectrum partition .....	28
Figure 19: Hybrid Active Tap.....	28
Figure 20: HAT block diagram .....	29
Figure 21: Simulation results of total AC power consumption vs number of amplifiers for N+2 network... 29	29
Figure 22: PA nonlinear distortion characterisation (source: Qorvo 3 GHz PA simulation data).....	31
Figure 23: Optimal Power Allocation for backwards compatible 3 GHz ESD .....	33
Figure 24: Comcast Model I.....	33
Figure 25: Capacity with and w/o transmit power optimization.....	34
Figure 26: Optimized SNR (solid) and SNR achieved with flat transmit PSD (dashed).....	35
Figure 27: Cyclic Prefix Overhead.....	38
Figure 28: CMTS Phase Noise PSD – DSB, referred to 1218 MHz.....	41
Figure 29: Efficiency Development of Output Stage Gain Blocks in so-called Power Doublers.....	43
Figure 30: Simplified model schematic of balanced design in cascode configuration employing GaAs semiconductor technology for FET1 and FET2 and GaN for FET3 and FET4.....	44
Figure 31: Multi-carrier distortion simulation test bench .....	45
Figure 32: Simulation RF input and output spectrum and the MER per subcarrier over frequency .....	46
Figure 33: Simulated compression characteristic of averaged MER versus TCP .....	47

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1: DOCSIS 3.1 Downstream Capacity .....	12
Table 2: DOCSIS 3.1 Upstream Capacity .....	13
Table 3: 3 GHz Premium Plan with Classic FDX .....	14
Table 4: 3 GHz Premium Plan with Extended FDX.....	14
Table 5: 3 GHz MoCA Plan with Classic FDX .....	16
Table 6: 3 GHz MoCA Plan with Extended FDX.....	16
Table 7: 3 GHz Legacy Plan with Classic FDX (Three Variations) .....	17
Table 8: 3 GHz Legacy Plan with Extended FDX .....	18
Table 9: 3 GHz Low Power Plan with Classic FDX.....	18
Table 10: 1.218 GHz DOCSIS 4.0 with Classic FDX .....	19
Table 11: 1.8 GHz DOCSIS 4.0 with No FDX.....	20
Table 12: Spectrum Plan Data Capacity Summary.....	21
Table 13: 3 GHz Tap Specifications.....	22
Table 14: DOCSIS 3.1 DS Micro-Reflection Bound (mask).....	36
Table 15: Coax Cable Losses for 0.5 $\mu$ s echo .....	36
Table 16: Coax Channel Model at 2 GHz and 3 GHz .....	37
Table 17: CMTS RMS Jitter Spec .....	39
Table 18: CMTS Phase Noise mask at 1002 MHz.....	40
Table 19: Combined Phase Noise Mask for CMTS.....	40
Table 20: ICI/dBc ratio in dB due to CMTS Phase Noise from Figure 28 .....	42

# Introduction

The path to 10 Gbps downstream has been laid out with DOCSIS 3.1 using spectrum up to 1.2 GHz. DOCSIS 4.0 adds full duplex operation for a 5 Gbps upstream or an extended spectrum high-split option using 1.8 GHz. Both support similar 10 Gbps throughputs. Yet, the PON world is already 10 Gbps downstream and upstream, and due to deliver a 25 and 50 Gbps standard in 2020. What will cable's response be?

The good news is that 25 Gbps DOCSIS can be built on the same wiring infrastructure of 10G DOCSIS 4.0 – the same digital Fiber to the Node, the **trunk** coax to the TAP, and drop coax to the Home. Furthermore, the transition from 10 Gbps speeds to 25 Gbps is not imposing that the fiber need to go any deeper. In DSL, the transition to higher speeds requires taking fiber deeper and making the copper link shorter whereas this is not the case for “25G” DOCSIS.


At the same time, extended spectrum DOCSIS (ESD) will require a complete product family refresh with CMTS/RPD - node - amp - tap - CM which will take time and investments from both vendors and operators. This paper shows how a roadmap to 3 GHz, 25 Gbps ESD is possible with 1.8 GHz ESD as a steppingstone. This paper explains the various techniques to get there and how the transition can be done. Considerations on power consumption, coexistence with FDX, spectrum plans and interference from MoCA, LTE and Wi-Fi are discussed.

## Concept

### 1. The Laws of Physics vs the Needs of the Market

The state of the art for nodes today is the DOCSIS 3.1 FDX N+0 specification with a 54 MHz to 1.218 GHz downstream spectrum with a total composite power (TCP) of 73.8 dBmV and tilt all the way up to 1218 MHz, or 71-72 dBmV with the appropriate stepping down of power for the last channel. To go beyond the 1.218 GHz barrier, something has to give or be changed. First, we start with two fundamental theorems and then lead to a course of action.

### 2. Theorem 1: Three technical levers - power, bit loading and distance



**Figure 1: Three Technical Levers**

There are three technical “levers” that can be changed, although pushing on one generally means pulling back on the other two. These levers are shown in Figure 1.


The first lever is power. Power is the most popular one. In DOCSIS 3.1, the TCP of the node was pushed farther than it had ever been before and there is not much more room to push it. Depending on the efficiency of the amplifier, this requires more node power. The output power stage of a node is built with Class A power amplifiers (PAs). If a class A amplifier has a 1% efficiency, then a 1 watt output would require 100 watts of power supply source input. Moving that amplifier to 2% efficiency would drop the power supply usage to 50 watts which is a huge difference.

As we will see later in the paper, newer silicon technologies have increased the efficiency of the PA. Technologies like digital pre-distortion (DPD) can correct for some of the non-linearities in the amplifier and allow it to be pushed a bit harder. The combination of all of these techniques is really not enough to expand the frequency range from 1.2 GHz to 1.8 GHz or 3.0 GHz.

The second lever is bit loading. Bit loading refers to how many bits are represented per hertz of bandwidth. It gets more complex in implementation with symbol rates and sub-carriers, each with its own constellation. But, in general, 10 bits per hertz uses  $2^{10} = 1024$ -QAM constellation. The ability to support this constellation depends on the difference between power level of the channel (signal) and the corresponding modulation error ratio (MER) which is basically a noise floor. If the power level goes down, say due to greater attenuation at higher frequencies, and the noise floor is flat, the transmission channel will still work but may require a lower modulation. This is the approach that the 1.8 GHz ESD specification is considering.

The third lever is distance. Distance is the space between the node and the first amp and then the distance between amps. Hybrid fiber/coax (HFC) plants are built by using higher-output power nodes; the RF power of the RF signals decreases as the signals travel through the span of cable and taps, until the RF power is so low that it needs to be amplified again. This plant layout is shown in Figure 2. Thus, the span of distance between nodes is dependent on the output power of the nodes and the cable and tap loss. If operation is required at a higher frequency where the cable and tap loss is greater, and the RF power output of the node stays the same, then a shorter distance between nodes and amplifiers is required. The 3 GHz ESD proposal in this paper leverages this principle but does so with an interesting twist.

### 3. Theorem 2: Coax lengths




**Figure 2: Reference HFC Architecture**

The DOCSIS specification [1] assumed maximum one-way transit delay in a cable network is 0.800 ms, which is the equivalent of about 100 miles of single-mode optical fiber. That plant today can run at 1.218 GHz at a modulation of 4K QAM. That's impressive. Of that plant length, though, most of it is the fiber run. The coax plant after the optical node is more on the order of about 5000 feet (1500 m). In fact, with the distributed access architecture (DAA), there is no longer an analog optical distance to consider.

The second theorem here is that there are no 5000 foot coax runs. Instead, the coax plant is made up of a series of very short pieces of coax separated by passives such as taps, and actives such as nodes and amps. As a thought experiment, if we replaced all passive and actives with 3 GHz amps, it would take very little amplification to drive that short piece of coax and the attenuation of that coax would be of little consequence. The frequency limitation would only be based on the quality of the coax. Most of the coax in the plant should be 3 GHz capable unless it has physical defect or is otherwise damaged.

## 4. 3 GHz ESD with Distributed Power Amplification



**Figure 3: 3 GHz ESD with Distributed Amplification**

If we combine these two concepts together, we have the answer, but with a twist. The twist is that there are two loss plans that are used, not one. The first loss plan is the one that already exists that has set the deployed spacing of the nodes and amplifiers. The second loss plan is the new high frequency loss plan. They both exist at the same time and the new high loss plan can be designed almost independent of the current loss plan.

Let's assume for now we limit the lower frequency loss plan up to 1002 MHz which is a common deployed maximum downstream frequency today. Then for the new loss plan from the 1002 MHz to 3000 MHz, we introduce small extended spectrum amplifiers (ESA) in the cable span between the established nodes and amplifiers. We are referring to this system as distributed power amplification (DPA).

It turns out that we do not have to put a lot of these ESAs between the larger amplifiers. Two or three per span should be sufficient. They could be co-located beside a tap or even within a tap housing. If they are within a tap housing, we are referring to that as a hybrid-active tap (HAT). If the ESA were to die, only the extended spectrum would be impacted, so the plant would continue to work but at diminished capacity.

We can also choose a lower power ESA by just defining more amps. The loss plan of the higher frequencies can be rebalanced independent of the lower loss plan. Also, as we will see in Section 15 on Distributed Power Amplification, that by changing the separation frequency between the higher band and the lower band, we can change the overall power consumption and efficiency of the HFC network.

We discuss these principles in subsequent sections. But first, let's do some spectrum planning and look at what other factors might influence the choice of the boundary between the lower and upper bands.



# Spectrum Plans for 3 GHz ESD

There are a variety of spectrum plans that could be chosen, depending upon the legacy frequency plan of the plant, what the spectrum is used for and how much power needs to be saved. These spectrum plans also show operation with and without FDX DOCSIS.

## 5. Methodology

### 5.1. Transition Bands

In the 3 GHz extended spectrum approach described in this paper, there is an extended spectrum transition band (ETB) required between the lower legacy and upper ESD spectrum plans for several reasons. This is explained in more detail in Section 15 and summarized here.

1. The ETB requirement is due to a diplexer that is located in the in-line amplifiers between the high and low frequencies. In this approach, the size of the ETB is set to the same size as the FDX transition band (FTB) which is 17.5% of the lower frequency of the transition band. This is described in Section 15.
2. The power level between the last carrier in the lower frequency band and the first carrier in the upper band may be different. This means the out-of-band spurious noise from the higher power carrier might impact the lower power carrier. This is described in Section 15.
3. There is a transition band at the top of the FDX band for FDX CMs and this transition band may or may not line up with the other two transition bands. For current FDX DOCSIS, this transition band is 17.5% and extends from 684 MHz to 804 MHz [4].

Part of the art of frequency planning is to maximize usage of the ETB if possible with some other services such as video, legacy DOCSIS, MoCA or other transition bands.

### 5.2. Coexistence with MoCA

Multimedia over Coax Alliance (MoCA) is a technology that provides Ethernet over coax within the residential environment. It is possible to physically isolate MoCA and DOCSIS within a home at installation time. When the drop cable terminates at the house, it would hit a two-way splitter. One leg of that splitter would go to the DOCSIS CM. The other leg of that splitter would go through a MoCA filter and then to the rest of the residential network. Note that this approach supports legacy STB MoCA network but would not support Wi-Fi extension ports from the CM over MoCA.

Unfortunately, many homes do not have a MoCA filter and thus the DOCSIS spectrum and the MoCA spectrum above 1 GHz may mutually interfere with each other. In this section, we describe an approach to manage that interference. But first, some background on MoCA.

MoCA 1.1 Band D is 400 MHz wide and extends from 1125 MHz to 1525 MHz. 50 MHz MoCA 1.1 channels can be placed starting at 1125 MHz in 50 MHz increments. MoCA 2.0 Band D is 550 MHz wide and extends from 1125 MHz to 1675 MHz. A single 100 MHz MoCA 2.0 channel can be placed starting at 1125 in 25 MHz steps. MoCA 2.0 also has a dual bonded channel that occupies 225 MHz.


MoCA 1.1/2.0 requires a 125 MHz guard band between adjacent unassociated MoCA channels. The SCTE 235 2017 operational practice [2] requires a 25 MHz guard band between any MoCA spectrum and DOCSIS spectrum, although Figure 2-7 in the MoCA 2.0/2.5 RF specification [3] would prefer more like 57 MHz or 81 MHz.

The idea here is to open up some spectrum and steer MoCA into it. In theory, that spectrum gap could also contain a separate DOCSIS 3.1 channel of which there is a selective membership that never interferes with MoCA on the plant, but that is an advanced algorithm not considered in this paper's bandwidth calculations.


There are several schemes that could be used. Some of these approaches starting at the lowest frequencies permitted are shown in Figure 4.

- (a) Single channel: This will support either one MoCA 1.1 channel with 100 MHz total bandwidth, or one MoCA 2.0 channel with 150 MHz total bandwidth (signal plus guard bands). This may be quite practical if most homes have only one common MoCA channel. This approach is used in Section 0.
- (b) Single channel: Similar to (a) but the DOCSIS guard bands have been increased to 75 MHz for MoCA 1.1 and 50 MHz for MoCA 2.0. This is more in line with the MoCA 2.0/2.5 RF spec co-existence examples [3] . This approach is described in Section 0.
- (c) Single channel: Similar to (b) but shifted by 100 MHz. This approach is described in Section 0.
- (d) Dual channel: The MOCA 2.0 specification requires 125 MHz between two unrelated MoCA channels. This mode is described in Section 10.


(a) Single Channel with MOCA 1.1 or 2.0 at lowest frequencies



(b) Single Channel with MOCA 1.1 or 2.0 at 1100 MHz



(c) Single Channel with MOCA 1.1 or 2.0 at 1218 MHz



(d) Dual Channel MOCA 1.1 and 2.0 at 1100 MHz




Figure 4: MoCA Spectrum for a DOCSIS System

In 2019, MoCA 2.5 was released. MoCA 2.5 supports bonding across five directly adjacent 100 MHz channels with a maximum throughput of 2.5 Gbps in the same extended Band D. MoCA 2.5 could exist above the 1218 MHz legacy spectrum but severely impacts a 1.8 GHz extended spectrum frequency plan.

This white paper includes MOCA interference in its spectrum recommendations.

### **5.3. Coexistence with Other RF signals**

With any new spectrum allocation comes previous tenants in the form of other spectrum. There are two fundamental ways that the old tenants and the new tenants do not get along:

1. Signal ingress when RF energy exterior to the coax plant comes into the plant and interferes with the cable signal.
2. Signal leakage (signal egress) where the RF energy from the HFC escapes and interferes with external RF spectrum.

In theory, the HFC plant is shielded and there is a high isolation between the two spaces. In practice, all it takes is a loose F connector somewhere. The challenge here, though, is that a defect in the coax plant that does not create ingress or leakage below 1 GHz may for plant that go up to 1.8 or 3.0 GHz.

Let's look at a few culprits. The problem can be split into three frequency zones of interest.

#### *Below 1218 MHz:*

In the over-the-air environment, there are all kinds of signal allocations that represent potential sources of ingress and direct pickup interference. Among them are UHF TV broadcast (470 MHz to 698 MHz), LTE (698 MHz to 806 MHz), various cellular and trunked radio services in the 800 MHz band, ISM in the 902 MHz to 928 MHz band, amateur radio service in the 420 MHz to 450 MHz range (as well as shared operation in the 902 MHz to 928 MHz ISM band). UHF broadcast TV services in the 600 MHz spectrum (just below the current LTE band) are in the process of being relocated to lower frequencies, as new mobile services begin operation there. And don't forget TV, FM broadcast, two-way radio, etc., below 470 MHz.

The 600 MHz spectrum has been auctioned off, and there is a multi-year transition period underway in which UHF TV broadcasters are vacating the 600 MHz spectrum (moving to lower frequencies). This transition is supposed to be complete in 2020. As such, in addition to the 698 MHz to 806 MHz LTE band, there will be new non-broadcast LTE-like services operating below 698 MHz that will be potential sources of ingress interference and susceptible to cable network leakage.

It is worth noting that with good due diligence and a lot of measurements, the HFC plant works despite this interference.

#### *From 1218 MHz to 1794 MHz:*

In the over-the-air spectrum one will find more LTE and cellular operation; GPS (not likely to cause interference to cable services, but could be interfered with by leakage from a cable network); amateur radio in the 1240 MHz to 1300 MHz spectrum (keep in mind that U.S. ham operators are allowed up to 1500 watts PEP transmitter power); and more mobile (cellular, etc.) services.

*From 1794 MHz to 3 GHz:*

One can find more mobile services, Wi-Fi, amateur radio, broadcast studio links (point-to-point), microwave ovens (2.45 GHz), and so on. Wi-Fi is probably the biggest challenge as it will co-exist with DOCSIS within the home gateway. Wi-Fi (801.11b/g/n/ax) in North America is from 2.401 MHz to 2.483 MHz. Zigbee and Bluetooth also share these frequencies.

This whitepaper notes this interference for further study.

#### **5.4. Simplified Bandwidth Accounting**

To keep bandwidth calculations in the following section simple but reasonably accurate, the downstream is calculated with 4K QAM (12 bits/Hz) at 80% efficiency, so 9.6 bits/symbol/Hz net. The upstream is calculated at 1K QAM (10 bits/symbol/Hz) with 80% efficiency, so 8 bits/Hz net. The results are approximate and rounded off for readability.

DOCSIS 3.1 OFDM downstream channels are 192 MHz maximum width and DOCSIS 3.1 OFDMA upstream channels are 96 MHz maximum width. Fractional OFDM channels are quoted for simplicity, but in reality, may be a combination of OFDM and SC-QAM. Also, the final frequency assignments may differ slightly from what is in this paper depending upon final channelization choices.

It should also be noted that legacy video in the downstream will reduce the spectrum available for DOCSIS and thus DOCSIS will have a lower throughput when video carriers are present. Also, if a 3 GHz DAA node is fed with 25 Gbps fiber, operation beyond 25 Gbps is not relevant.

## 5.5. Rigorous Bandwidth Accounting

Here is the detailed version and how the 20% overhead was observed. The results are close to the same.

**Table 1: DOCSIS 3.1 Downstream Capacity**

Downstream			
BW	192	192	MHz
Guardband	2	2	
FFT size (4K or 8K FFT)	4096	8192	subcarriers
Subcarrier spacing	50	25	kHz
FFT duration (useful symbol duration)	20	40	μs
Cyclic prefix (CP)	1.25	1.25	μs
Effective symbol duration	21.25	41.25	μs
Number of active subcarriers	3800	7600	subcarriers
Pilot overhead	30	60	
PLC overhead (number of subcarriers)	8	16	subcarriers
Num of NCP	10	10	
QAM order of NCP	4	4	
NCP overhead	120	120	
FEC overhead	12%	12%	
Data QAM order (bits per symbol)	12	12	
Total data bit	38351	77966	
<b>L1 Throughput (Gbps)</b>	<b>1.80</b>	<b>1.89</b>	<b>Gbps</b>
Efficiency	78%	82%	
Equivalent bits/Hz	9.4	9.8	

**Table 2: DOCSIS 3.1 Upstream Capacity**

Upstream			
BW	96	96	MHz
Guardband	1	1	
FFT size (2K or 4K FFT)	2048	4096	subcarriers
Subcarrier spacing	50	25	kHz
FFT duration (useful symbol duration)	20	40	μs
Cyclic prefix (CP)	1.25	1.25	μs
Effective symbol duration	21.25	41.25	μs
Number of active subcarriers	1900	3800	subcarriers
Frame length	18	9	
Mini-slot height	8	16	
Pilot pattern	2	2	
Pilot overhead	4%	4%	
FEC overhead	11%	11%	
Data QAM order (bits per symbol)	10	10	
Total data bit	16185	32370	
<b>Throughput (Gbps)</b>	<b>0.76</b>	<b>0.78</b>	<b>Gbps</b>
Efficiency	79%	82%	
Equivalent bits/Hz	7.9	8.2	

## 6. 3 GHz Premium Plan with a 1218 MHz Cross-over with 10 Gbps US

If the goal is to get to a 10 Gbps upstream, then lots of upstream spectrum will be needed, and that spectrum should be at the lowest frequency possible so that the minimum power will be required to drive it. Figure 5 shows a 3 GHz frequency plan where the cross-over between the legacy spectrum and the extended spectrum starts at 1218 MHz.



**Figure 5: 3 GHz Premium Plan**

The current HFC equipment is designed to go to 1218 MHz, although in practice, most deployed HFC plants only have spectrum plans to 1002 MHz or less. The ETB in this approach is placed at 1218 MHz.

Let's do the numbers.

**Table 3: 3 GHz Premium Plan with Classic FDX**

Freq Range MHz	Comments	US Gbps	DS no FDX Gbps	DS with FDX Gbps
15 to 85	70 MHz, 4 ATMDA + 0.5 OFDMA	0.5		
108 to 684	576 MHz, 6 OFDMA, 3 OFDM	4.5		5.5
684 to 1218	534 MHz, 2.8 OFDM		5	5
1431 to 3000	1569 MHz, 8.2 OFDM		15	15
	<b>Total Data Capacity</b>	<b>5</b>	<b>20</b>	<b>25</b>

In Table 3, with a 684 MHz return path, the upstream data capacity is 5 Gbps. The downstream data capacity is 25 Gbps with FDX enabled and 20 Gbps with FDX not enabled.

**Table 4: 3 GHz Premium Plan with Extended FDX**

Freq Range MHz	Comments	US Gbps	DS no FDX Gbps	DS with FDX Gbps
15 to 85	70 MHz, 4 ATMDA + 0.5 OFDMA	0.5		
108 to 684	576 MHz, 6 OFDMA, 3 OFDM	4.5		5.5
684 to 1218	534 MHz, 5.6 OFDMA, 2.8 OFDM	4.3		5
1431 to 3000	1569 MHz, 8.2 OFDM		15	15
	<b>Total Data Capacity</b>	<b>9.3</b>	<b>15</b>	<b>25</b>

In Table 4, an extended FDX upstream path, one that has not been defined yet in the standards, is presumed. In this example, the return path is the full 15 MHz to 1218 MHz (omitting 85 MHz to 108 MHz). The upstream data capacity with 1K QAM is 9.3 Gbps. The downstream data capacity is 25 Gbps with FDX and 15 Gbps without FDX.

The 9.6 Gbps upstream is very impressive, but it is not quite 10 Gbps. There are several refinements that can provide 10 Gbps upstream:

1. If the US were run at 2K QAM which provides 10% more data capacity (currently required at the CM but not at the CMTS), the upstream throughput would increase to  $\approx 10.3$  Gbps.
2. If the US were taken to 1300 MHz with 1K QAM, it would provide 10 Gbps.
3. If the upstream were taken to 1260 MHz, that would exactly be 12 OFDMA channels and 9.7 Gbps, which is a convenient design point and should be close enough to 10 Gbps.
4. If bonding across ATDMA legacy is a problem, and a 10 Gbps upstream service is all above 108 MHz, then 13 OFDMA channels would be needed which would push the return path upper bound to 1356 MHz.

The advantage of this approach is that it would allow for growth of the upstream plant to go to 10 Gbps. The ETB should be usable for MoCA co-existence using Figure 4 option (c), but is not usable for legacy MPEG-TS video. The ETB would also provide the FDX transition band for the extended FDX CMs. Also, if the ESA failed, the passive path that would remain would contain the entire 1218 MHz spectrum.

The disadvantage of this plan is that it uses maximum power. There is the maximum TCP that is already used for the legacy plan and then there is the additional power required for the extended spectrum. Of course, if the power is there, then this plan provides the most upstream spectrum with MoCA protection.

Note that at this time, that there is no extended FDX DOCSIS specification.



## 7. 3 GHz MoCA Plan with a 1100 MHz Cross-over

This is a 3 GHz frequency plan where the cross-over between the legacy spectrum and the extended spectrum starts at 1100 MHz to facilitate co-existence with MoCA and where some HFC power can be saved. This is shown in Figure 6.

M  
O  
C  
A

**Figure 6: 3 GHz MoCA Plan**

Let's do the numbers.

**Table 5: 3 GHz MoCA Plan with Classic FDX**

Freq Range MHz	Comments	US Gbps	DS no FDX Gbps	DS with FDX Gbps
15 to 85	70 MHz, 4 ATMDA + 0.5 OFDMA	0.5		
108 to 684	576 MHz, 6 OFDMA, 3 OFDM	4.5		5.5
684 to 1100	416 MHz, 2.2 OFDM		4	4
1300 to 3000	1700 MHz, 8.9 OFDM		16	16
<b>Total Data Capacity</b>		<b>5</b>	<b>20</b>	<b>25</b>

In Table 5, with a 684 MHz return path, the upstream data capacity is 5 Gbps. The downstream data capacity is 25 Gbps with FDX enabled and 20 Gbps with FDX not enabled.

**Table 6: 3 GHz MoCA Plan with Extended FDX**

Freq Range MHz	Comments	US Gbps	DS no FDX Gbps	DS with FDX Gbps
15 to 85	70 MHz, 4 ATMDA + 0.5 OFDMA	0.5		
108 to 684	576 MHz, 6 OFDMA, 3 OFDM	4.5		5.5
684 to 1100	416 MHz, 4 OFDMA, 2.2 OFDM	3.3		4
1300 to 3000	1700 MHz, 8.9 OFDM		16	16
<b>Total Data Capacity</b>		<b>8.4</b>	<b>16</b>	<b>25</b>

In Table 6, an extended FDX upstream path, one that has not been defined yet in the standards, is presumed. In this example, the return path is 1100 MHz. The upstream data capacity is 8.4 Gbps and the downstream data capacity is 25 Gbps with extended FDX enabled and 16 Gbps without extended FDX enabled.

The advantages of this plan are a friendly accommodation of MoCA with slightly less power usage on the HFC plant than the 1218 MHz plan. There is no real disadvantage with this plan unless the total power consumption is still greater than what the HFC plant budget permits.

## 8. 3 GHz Legacy Update Plan with a 1002/862/750 MHz Cross-over

Earlier we stated that very little HFC plant has been upgraded to 1218 MHz. Well, there is a lot of plant that is either at 1002 MHz or still at 862 MHz or 750 MHz. If that legacy frequency band determined the channel line-up and the HFC rebuild worked with that, what would it look like? These are just more variations of the 1100 MHz approach.

ETB

**Figure 7: 3 GHz Legacy Plan**

Figure 7 is a frequency plan built on the legacy frequencies of 750/862/1002.

Let's do the numbers.

**Table 7: 3 GHz Legacy Plan with Classic FDX (Three Variations)**

Freq Range MHz	Comments	US Gbps	DS no FDX Gbps	DS with FDX Gbps
15 to 85	70 MHz, 4 ATDMA + 0.5 OFDMA	0.5		
108 to 684	576 MHz, 6 OFDMA, 3 OFDM	4.5		5.5
684 to 750	66 MHz, 0.3 OFDM		0.6	0.6
684 to 862	178 MHz, 0.9 OFDM		1.7	1.7
684 to 1002	318 MHz, 1.7 OFDM		3.1	3.1
881 to 3000	2119 MHz, 11 OFDM		20.3	20.3
1013 to 3000	1987 MHz, 10.3 OFDM		19.1	19.1
1177 to 3000	1823 MHz, 9.5 OFDM		17.5	17.5
<b>Total Data Capacity</b>		<b>5</b>	<b>20</b>	<b>25</b>

In Table 7 with a 684 MHz return path, the upstream data capacity is 5 Gbps. The downstream data capacity is 25 Gbps with FDX enabled and 20 Gbps with FDX not enabled.

In Table 8, an extended FDX upstream path, one that has not been defined yet in the standards, is presumed. In this example, the return path is 750/862/1002 MHz. The upstream data capacity is 5.5 Gbps to 7.5 Gbps. The downstream data capacity is 25 Gbps with extended FDX enabled and 17 Gbps to 20 Gbps without extended FDX enabled.

The advantage of these legacy plans is progressively less HFC power required as the transition band is lowered. Also, if the ESA fails, the entire legacy band would remain operational.

The disadvantage of these plans is that the ETB no longer overlaps a MoCA band, so if MoCA is present, it would further reduce the downstream bandwidth.

**Table 8: 3 GHz Legacy Plan with Extended FDX**

Freq Range MHz	Comments	US Gbps	DS no FDX Gbps	DS with FDX Gbps
15 to 85	70 MHz, 4 ATMDA + 0.5 OFDMA	0.5		
108 to 684	576 MHz, 6 OFDMA, 3 OFDM	4.5		5.5
684 to 750	66 MHz, 0.7 OFDMA, 0.3 OFDM	0.5		0.6
684 to 862	178 MHz, 1.9 OFDMA, 0.9 OFDM	1.4		1.7
684 to 1002	318 MHz, 3.3 OFDMA, 1.7 OFDM	2.5		3.1
881 to 3000	2119 MHz, 11 OFDM		20.3	20.3
1013 to 3000	1987 MHz, 10.3 OFDM		19.1	19.1
1177 to 3000	1823 MHz, 9.5 OFDM		17.5	17.5
<b>Total Data Capacity</b>		<b>5.5 to 7.5</b>	<b>17 to 20</b>	<b>25</b>

## 9. 3 GHz Low Power Plan with a 684 MHz Cross-over

There is an FDX transition band (FTB) located above the current FDX upstream band that FDX cable modems observe. This transition band is usable by non-FDX CMs as well as video. This could be utilized as the extended spectrum transition band.



**Figure 8: 3 GHz Low Power Plan**

Figure 8 shows the frequency plan with a transition band just above 684 MHz. In this example, the ETB is identical to the FTB. This frequency range is also in the LTE band. This means that in a system that was all DOCSIS 3.1/4.0 with no video, then this band potentially could be left empty and there would be no interference from LTE and no plant leakage into the LTE frequencies.

Let's do the numbers.

**Table 9: 3 GHz Low Power Plan with Classic FDX**

Freq Range MHz	Comments	US Gbps	DS no FDX Gbps	DS with FDX Gbps
15 to 85	70 MHz, 4 ATMDA + 0.5 OFDMA	0.5		
108 to 684	576 MHz, 6 OFDMA, 3 OFDM	4.5		5.5
804 to 3000	2196 MHz, 11.4 OFDM		21	21
<b>Total Data Capacity</b>		<b>5</b>	<b>21</b>	<b>25</b>

In Table 9 with a 684 MHz return path, the upstream data capacity is 5 Gbps. The downstream data capacity is 25 Gbps with FDX enabled and 21 Gbps with FDX not enabled.

There is no extended FDX scenario for this plan since the ETB is at 684 MHz.

The advantages of this plan are the absolute lowest power as the most downstream spectrum is shifted to the distributed amplification. The downstream extended cross-over band can be shared with the FDX guard band.

The disadvantage of this solution is that in if the extended spectrum amplifiers fail, there is only FDX downstream spectrum available, so non-FDX CMs and MPEG video STBs would lose their connection. This could be mitigated by only running FDX up to 492 MHz.

## 10. Comparison with DOCSIS 4.0 with FDX

DOCSIS 4.0 describes a DOCSIS 3.1 OFDM/OFDMA system with FDX operation. That spectrum plan is show in Figure 9.

M  
O  
C  
A

**Figure 9: 1.2 GHz Frequency Plan with Extended FDX**

Let's do the numbers.

**Table 10: 1.218 GHz DOCSIS 4.0 with Classic FDX**

Freq Range MHz	Comments	US Gbps	DS no FDX Gbps	DS with FDX Gbps
15 to 85	70 MHz, 4 ATMDA + 0.5 OFDMA	0.5		
108 to 684	576 MHz, 6 OFDMA, 3 OFDM	4.5		5.5
684 to 804	120 MHz, 0.6 OFDM (1.1 Gbps)		0 or 1.1	0 or 1.1
804 to 1218	414 MHz, 2.2 OFDM		4	4
	<b>Total Data Capacity</b>	<b>5</b>	<b>4 or 5</b>	<b>9.5 or 10</b>

In Table 10, with a 684 MHz return path, the upstream data capacity is 5 Gbps. The aggregate downstream data capacity is about 10 Gbps with FDX enabled and 5 Gbps with FDX not enabled. A single FDX CM will receive slightly less due the FTB, making the throughput 9.5 Gbps down with FDX and 4 Gbps without FDX.

## 11. Comparison with DOCSIS 4.0 with 1.8 GHz Extended Spectrum

The DOCSIS 4.0 specifications for 1.8 GHz extended spectrum are not written at this time, so this paper makes some assumptions.

- Return path choices are 204 MHz, 300 MHz, 492 MHz and 684 MHz. For this example, a return path of 492 MHz has been chosen. For 1.8 GHz extended spectrum, this allows for a 3:1 ratio of DS to US bandwidth. For a 3.0 GHz extended spectrum, this allows for a 6.4:1 ratio, both of which are respectable.
- A passive transition band was chosen without reduced guard band (RGB) and thus no active echo cancellers. This allows an inexpensive amplifier to be built.
- A transition band of 96 MHz (492 to 588 MHz) was chosen. That is a ratio of 1.95:1 which is more conservative than the 1.175:1 ratio used in the other scenarios but is what is close to what is in current discussions. 96 MHz allows sixteen 6 MHz video channels or twelve 8 MHz video channels.
- In case MoCA is a problem, the simplest MoCA solution (a) of 150 MHz was chosen.

MoCA

**Figure 10: 1.8 GHz Static FDX Frequency Baseline**

Let's do the numbers.

**Table 11: 1.8 GHz DOCSIS 4.0 with No FDX**

Freq Range (MHz)	Comments	US (Gbps)	1.8 GHz (Gbps)	3 GHz (Gbps)
15 to 85	70 MHz, 4 ATMDA + 0.5 OFDMA	0.5		
108 to 492	384 MHz, 4 OFDMA	3.1		
588 to 1100	512 MHz, 2.7 OFDM		4.9	4.9
1100 to 1250	150 MHz, 1.4 OFDM, (1.4 Gbps)		0 or 1.4	0 or 1.4
1250 to 1800	550 MHz, 5.3 OFDM		5.3	5.3
1800 to 3000	1200 MHz, 11.5 OFDM			11.5
	<b>Total Data Capacity</b>	<b>3.6</b>	<b>10 or 11.5</b>	<b>21.5 or 23</b>

In Table 11, with a 492 MHz return path, the upstream data capacity is 3.6 Gbps. The downstream data capacity for a 1.8 GHz system is about 11.4 Gbps without MoCA and 10 Gbps if allowing for MoCA and the DS:US ratio would be 3:1. If this system were later extended to 3 GHz, the downstream capacity would be 23 Gbps without MoCA and 21.5 Gbps with MoCA and the DS:US ratio would be 6.3:1.

The advantage of this system is that it is simple to understand and build. It is just an ultra-high split. There are no echo cancellers. The amplifier can be inexpensive, and the node will be less expensive. Single channel MoCA can also be accommodated by giving up some downstream bandwidth. There is enough bandwidth to support legacy video and legacy DOCSIS. CMs that have a 684/804 FDX transition band could coexist in this spectrum plan.

The disadvantage of this system is less upstream bandwidth than an FDX system. However, at ratios of 3:1 and 6.3:1, one could argue that is enough bandwidth. Still, this system will never meet the 10 Gbps upstream goals set by the 10G initiative.

## 12. Summary

The throughput of the various spectrum plans are shown in Table 13.

**Table 12: Spectrum Plan Data Capacity Summary**

Plan	ETB	Classic FDX (Gbps)			Extended FDX (Gbps)		
		US	DS FDX off	DS FDX on	US	DS FDX off	DS FDX on
3G Premium	1218 to 1431	5	20	25	9.3	15	25
3G MoCA	1100 to 1300	5	20	25	8.4	16	25
3G Legacy	1002 to 1177	5	20	25	7.5	20	25
3G Min Power	684 to 804	5	21	25	-	-	-
1.2G D4.0 FDX	684 to 804	5	5	10	-	-	-
1.8G ESD	No FDX	3.6	10	-	-	-	-

## 3 GHz Passive Taps

### 13. 3 GHz Tap specifications

The coaxial network consists of three main components: coaxial cables, taps and amplifiers. The propagation loss of coaxial cables is well studied and understood for the frequency range of interest (5 MHz to 3 GHz). The study of 3 GHz amplifier technology is discussed later in this paper. This section presents the study and prototype of 3 GHz taps.

Cisco worked with a tap ODM to prototype a series of 3 GHz taps. Three tap values were selected for the prototype development: tap14, tap20 and tap26. The targeted 3 GHz tap specifications are given in Table 13.

**Table 13: 3 GHz Tap Specifications**

Targeted 3GHz Tap Specifications					
Model 4-Way			4-14	4-20	4-26
Parameter		Value	14dB	20dB	26dB
Tap Loss (dB)	Tap Tol.	Frequency	Nominal	Nominal	Nominal
			Tap Value	Tap Value	Tap Value
	±1.5	5-10	12.5	20.5	26.5
		11-50	12.5	20.5	26.5
		51-450	12.5	20.5	26.5
		451-750	12.5	20.5	26.5
		751-870	12.5	20.5	26.5
		871-1003	12.5	20.5	26.5
		1004-1250	12.5	21	27
	±2.0	1251-1950	13	21	27.5
		1951-2250	13	21	27.5
		2251-2500	14.5	22	28
		2501-2750	15.5	23	28
		2751-3000	16	23.5	28
		Frequency	Typical	Typical	Typical
		5	2.3	1.3	0.6
		50	3.8	1.3	0.6
		450	3.8	1.5	0.7
		750	3.8	1.6	0.8
TYPICAL		870	3.6	1.7	0.9
Insertion Loss (dB)		1003	3.6	1.7	1
		1250	3.8	1.8	1.1
		1950	4.2	2.2	1.2
		2250	4.7	2.3	1.3
		2500	4.8	2.6	1.5
		2750	5	3	1.8
		3000	5.5	3.5	2.3

The prototype includes the following aspects:

1. The coupler that operates from 5 MHz to 3 GHz with the designed coupling coefficients
2. The pin seizure that operates from 5 MHz to 3 GHz. The insertion loss of the pin seizure is a part of the overall tap insertion loss which is specified in Table 13.
3. The AC choke (AC bypass). Its loss is also a part of the overall tap insertion loss which is specified in Table 13.
4. The DC coupler. Its loss is also a part of the overall tap insertion loss which is specified in Table 13.
5. The bypass beam.

All the components are included in the tap housing which is slightly larger than the current tap housing (see Figure 11)



**Figure 11: 3 GHz tap prototype**


One key aspect not covered in this paper is tap equalization, also referred to as signal conditioning. Signal conditioning plug-ins are inserted into taps to equalize tilt in the drop signal without impacting main feeder path. Current Plug-ins work for frequency range up to 1.2 GHz. When extending these plug-ins to 3 GHz DOCSIS, we have to consider how the signal need to be conditioned for optimal 3 GHz DOCSIS reception. Here we have to take into account multitude of factors, such as the original tilt of transmit signals, the network induced down tilt over the frequency range of interest, receiver capability in handling tilt in input RF signal. The topic of optimal transmit power allocation, which in some cases lead to an up-tilted transmit signal, is discussed later in this paper.

## **14. 3 GHz Tap Prototype Test Results**

Figure 12 to Figure 17 show the test data of the insertion losses and tap losses for tap14, tap20 and tap26. The curves labeled as JL are the original design specs, and the curves labeled as CV are the revised design specs.


In the world of taps, insertion loss is the loss created along the main cable from the input to the output of the tap. It is additive with each tap and additive to the loss of the main cable. The tap loss is the loss from the input of the tap to the tap port output that feeds the drop cable. The tap loss is by design to help create a consistent loss plan for the HFC plant where each CM is presented with a similar power level. The insertion loss is a by-product and ideally is as low as possible.






**Figure 12: Test data for Tap14 insertion loss**

In the 14 dB tap, there is up to 5.5 dB of attenuation at 3 GHz. However, at 1.2 GHz, there is 3.5 dB attenuation, so the extra attenuation going up to 3 GHz is only 2 dB.




**Figure 13: Test data for Tap14 tap loss**

The deviation of tap14 loss from the nominal 14 dB value increases with the frequency. For frequencies up to 1.2 GHz, the tap loss is within 2 dB of the nominal value of 14 dB. As the frequency reaches 3 GHz, the deviation reaches as high as 7 dB.




**Figure 14: Test data for Tap20 insertion loss**

In the 20 dB tap, there is up to 3.5 dB of attenuation at 3GHz. However, at 1.2 GHz, there is 1.8 dB attenuation, so the extra attenuation going up to 3 GHz is only 1.2 dB.




**Figure 15: Test data for Tap20 tap loss**

As in tap14, the deviation of tap20 loss from the nominal 20 dB value increases with the frequency. For frequencies up to 1.2 GHz, the tap loss is within 1.2 dB of the nominal value of 20 dB. As the frequency reaches 3 GHz, the deviation increases to 6 dB.



**Figure 16: Test data for Tap26 insertion loss**

In the 26 dB tap, there is up to 3 dB of attenuation at 3GHz. However, at 1.2 GHz, there is 1.2 dB attenuation, so the extra attenuation going up to 3 GHz is only 1.8 dB.



**Figure 17: Test data for Tap26 tap loss**

As in the previous two cases, the deviation of tap26 loss from the nominal 26 dB value increases with the frequency albeit at the lower average slope. For frequencies up to 1.2 GHz, the tap loss is within 1.2 dB of nominal value of 26 dB. As the frequency reaches 3 GHz, the deviation reaches 2 dB.

In summary, for the 3 tap cases discussed above, the extra insertion loss going from 1.2 GHz to 3 GHz is in the range of 2 dB to 1.2 dB. These losses integrate as the signal travel further down the trunk coax network, passing taps along the way, leading to additional down-tilt in the signal as it travels further away

from the point of transmission. Similarly, the deviation of tap losses from the nominal tap values generally increases with frequency, leading to more losses than intended by the tap. Additional losses here can be as high as 7 dB, but fortunately the tap losses from a particular tap only impact devices connected to that tap. In other words, these losses do not integrate along the trunk coax network. Both these loss factors and the range of losses discussed here are manageable within the 3 GHz DOCSIS framework developed here. Applying an up-tilt to transmit signal is an effective way to pre-equalize the signal to counter expected down-tilt of coax channel. This topic is covered extensively in the next section. Furthermore, the tap equalizers mentioned before is another tool available to MSOs to manage the tilt in signal spectrum in tap output going into drop network in a more localized manner.

## Power Plans


### 15. Distributed Power Amplification

When the operating frequency range is extended from 5 MHz to 1.2 GHz to 5 MHz to 3 GHz, the network needs to provide extra AC power to support the extra spectrum and related services. Just to support the existing operating frequency range (5 MHz to 1.2 GHz), the AC power grid is configured to run at its full capacity. To provide extra AC power for the spectrum between 1.2 GHz to 3 GHz will require a potentially expensive upgrade to the network power grid. The ideal case is to maintain the same AC power to avoid an expensive power grid upgrade, while extending the operating frequency range to 5 MHz to 3 GHz.

In general, when we look at the overall network power consumption, there are two factors that need be considered: the number of active devices (amplifiers) and the power consumption of each active device. The total power consumption will be the sum of the power consumptions of all devices. In one extreme case, we could place an amplifier in each tap, and the TCP of each amplifier is very low since it only needs to overcome the path loss of a very short section of coaxial network.

In this case, the total power consumption of the network will be low, but the number of the active devices is high. In the other extreme case, we place amplifiers with the same interval as the legacy network (4-5 tap interval). In this case, the number of active devices remains the same, but the power consumption is high: the amplifier needs to deliver the same RF power for the legacy 5 MHz to 1.2 GHz spectrum, and at the same time needs to deliver RF power for the 1.2 GHz to 3 GHz extended spectrum, which experiences much higher path loss. The optimal solution is somewhere in between.


The optimal power allocation scheme is nicknamed a Robin Hood power scheme: the whole DS spectrum is partitioned into the low spectrum from 5 MHz to 684 MHz, which is denoted as LS, and the extended spectrum is 804 MHz to 3 GHz, denoted as ES. 684 MHz to 804 MHz will be the cross over band (Figure 18)



**Figure 18: The spectrum partition**


Compared to the legacy spectrum (108 MHz to 1.2 GHz), the new low spectrum (108 MHz to 684 MHz) reduced the power consumption of the legacy system. The saved power will be used to support the power required for the extended spectrum while keeping the same or even lower the total power consumption of the network. LS and ES can run as two independent networks with different amp intervals.

One approach to accomplish this is with the use of a hybrid active tap (HAT) as shown in Figure 19.



**Figure 19: Hybrid Active Tap**


The HAT is constructed as follows: The incoming signal is split into LS and ES at the diplexers, and ES will go through the top branch where it is amplified. LS will go the bottom branch without amplification. The coupler coefficients, such as the EQ values, could be different for LS and ES, depending on the link budget and design. This is shown in Figure 20. To prevent feedback from occurring on the ESA within the tap, a transition band is required that will be implemented in the diplexers on either side of the ESA.



**Figure 20: HAT block diagram**

One needs to trade-off between the number of active devices and total power consumption of the network. Figure 21 presents the simulation results of total AC power consumption vs number of amplifiers for N+2 network:

- 200 HHP
- 24 HHP per tap
- Link budget:  $>-8$  dBmV/6 MHz at CM @3 GHz
- 3% power efficiency (AC $\rightarrow$ RF)



**Figure 21: Simulation results of total AC power consumption vs number of amplifiers for N+2 network**

## 16. Optimization of Transmit Power for Capacity

One of the key technology challenges in extending the DOCSIS spectrum beyond the current 1.2 GHz spectrum is the limited total composite power of transmit power amplifiers as explained previously. Silicon technology indicates that the TCP of 3 GHz PAs would need to be around the same mark as current 1.2 GHz PAs if we are to achieve the level of signal fidelity required to target higher order modulations.

For 1.2 GHz DOCSIS FDX, downstream transmit reference power spectral density (PSD) is defined in the standard [1] as having an up tilt of 21 dB, 37 dBmV/6 MHz at 108 MHz to 58 dBmV/6 MHz at 1218 MHz. The up tilt is there to counter the cable losses that monotonically increase with frequency.

Given TCP constraint mentioned above, we do not have excess headroom in PAs to allocate any power for the extended spectrum in the 1.2 GHz to 3 GHz range, let alone maintain the current level of up tilt. Hence, we need to rethink the transmit power allocation for 3 GHz ESD.

Firstly, we outline a theoretical framework to calculate the optimal power allocation given PA-related constraints and the cable plant characteristics. Following that, we show that additional constraints, such as backwards compatibility, can be incorporated in this framework to devise a power allocation strategy that is both non-disruptive to existing devices in the network and optimal for ESD devices.

### 16.1. Theoretical Framework

The capacity of HFC network for 3 GHz ESD is impacted by the propagation channel characteristics and capacity limiting factors in the transmitter and receiver. Our aim here is to optimize the transmit power distribution taking into account capacity limiting factors related to link budget. In that regard, the dominant capacity limiting factors in the DOCSIS transmitter is PA distortion. At the receiver end, additive white Gaussian noise as well as receiver distortion and noise due to analog-to-digital conversion limit capacity.

Other performance limiting factors, such as transmit and receive phase noise, are not directly linked to transmit power and considered outside the scope of this analysis. In this subsection, we outline the theoretical framework for optimizing capacity of the network subjected to the above factors and the constraint on TCP.

HFC transmission schemes such as DOCSIS 3.1 [1] use OFDM, where the channel is partitioned into  $K$  narrowband subcarriers  $k = 1, \dots, K$  with a subcarrier spacing  $\Delta f$ . Those orthogonal subcarriers 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 overall data rate  $R$  where  $k$ 'th subcarrier signal-to-noise ratio  $SNR(k)$  is given by,

$$R = \eta \Delta f \sum_{k=1}^K \min \left\{ \log_2 \left( 1 + \frac{SNR(k)}{\Gamma} \right), b_{max} \right\} \quad (1)$$

Where  $\Gamma$  is a scalar, whose value is greater than unity, representing the SNR gap to Shannon capacity for the modulation and coding scheme used. This equation also captures the impact of the limit on the number of bits per subcarrier,  $b_{max}$ . The efficiency factor,  $\eta < 1$ , captures overhead due to factors such as cyclic extension in OFDM and redundancy in forward error correction (FEC). Using  $\eta = 1$ ,  $\Gamma = 1$  and  $b_{max} \rightarrow \infty$  gives the capacity without coding and modulation limitations.

To formulate the constrained capacity optimization problem mentioned above, we expand the SNR in equation (1) and also lay out the constraints in mathematical form as follows, leading to a smooth (differentiable) form for capacity,

$$C = \max_{x(k)} \sum_{k=1}^K \log_2 \left( 1 + \frac{H(k)x(k)}{\Gamma \sigma^2(k)} \right) \quad (2)$$

where:

$H(k)$  is the channel coefficient, (attenuation, phase), for subcarrier  $k$

$\sigma^2(k)$  is the additive white Gaussian noise variance on subcarrier  $k$

subjected to:

TCP constraint:  $\sum_{k=1}^K x(k) \triangleq TCP$

Spectral mask constraint:  $0 \leq x(k) \leq p_{mask}(k)$

Spectral mask constraint captures the limitation of the modulation alphabet as follows,


$$p_{mask}(k) = \Gamma \frac{(2^{b_{max}} - 1) \sigma^2(k)}{|H(k)|^2} \quad (3)$$

Overall noise variance,  $\sigma^2(k)$  need to include both receiver contributions as well transmitter contributions. Transmitter contributions are dominated by the PA nonlinear floor, which is a function of TCP,  $p_{sum}$ . Furthermore, both receiver noise and transmit nonlinearity contribution may have a frequency dependency. Hence we can write noise power in general from as,

$$\sigma^2(k) = \sigma_n^2(k) + \sigma_t^2(k, x(1), \dots, x(K)) \quad (4)$$

where  $\sigma_n^2(k)$  represents receiver noise and  $\sigma_t^2(k, x(1), \dots, x(K))$  represent transmit distortion referred to the receiver end. PA nonlinear noise floor at the transmitter output is a function of TCP and it gets scaled by channel response before reaching the receiver.

The following figures shows simulation results for how the nonlinear distortion floor varies across frequency for given TCP and power distribution – flat and non-flat PSD, and the average MER vs TCP.



**Figure 22: PA nonlinear distortion characterisation (source: Qorvo 3 GHz PA simulation data)**



Simulation results shows that the shape of nonlinear distortion PSD is largely unaffected by the shape of transmit PSD for a given TCP. Hence, we can write  $\sigma_t^2(k)$  in following form to accurately model channel scaling and PA behavior shown above,

$$\sigma_t^2(k) = \sigma_d^2(k, TCP)|H(k)|^2 \quad (5)$$

where  $\sigma_d^2(k, TCP)$  is the nonlinear distortion in subcarrier  $k$  at the output of the transmitter. Note that  $\sigma_t^2$  here is a function of  $TCP$ , but not individual subcarrier power levels as given in (4).

For the above optimization problem, we consider  $TCP$  to be fixed. Hence the overall noise floor seen at the receiver end,  $\sigma^2(k)$ , is independent of power allocation. More generalized case of flexible  $TCP$ , where the received noise floor varies with  $TCP$  is outside the scope of this paper and will be published in a conference paper soon.

The solution for the optimization problem described in equations (2) to (5) can be shown to be a form of water filling solution [5].

$$\frac{|H(k)|^2}{\Gamma\sigma^2(k) + |H(k)|^2 x(k)} - \frac{1}{\mu} = 0 \quad \forall k, \quad 0 \leq x(k) \leq p_{mask}(k) \quad (6)$$

where the relationship between  $\mu$  and TCP is given by,

$$\frac{1}{\mu} = \frac{1}{|I_{fill}|} \left( TCP + \sum_{k \in I_{fill}} \frac{\Gamma\sigma^2(k)}{|H(k)|^2} - \sum_{k \in I_{mask}} p_{mask}(k) \right) \quad (7)$$

where  $I_{fill}$  is the set of subcarriers where water filling condition in (6) is met, and  $I_{mask}$  is set of subcarriers where the spectral mask is hit before the water filling level is reached.  $|I_{fill}|$  denotes the cardinality of set  $I_{fill}$ .


Equation (7) leads to a practical implementation of the algorithm: First calculate the noise variance per subcarrier referred to the transmitter,  $\frac{\Gamma\sigma^2(k)}{|H(k)|^2}$ . Then the transmit power is filled into subcarriers to maintain constant  $\frac{\Gamma\sigma^2(k)}{|H(k)|^2} + x(k)$  level while keeping an eye on per-subcarrier transmit power mask,  $p_{mask}(k)$ . Subcarriers that reaches  $p_{mask}(k)$  power level are stopped from receiving any more power. This filling process continues until total power allocation reaches the TCP.

## 16.2. Backwards Compatible and Optimal Power allocation

The optimal power allocation algorithm given above can be applied directly if there were no requirement for backwards compatibility. However, in DOCSIS, cable operators generally like to do upgrades while maintaining backwards compatibility to allow gradual phasing out of existing devices (CMs in particular) in the network.

We can apply the power optimization algorithm with the additional constraint of backwards compatibility. Assume we maintain the uplink in PSD up to 750 MHz to maintain backwards compatibility with legacy devices, which includes FDX DOCSIS devices. We can apply an optimal algorithm to optimize capacity over the 750 MHz to 3 GHz spectrum using the remaining power.

With the added requirement of backwards compatibility with DOCSIS-FDX and legacy DOCSIS, we end up with two options for the downstream power allocation for 3 GHz ESD as shown in Figure 23.




**Figure 23: Optimal Power Allocation for backwards compatible 3 GHz ESD**

The two options are as follows:

- Option 1: Maintain DOCSIS FDX reference PSD for the FDX band (up to 684 MHz). Then allocate the remaining power optimally.
- Option 2: Maintain DOCSIS FDX reference PSD not just for FDX band but to cover the frequency range for any legacy devices in the network (e.g., 750 MHz used here). Then allocate the remaining power optimally.

### 16.3. Results and Conclusions

An end-to-end simulation model is used to compare capacity for the N+0 Comcast Model I [12] network, as shown in Figure 24, with and without transmit power optimization.



- Trunk cables are 175 feet of type QR540
- Drop cables are 100 feet of type RG6
- 6 taps of 29 dB, 29 dB, 26 dB, 20 dB, 14 dB, 8 dB

**Figure 24: Comcast Model I**

PA nonlinearity is modelled based on Qorvo data and network taps are assumed to be upgraded to 3 GHz. Cable modem point of entry (PoE) installation, where the CM is professionally installed at the point where the drop cable enters the customer premises, and deep home run (additional 100 feet cable from PoE) self-install scenario is shown in **Figure 25**. The spectrum below 1.2 GHz is assumed to be as defined in DOCSIS 4.0 FDX. The spectrum beyond 1.2 GHz is considered all downstream. **Figure 25** shows downstream and aggregate (downstream + upstream) capacities of the network.




Figure 25: Capacity with and w/o transmit power optimization

In summary, optimal power allocation follows the water filling principle with added complexity having to deal with sum of all noise and distortion sources. The solution can be found iteratively. Generally, the optimal solution has an up-tilted spectrum at lower frequencies and flat power allocation in higher end frequencies. Power optimization improves capacity by 3-5% for PoE and up to 8% for deep home run (home wiring of 100' RG6).

## DOCSIS PHY Optimizations

### 17. Introduction

OFDM with a conventional cyclic prefix (CP) [6] is an elegant multicarrier modulation scheme, which offers all the advantages associated with multicarrier systems, such as SNR vs frequency dependent bit loading, frequency domain one tap equalization, for a small overhead of cyclic prefix. There are other multicarrier options found in the literature that reduce or eliminate the cyclic prefix [7], but these come with added complexity and hardware resources, such as memory. In a nutshell, we can summarize reasons why we should continue with OFDM for 3 GHz ESD as follows,

- With time and frequency interleaving in place, OFDM is very robust to both burst and ingress interference
- Simplified transceiver architecture: e.g., efficient modulation/demodulation with FFTs, one tap equalization
- Robust against multipath - easily dealt with guard interval (cyclic prefix)

- d) Sensitivity to phase noise and time variations more than single carrier modulation, but these can be mitigated by well-known algorithms – CPE correction, adaptive channel estimation.
- e) With narrow subcarriers (e.g., 50 kHz), bit loading can be used to optimize throughput with fine frequency resolution
- f) Classic cyclic prefix OFDM as in DOCSIS 3.1/FDX is recommended. We can reduce CP overhead for ESD. More complex multicarrier schemes lose some of the above advantages – b, c


In this section, we show that DOCSIS 3.1 OFDM parameters are well suited for 3 GHz ESD with potentially reduced cyclic prefix overhead for the extended part of the spectrum. We also explore potential extensions to DOCSIS 3.1 OFDM to support 3 GHz.

## 18. A Systematic Approach to Selecting OFDM Parameters

The two key parameters for OFDM are the useful symbol duration,  $T_u$ , and the cyclic prefix,  $T_{cp}$ . Once these are determined, the number of subcarriers in an OFDM channel are defined based on individual OFDM channel bandwidths required to give per-channel throughput. Channel bandwidth is more of a high-level PHY decision to be made based on how you'd like to organize the DOCSIS spectrum into channels. This could also have hardware implications that need to be taken into account.

### 18.1. Target MER

When looking for suitable OFDM parameters, target MER is the key metric of consideration. The upper bound for MER for 3 GHz ESD can be worked out from a link budget analysis of the representative ESD systems. Figure 26 shows MER results of a full system (transmitter, network, and receiver) simulation. The MER values achieved for CMs connected to different taps in the Comcast Model I are shown here.



**Figure 26: Optimized SNR (solid) and SNR achieved with flat transmit PSD (dashed)**

At lower frequencies, we are targeting MER of 37.5 dB and at higher frequencies 30 dB looks like a realistic target (at least for some CMs). With many such system analysis, we can determine a target MER mask for 3 GHz ESD. In the following sections, we assume we are targeting a 35 dB MER in the ESD spectrum.

## 18.2. Cable Propagation Channel

DOCSIS 3.1 [1] defines individual micro-reflection masks for the downstream channel shown in Table 14.

**Table 14: DOCSIS 3.1 DS Micro-Reflection Bound (mask)**

Echo Delay	dBc Level
$\leq 0.5 \mu\text{s}$	-20
$\leq 1.0 \mu\text{s}$	-25
$\leq 1.5 \mu\text{s}$	-30
$> 2.0 \mu\text{s}$	-35
$> 3.0 \mu\text{s}$	-40
$> 4.0 \mu\text{s}$	-45
$> 5.0 \mu\text{s}$	-50

This channel spec is based on return losses and propagation losses in passive elements, including the cable itself, in the network. Surveying return losses for various passives in network, 10-12 dB seems reasonable worst case we can expect. There are at least two reflections for each echo in the forward channel. Hence, -20 dB for very short echoes, as given in Table 14, is a reasonable mask. Even higher losses for longer echoes seen in Table 14 can be attributed to cable losses and other insertion losses in the cable channel.

There are two broad types of cables in network - drop cables and hardline distribution cables, that need to be considered when calculating the cable attenuation in the network. We used an extensive collection of cable S-parameter data from CableLabs [11] to create a summary of cable loss vs frequency for a 0.5  $\mu\text{s}$  individual echo given in Table 15.

**Table 15: Coax Cable Losses for 0.5  $\mu\text{s}$  echo**

	100 MHz	1 GHz	2 GHz	3 GHz
<b>Drop</b>	6.5 dB	20 dB	30 dB	36.7 dB
<b>Distribution</b>	3 dB	7.6 dB	10.4 dB	13.3 dB

Note that for each increment of 0.5  $\mu\text{s}$  delay, the echo level in Table 14 decreased by 5 dBc for echoes up to 2  $\mu\text{s}$ . This matches with the sum loss of distribution cable at 100 MHz, 3 dB (Table 15), and an additional 2 dB insertion loss due to other passive elements, such as taps, along the way.

Based on the above observations, we make the following very conservative estimate of loss per 0.5  $\mu\text{s}$  increment in micro-reflection delay at 2 GHz and 3 GHz.

- At 2 GHz, distribution cable loss ~10 dB + 2 dB other insertion losses for other passive giving total of = 12 dB loss/0.5  $\mu\text{s}$
- At 3 GHz, ~13 dB (trunk) + 2 dB (other) = 15 dB

This leads to the following modified micro-reflection mask at 2 GHz and 3 GHz.

**Table 16: Coax Channel Model at 2 GHz and 3 GHz**

Echo Delay	DOCSIS 3.1 Downstream	2 GHz	3 GHz
$\leq 0.5 \mu\text{s}$	-20 dBc	-20 dBc	-20 dBc
$\leq 1.0 \mu\text{s}$	-25 dBc	-32 dBc	-35 dBc
$\leq 1.5 \mu\text{s}$	-30 dBc	-44 dBc	-50 dBc
$> 2.0 \mu\text{s}$	-35 dBc	-56 dBc	-65 dBc
$> 3.0 \mu\text{s}$	-40 dBc		
$> 4.0 \mu\text{s}$	-45 dBc		
$> 5.0 \mu\text{s}$	-50 dBc		

Using distribution cable data for the micro-reflection mask (upper bound) here is well justified as this gives the worst-case echoes, leading to a reliable if not pessimistic micro-reflection mask. The calculations here need to be validated using real HFC network measurements as part of network characterization for 3GHz DOCSIS.

This indicates channel impulse response delay spread for significant echoes, i.e., stronger than (Target\_MER+10 dB), for signal in extended spectrum could be significantly smaller compared to lower frequencies.

### 18.3. Cyclic Prefix for OFDM

The OFDM cyclic prefix length (guard interval) needs to be long enough to prevent any significant performance degradation due to micro-reflections. To achieve this objective, the delay spread for all significant echoes needs to be smaller than the cyclic prefix with enough margin for transmitter and receiver windowing. The time margin needed for windowing can be significantly reduced or eliminated altogether for extended spectrum as we are dealing with a clean part of spectrum with no legacy channels. For example, by forcing all OFDM channels in extended spectrum to have the same OFDM parameters and synchronizing their timing and frequency, we can eliminate any leakage between OFDM channels without a need for any windowing. For the following analysis we assume there is no windowing.

Given a target MER of 35 dB for extended spectrum, we can aim for a target carrier-to-interference ratio (CIR), which is based on the sum of inter-symbol interference (ISI) and inter-carrier interference (ICI) level, of 50 dB for individual echoes. The idea is to keep aggregate ISI + ICI impact of all echoes outside of the guard interval to be below -45 dBc (i.e., only 0.4 dB impact on 35 dB MER point).

The micro-reflection masks given in Table 16 shows echoes above 2  $\mu\text{s}$  are not relevant for added ESD spectrum as echo amplitudes are well below the -50 dBc threshold. We can restrict this even further because only the portion of echo outside the guard interval contribute to ISI and ICI. More precisely, for an echo longer than the guard interval, the ICI+ISI contribution is given by,

$$ICI\_ISI = A_{dB} + 3 + 10 * \log_{10}((\tau - T_g)/T_u) \quad (7)$$


Where  $A_{dB}$  is the echo power,  $\tau$  echo delay,  $T_g$  is guard interval and  $T_u$  is OFDM useful symbol duration.

For example, a -40 dBc echo that is 1  $\mu\text{s}$  longer than the cyclic prefix in the 20  $\mu\text{s}$  OFDM case causes ISI + ICI = -50 dBc.

We have two downstream OFDM modes within the DOCSIS 3.1 192 MHz channel,

- 8K mode: Symbol length 40  $\mu\text{s}$
- 4K mode: Symbol length 20  $\mu\text{s}$

Furthermore, we have five cyclic prefix lengths, 5  $\mu\text{s}$ , 3.75  $\mu\text{s}$ , 2.5  $\mu\text{s}$ , 1.25  $\mu\text{s}$ , and shortest of which is 0.9375  $\mu\text{s}$  (192 samples at 204.8 MHz). Given the significant reduction in echo levels in extended spectrum, it could be beneficial to introduce shorter CP lengths for 3 GHz ESD. Potential efficiency gains from a reduced cyclic prefix are shown in Figure 16 for the current 4K and 8K OFDM modes and a hypothetical (potential new) 2K OFDM mode.



**Figure 27: Cyclic Prefix Overhead**

The shortest cyclic prefix in DOCSIS 3.1 is 0.9375  $\mu\text{s}$ . We should consider introducing shorter CP lengths to improve efficiency for ESD. However, a new 2K mode is not a good idea from an efficiency point of view.

#### 18.4. OFDM Symbol Length

Recall that we have two OFDM modes for a DOCSIS 3.1 channel, giving two OFDM symbol length options:

- 8K mode: Symbol length 40  $\mu\text{s}$  -> subcarrier spacing of 25 kHz
- 4K mode: Symbol length 20  $\mu\text{s}$  -> subcarrier spacing of 50 kHz

Factors we need to consider in optimizing OFDM symbol length

- **Latency:** Longer symbols gives rise to longer latency. With new applications such as 5G front haul, we may want to optimize latency as much as possible. However, the lion's share of latency comes not from PHY, but due to scheduling related delays (buffering, etc.). We'd like to keep symbol length shorter, but the current choices of symbol lengths are good enough for this purpose.
- **Time variations in overall channel:** Variations here include outside plant time variations due to factors such as temperature as well as variations internal to transceivers, such as phase noise, AGC variations, etc. We'd like to keep symbol lengths shorter. This is analyzed in the following sections.
- **Efficiency:** We need to make OFDM symbols as long as possible to reduce cyclic prefix overhead. As discussed previously, current OFDM symbols lengths offer acceptable efficiency with the required cyclic prefix lengths. There is further room for efficiency gains in 3 GHz ESD by introducing shorter cyclic prefix options.
- **Frequency resolution for bit loading:** Bit loading frequency resolution should be good enough to follow the frequency dependent SNR of a DOCSIS channel due to amplitude tilt and frequency selectivity of multipath channel. This is likely to be in the order of 100's of kHz if not more. Current frequency resolution of 50 kHz and 25 kHz offered by 4K and 8K OFDMs are more than enough to enable optimal bit loading to closely fit the SNR profile across ESD Spectrum.

## 18.5. OFDM Time variations

Based on limited studies in the FDX workgroup, the outside plant variations occur in time scale of seconds. These are too slow to impact decision on OFDM symbol length.

As for variations inside transceivers, phase noise is a major factor. We need phase noise variations within an OFDM symbol to be contained. This is because any variations within the OFDM symbol leads to inter-carrier-interference. It is desirable to keep ICI due to phase noise 10-15 dB below the MER target to limit the impact on performance.

The DOCSIS 3.1 specification [1] mandates that the CMTS adheres to the following clock jitter requirements for the downstream OFDM symbol clock jitter mask over the specified frequency ranges as shown in Table 17:

**Table 17: CMTS RMS Jitter Spec**

Frequency Range	RMS Jitter	Equivalent phase noise referred to $f_{DS}$ - SSBError! Bookmark not defined. <sup>1</sup>
10 Hz to 100 Hz	< 0.07 ns	$-21+20*\log(f_{DS}/204.8)$ dBc
100 Hz to 1 kHz	< 0.07 ns	$-21+20*\log(f_{DS}/204.8)$ dBc
1 kHz to 10 kHz	< 0.07 ns	$-21+20*\log(f_{DS}/204.8)$ dBc
10 kHz to 100 kHz	< 0.5 ns	$-4+20*\log(f_{DS}/204.8)$ dBc
100 kHz to $(f_{DS}/2)$ ,	< 1 ns	$2+20*\log(f_{DS}/204.8)$ dBc

<sup>1</sup> Equivalent phase noise =  $20\log_{10}(RMS\_Jitter \times 2\pi f_{DS})$  dBc



In addition to meeting the above clock jitter requirements, the CMTS is required to meet the following phase noise requirements [1] as shown in Table 18.

**Table 18: CMTS Phase Noise mask at 1002 MHz**


Frequency Range	Integrated Phase Noise (@1002 MHz) -SSB
1 kHz - 10 kHz	-48 dBc
10 kHz - 100 kHz	-56 dBc
100 kHz - 1 MHz	-60 dBc
1 MHz - 10 MHz	-54 dBc
10 MHz - 100 MHz	-60 dBc

In the event of a conflict between the clock jitter and the phase noise requirement, the CMTS MUST meet the more stringent requirement [1]. Based on above rule, we can define a combined time jitter phase noise PSD for the CMTS @1218 MHz. This is shown in Table 19.

**Table 19: Combined Phase Noise Mask for CMTS**

Frequency Range	Integrated Phase Noise (@1002 MHz) - SSB
10 Hz to 100 Hz	-7.2 dBc
100 Hz to 1 kHz	-7.2 dBc
1 kHz to 10 kHz	-48 dBc
10 kHz to 100 kHz	-56 dBc
100 kHz to 1 MHz	-60 dBc
1 MHz to 10 MHz	-54 dBc
10 MHz to 100 MHz	-60 dBc

Given double sideband (DSB phase noise PSD  $\Phi(f)$ ), the integrated phase noise in  $f_1 \text{ Hz} - f_2 \text{ Hz}$  range is given by  $\int_{f_1}^{f_2} 2\Phi(f) df$ . CMTS phase noise below 1 kHz is perhaps too relaxed. In Figure 28, let's look at the DOCSIS 3.1 phase noise profile with the phase noise below 1 kHz represented with -10 dB/decade tilt.



Int\_Pha\_Power (10 Hz to 100 Hz) -Inf dB  
 Int\_Pha\_Power (100 Hz to 1 kHz) -33.777828 dB  
 Int\_Pha\_Power (1 kHz to 10 kHz) -44.668151 dB  
 Int\_Pha\_Power (10 kHz to 100 kHz) -54.262065 dB  
 Int\_Pha\_Power (100 kHz to 1 MHz) -58.151684 dB  
 Int\_Pha\_Power (1 MHz to 10 MHz) -52.337972 dB  
 Int\_Pha\_Power (10 MHz to 100 MHz) -58.071943 dB

ICI\_mid = 49.693168 dB, ICI\_edge = 52.682889 dB  
 CPE power = 6.031562e-04 radians radians, RMS C  
 ICI power <= 100 Hz = 90.112047 dB  
 ICI power <= kHz = 72.709654 dB  
 ICI power kHz - 10kHz = 68.517150 dB  
 ICI power 10kHz - 100kHz = 61.384878 dB  
 ICI power 100kHz - MHz = 61.167415 dB  
 ICI power MHz - 10MHz = 55.348498 dB  
 ICI power 10MHz - 100MHz = 61.082261 dB  
 ICI power 1kHz - 100MHz = 52.732882 dB

**Figure 28: CMTS Phase Noise PSD – DSB, referred to 1218 MHz**

The impact of phase noise on OFDM can be worked out using the method described in [5]. RMS common phase error (CPE) for OFDM symbol with subcarrier spacing  $f_u = 1/T_u$  and number of subcarriers  $N$  is given by,

$$RMS\_CPE = \sqrt{\int_0^{Nf_u/2} 2 \text{sinc}^2\left(\frac{f}{f_u}\right) \Phi(f) df} \quad (8)$$

CPE introduces the same phase rotation to all subcarriers. Digital demodulation can correct for CPE. This can be done accurately in DOCSIS with available continuous pilots. For example, using the minimum of 8 boosted power (by 6 dB) continuous pilots, we can estimate CPE with an accuracy of  $6+10*\log_{10}(8) = 15$  dB below the noise floor.

Inter-carrier interference to signal ratio (ICI/S) for edge and middle subcarrier of OFDM symbol is given in below [5].

<b>Mid Subcarrier</b>	<b>Edge Subcarrier</b>
$\int_0^{Nf_u/2} 2 \left(1 - \text{sinc}^2\left(\frac{f}{f_u}\right)\right) \Phi(f) df$	$\int_0^{Nf_u} \left(1 - \text{sinc}^2\left(\frac{f}{f_u}\right)\right) \Phi(f) df \quad (9)$

ICI can also be corrected to a certain degree [9][10] with advanced signal processing techniques. However, it is desirable to keep the ICI level well below quasi error free (QEF) noise floor, if possible, to keep demodulator complexity in check.

ICI due to CMTS phase noise profile analyzed here is given in Table 20.

**Table 20: Subcarrier Power to ICI ratio in dB due to CMTS Phase Noise from Figure 28**

	25 kHz Subcarrier (8K mode)		50 kHz Subcarrier (4K mode)		100 kHz Subcarrier (2K mode)	
Upper Edge	Middle Carrier	Edge Carrier	Middle Carrier	Edge Carrier	Middle Carrier	Edge Carrier
<b>500</b>	56.5	59.5	57.4	60.4	57.9	60.9
<b>1218</b>	48.8	51.8	49.7	52.7	50.2	53.2
<b>2000</b>	44.5	47.5	45.5	48.4	45.9	48.9
<b>3000</b>	41.0	44.0	41.9	44.9	42.4	45.4

ICI for 8K mode is only 0.9 dB worse than 4K mode because of broad phase noise PSD compared to subcarrier spacing. The hypothetical 2K mode is only 0.5 dB better.

CMTS ICI numbers are not good enough for us to have 15 dB margin over target MER (40 dB at lower frequencies to 30 dB at higher frequencies) given in Figure 26. We need about 5 dB improvement. In practice, the phase noise profiles are better than what is given in the specifications and in this regard, the DOCSIS PHY specification needs updating.

Going for a new shorter symbol length (2K OFDM mode) to ease the phase noise spec is not justifiable as it does not give enough improvement in ICI levels and ends up costing in efficiency as explained in the cyclic prefix discussion. Non-performance related reasons, such as to reduce number of subcarriers to lower memory requirements, etc., could be considered in the spec process.

## 19. Conclusion

In conclusion, DOCSIS 3.1 OFDM symbol lengths, 20  $\mu$ s and 40  $\mu$ s, look like good candidates for 3 GHz ESD. New shorter cyclic prefix options could be considered to further improve efficiency. We could also reduce RX and TX window length by using synchronized OFDM channels – i.e., no ICI/leakage between channels. The impact of potential external interferers also needs to be considered. This potentially opens possibility for very small CP lengths leading to improved OFDM efficiency.

Node and CM phase noise specs need to be reassessed for 3 GHz ESD and improve in line with what is possible with technology.


OFDM channel bandwidth choice has little to do with PHY performance. We can get to any reasonable bandwidth by changing the number of subcarriers, for instance, double the number of subcarriers to get to twice the existing DOCSIS 3.1 maximum channel bandwidth, to 384 MHz. Unless there is a good argument for doing this, we recommend keeping the maximum channel bandwidth as it is now in DOCSIS 3.1: 192MHz.

# Power Amplifier Circuit Design Development

## 20. Output Stage Gain Blocks for HFC Amplifiers and Nodes

The performance and specifically the linear output power of the output stage gain blocks, the so-called power doublers (PD) inside cable amplifiers and nodes for HFC networks, have a substantial impact on the design of such system architectures. Therefore, to define new network systems, it is required to understand the linear output capability of cable amplifiers and nodes by characterizing the output stage gain blocks for their specific capabilities under the required loadings. On top of this, it needs to be considered that there is about 2 dB to 3.5 dB of loss between the output stage gain block and the amplifier or node housing output.

In the past, new semiconductor technologies enabled new generations of systems with higher bandwidth and higher output capability. Figure 29 shows the efficiency of the output stage gain block as it developed and increased over time while introducing new semiconductor technologies and circuit designs. For example, in the 1990s the introduction of gallium arsenide (GaAs) semiconductor technology to replace silicon (Si) bipolar transistors significantly increased the linear output power of the output stage gain block and therefore the efficiency of the power amplifier. About 10 years later, another semiconductor technology enabled even better efficiencies and higher linear output power with the combination of GaAs and gallium nitride (GaN) process technology in one gain block.



**Figure 29: Efficiency Development of Output Stage Gain Blocks in so-called Power Doublers**

Since then, this combination of semiconductor technologies was further developed and optimized to provide higher bandwidth with the expansion to 1.2 GHz systems and additional higher linear output power.

50 MHz to 1.2 GHz gain blocks are state of the art today. There is some hardware available for systems up to 3 GHz or even higher, but this is limited to single-ended devices that are intended to be used as

To start, defining a new 3 GHz HFC network systems requires a detailed perform system simulation. To support this, output stage gain block models had to be developed that can be used to perform multi-carrier linearity simulations that then can be used to understand the limitations of the power amplifiers in 3 GHz systems.




Figure 31 shows a simplified schematic of a balanced design in cascode configuration employing GaAs pseudomorphic high electron mobility transistor (pHEMT) technology for the cascode bottom device (see field effect transistor 1 (FET1) and FET2 in Figure 31) and GaN high electron mobility transistor (HEMT) semiconductor technology for the cascode top device (see FET3 and FET4 in Figure 31).

The ability to combine various semiconductor technologies into one gain block offers the possibility to select the technology that provides the best properties for each stage in the gain block. GaAs pHEMT provides high transconductance and high  $f_T$  to enable high gain and a wide bandwidth for the amplifier.

finally defines the linear output power of the gain block. Baluns and transformers (TF) (see TF1 and TF2 in Figure 31) are used to convert the balanced circuit to single ended input and output ports and to match the transistor circuit to 75 ohms.

## 21. 3 GHz Output Stage Gain Block Simulation Model

To extend the bandwidth from today's standard 1.2 GHz to 3 GHz, it is essential to investigate the RF properties and performance of the employed semiconductor technology. Specifically, for the GaN based top devices of the cascode, various available process technologies were modeled and characterized to

understand the capabilities and limitations for this application. Historically, GaN based gain blocks used in cable amplifiers primarily employed 0.5  $\mu\text{m}$  or 0.25  $\mu\text{m}$  gate length GaN processes. To accommodate the higher frequencies requirements for 3 GHz applications, Qorvo's GaN process GAN15 was selected for the model. The GAN15 GaN HEMT process comes with a 0.15  $\mu\text{m}$  gate length enabling high frequency applications even in the mm-wave frequency range today.

To derive an exact gain block model, special non-linear models for the GaN stage were generated by taking load pull data from actual GAN15 FET devices at the specific bias and frequency conditions applying for this application. This non-linear model was then used to develop and optimize the output stage gain block model.

The developed model of the gain block provides a bandwidth of 50 MHz to 3 GHz, with 19 dB gain at 100 MHz and 21 dB gain at 3 GHz. The bias condition was selected to be  $V_+ = 32\text{ V}$  and  $I_{DC} = 560\text{ mA}$  to achieve the highest linear output power with this configuration.

To simulate the linear performance of gain blocks today, the simulation benches with two-tone tests like second order intercept point (IP2) or third order intercept point (IP3) are used in most cases. These simulation results can be used to optimize the gain block for linearity, and measurements results can be correlated with the simulation results once the amplifier is taped out, processed and finally measured. However, the two-tone tests provide only an indication about the broadband linearity performance under multi-carrier loadings. There is no strong correlation between two-tone tests and multi-carrier distortion tests.

For a 3 GHz network system, it is not sufficient to perform two-tone tests only for the simulations. Therefore, it was required to develop a simulation bench that can apply a multi-carrier or broadband input signal to the gain block model and to analyze the output of the gain block for distortion products generated by the gain block.

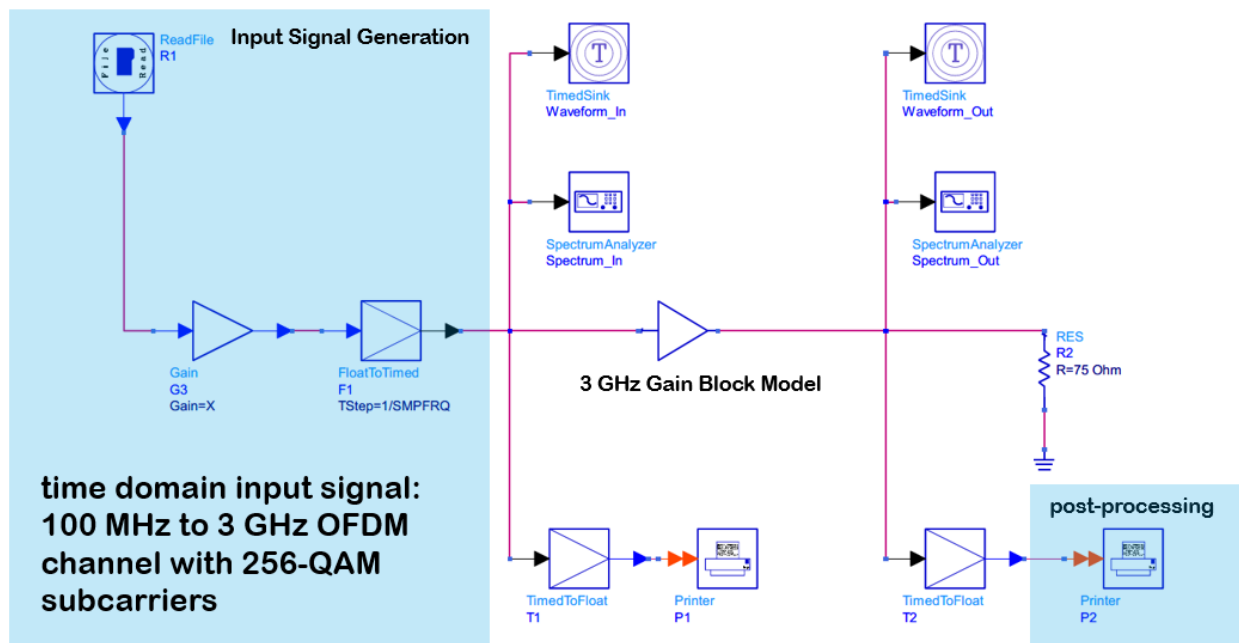
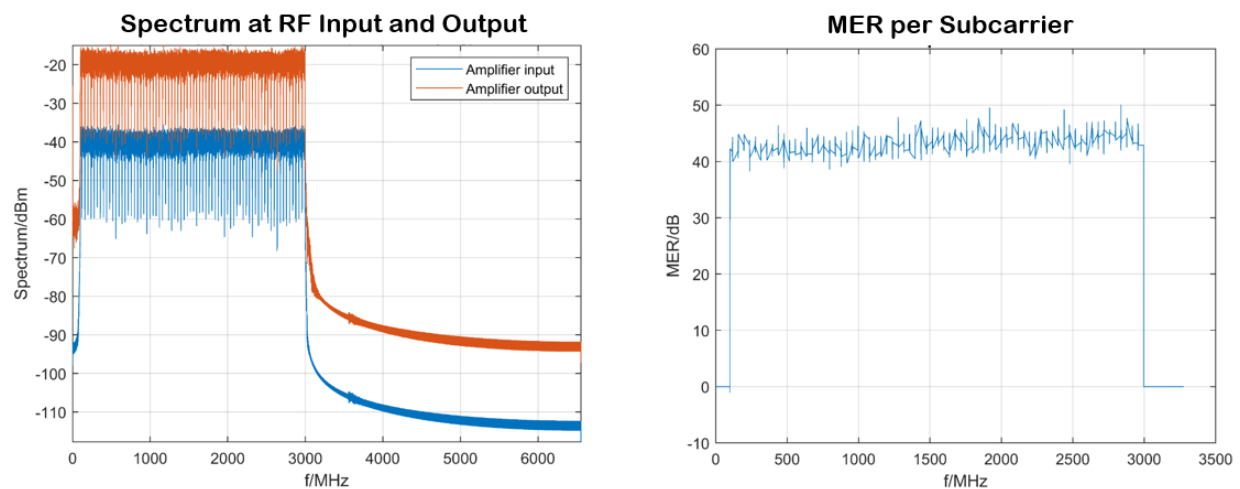


Figure 31: Multi-carrier distortion simulation test bench


Figure 31 shows the simulation bench that was used to simulate the distortion performance and the linear output capability of the developed output stage gain block model. For the input signal, a 100 MHz to 3 GHz OFDM channel with 256 state quadrature amplitude modulation (256-QAM) subcarriers was selected that is also used for 3 GHz network system simulations. This time domain input signal is converted into the RF domain and applied to the gain block model to perform the non-linear large signal simulation of the gain block. Input and output signals are monitored in the time and RF domain for external post-processing to derive MER data over frequency and level.

It is specifically challenging to make sure such simulations converge and provide results when applying a wideband large signal to a gain block model. It required several iterations of semiconductor models and circuits design model updates to finally successfully complete the non-linear simulations.



**Figure 32: Simulation RF input and output spectrum and the MER per subcarrier over frequency**

Today, the exact shape of the signals, levels and split between upstream and downstream is not yet specified for future 3 GHz network systems. The intention of this simulation task, however, is to derive an idea about the linear output power or (TCP) that can be achieved with currently available semiconductor technology. Therefore, a full loading between 50 MHz and 3 GHz with no tilt was selected, applied to the gain block model and characterized for MER over frequency as shown in Figure 32.



**Figure 33: Simulated compression characteristic of averaged MER versus TCP**

This is the simulation results for a specific input and therefore the output level (TCP = 57.6 dBmV RF input and TCP = 78.2 dbmV RF output). The output signal derived from the simulation was post-processed and resulted in an average MER of 45 dB. Additionally, the simulations showed that the MER versus frequency signature basically follow the input and output signal shape.

To understand the gain block characteristic at various loading levels, the previously described input signal was swept and an MER compression curve was simulated. Figure 34 shows the averaged MER over the output level (TCP). The gain block model provides a linear degradation of MER of about 1 dB per 1 dB increase in level up to a TCP of 77 dBmV. Above 78 dBmV, the gain block is compressing with more than 2 dB MER degradation per 1 dB increase in level. Also, MER drops below 40 dB MER above 80 dBmV TCP.

This linear output power capability basically matches the performance that can be measured on state of the art 1.2 GHz output stage gain blocks under DOCSIS 3.1 loadings. The compression characteristic results derived from the simulation were used for further HFC network system simulation to design and define the future system architecture.

## 22. Conclusion

A 3 GHz output stage gain block model was developed and characterized for averaged MER versus output TCP. Similar linear output power was achieved as with measurements on state of the art 1.2 GHz output stage gain blocks. Therefore, the performed simulation tasks proved that semiconductor technologies available today are capable of supporting 3 GHz gain block developments for 3 GHz HFC applications with respect to bandwidth and linearity. Future investigations have to characterize the performance of the gain blocks when the exact shape of the loading and the split between the upstream and downstream is defined.



## Additional Deployment Considerations

This technology will not be available all at once. This is a likely 1.8 GHz upgrade phase that will happen prior to the 3 GHz upgrade phase. In addition, the upgrade for any increased spectrum impact the both the HFC and the DOCSIS equipment, so a joint upgrade has to occur. Specifically, an upgrade for 1.8 GHz and/or 3 GHz impacts:

1. The CM
2. Passive splitters in the home network and HFC plant
3. Taps
4. Amps
5. Nodes
6. RPD
7. CMTS Core

That is a lot of coordination required. Some practical interim steps could help speed time to market. Most deployed taps in the field will not get to 1.8 GHz due to the pin seizure design. This white paper looked at a redesign of the pin seizure that allowed that tap to extend its performance to 3 GHz. As such, all tap upgrades should be 3 GHz to allow for future planning, even if the near-term plant usage is 1.8 GHz.

If or when the node/amp/line extender housing are upgraded for extended spectrum, they should be 3 GHz capable even though the initial electronics may only be 1.8 GHz capable. The CM could be designed with a 3 GHz front end, but only have enough OFDM channels to support 10 Gbps. CMs could be frequency stacked.

# Summary

This white paper discussed in detail how to move from the 1.2 GHz systems of today that are capable of 10 Gbps in the downstream to a 3 GHz system that is capable of 25 Gbps in the downstream. This would match the speed of fiber, either competitively, or fiber that is used in a DAA architecture to backhaul a DAA node.

The paper discussed how extended spectrum required either more power, less modulation, or less distance between actives. The proposal in this paper was for less distance between actives but only for the extended spectrum. This was achieved by putting one or two 3 GHz extended spectrum amplifiers (ESA) between existing amps. If the ESA was co-located with a tap, that would be a hybrid active tap (HAT). Between the legacy spectrum and the extended spectrum, there would be an extended spectrum transition band (ETB) for the diplexers in the HAT.

The power that is used to operate the extended spectrum is less than the legacy spectrum as there are more amps with less power per amp. If the transition band is moved down, the power to run the legacy spectrum is also reduced. It may be possible to take enough power from the legacy band and use it to run the extended band without increasing the overall plant power. This was referred to as the Robin Hood scheme.

There are many different spectrum plans that could be used. This paper looked at four 3 GHz plans

1. *Premium plan* with 1218 MHz ETB with extended FDX (not defined yet). This could support 25 Gbps downstream and 10 Gbps upstream.
2. *MoCA plan* with 1100 MHz ETB. The ETB is chosen to align with the MOCA band.
3. *Legacy plan* with 1002/862/750 ETB. The ETB would line up with previous older generations of plant.
4. *Lowest power plan* with a 684 MHz cross-over. This is the maximum Robin Hood scheme.

This was compared to DOCSIS 3.1 at 1.218 GHz and DOCSIS 4.0 at 1.8 GHz which are both 10 Gbps downstream systems.

Amplifier tilt is another design consideration. Tilt is needed to match the increase attenuation at high frequencies. If the amplifier components in the node, amps, and line extenders are full spectrum, then they will need a common tilt. However, if the RPD feeding the node came out at different flat power levels at different frequencies, then less power could be put into the extended band where the ESA exists. If there are separate amps for legacy and extended spectrum, then the tilt values could be different. The final specification may have separate power rules for below 1.218 GHz, 1.218 GHz to 1.8 GHz and 1.8 GHz to 3.0 GHz.

The PHY may require some tweaking to get to 3 GHz. There may be an impact to the cyclic prefix for efficiency. Phase noise is much harder at 3 GHz than at 1.2 GHz by a factor of 9. Finally, a complete silicon simulation was performed to prove that a 3 GHz amplifier component could be built.

## Acknowledgements

The authors would like to thank Shaul Shulman and Rainer Strobel from Intel and Nguyenvu Chu from Qorvo for their contributions to this project. We would also like to thank Ron Hranac for his excellent review of our paper.

# Abbreviations

bps	bits per second
CIR	carrier-to-interference ratio
CM	cable modem
CP	cyclic prefix
CPE	common phase error
DOCSIS	Data-Over-Cable Service Interface Specifications
DPA	distributed power amplification
DS	downstream
DSB	double sideband
ESA	extended spectrum amplifier
ESD	extended spectrum DOCSIS
ETB	extended spectrum transition band
FDX	full-duplex
FEC	forward error correction
FET	field effect transistor
$f_t$	transition frequency
FTB	full duplex transition band
GaAs	gallium arsenide
GaN	gallium nitride
HAT	hybrid active tap
HEMT	high electron mobility transistor
HFC	hybrid fiber/coax
Hz	hertz
dBmV	decibel millivolt
ICI	inter-carrier interference
ICI/S	inter-carrier interference to signal ratio
IDC	DC current
IP2	second order intercept point
IP3	third order intercept point
ISBE	International Society of Broadband Experts
ISI	inter-symbol interference
OFDM	orthogonal frequency division multiplexing
PA	power amplifier
PD	power doubler
pHEMT	pseudomorphic high electron mobility transistor
PoE	point of entry
PSD	power spectral density
MER	modulation error ratio
QAM	quadrature amplitude modulation
QEF	quasi error free
RGB	reduced guard band
SCTE	Society of Cable Telecommunications Engineers
Si	silicon
TCP	total composite power
TF	transformer
US	upstream

# Bibliography & References

- [1] *MoCA 2.0 Specification for Device RF Characteristics*, MoCA Alliance, April 6, 2015
- [2] *SCTE 235 2017: Operational Practice for the Coexistence of DOCSIS 3.1 Signals and MoCA Signals in the Home Environment*, SCTE, 2017
- [3] *MoCA 2.0/2.5 Specification for Device RF Characteristics*, MoCA Alliance, August 8, 2016
- [4] *Data-Over-Cable Service Interface Specifications DOCSIS® 4.0*, CableLabs, Physical Layer Specification, CM-SP-PHYv4.0-D01-190628
- [5] Rainer Strobel, *Channel Modeling and Physical Layer Optimization in Copper Line Networks*, Springer, 2019
- [6] Bingham, A. C., *Multicarrier modulation for data transmission: an idea whose time has come*, IEEE Communications magazine, May 1990
- [7] Sahin, A., et al. *A Survey on multicarrier communications: prototype filters, lattice structures, and implementation aspects*, July 2013
- [8] Stott, J., *The effect of phase noise in COFDM*, BBC Research and Development
- [9] Hewavithana Thushara, et al., *Method and apparatus for phase noise mitigation*, US Patent 8897412
- [10] Hewavithana Thushara, et al., *Computationally Efficient Algorithm for Mitigating Phase Noise in OFDM Receivers*, US Patent 10171272
- [11] *HFC network measurements*, CableLabs, 2017
- [12] *FDX Channel Models From Fiber Deep Designs*, Comcast, 2017

# **R-PHY with Remote Upstream Scheduler**

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

170 W Tasman Dr, San Jose, CA 92677

408-526-7651

jchapman@cisco.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Content .....	5
1. R-PHY US Scheduler Location Options.....	5
1.1. The Case for a Centralized US Scheduler.....	6
1.2. The Case for a Remote US Scheduler .....	8
1.2.1. REQ-GNT Protocol Tightening .....	8
1.2.2. Long Distance R-PHY Deployment.....	9
1.2.3. CCAP Realtime Performance Acceleration.....	9
1.2.4. Easy to Add and Flexible to Adapt.....	10
2. Remote US Scheduling Services.....	10
2.1. Remote REQ-GNT Service.....	12
2.2. Remote MAP Builder Service.....	13
2.3. Remote MAP and UCD Replication Service.....	13
3. R-UEPI in Remote US Scheduling.....	14
4. R-DEPI in Remote US Scheduling.....	15
5. Remote US Scheduling APIs.....	16
5.1. Top-Level RPD Operational Configuration Objects.....	17
5.2. Remote Upstream Scheduler Data Object Tree .....	18
5.2.1. MAC Domain REQ-GNT Scheduler .....	19
5.2.2. MAC Domain MAP Builder .....	20
5.2.3. MAC Domain MAP-UCD Replicator.....	22
6. Locality Optimization with Remote US Scheduling.....	23
Conclusion .....	24
Acknowledgements .....	24
Abbreviations.....	25
Bibliography & References .....	25

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - DAA Options Including R-PHY with Remote US Scheduler .....	5
Figure 2 - Centralized vs. Remote US Scheduler Deployment Scenarios. ....	6
Figure 3 - REQ-GNT Delay Elements and the Impact of the MAP Interval and CIN Distance.....	7
Figure 4 - CIN Delay Impact on the REQ-GNT Latency at Short MAP Intervals. ....	9
Figure 5 - R-PHY Remote Scheduling Model .....	11
Figure 6 - Remote US Scheduling Service Model.....	12
Figure 7 - REQ-GNT Operation Modes .....	13
Figure 8 - UEPI Architecture with Centralized US Scheduling.....	14
Figure 9 - UEPI Architecture with Remote US Scheduling .....	15
Figure 10 - R-PHY DEPI connection with Centralized US Scheduling.....	16
Figure 11 - R-DEPI Connectivity with Remote US Scheduling.....	16

Figure 12 - Remote US Scheduling Service API - Scope and Choices .....	17
Figure 13 - Remote US Scheduler Module in Relation with Top Level RPD Configuration Objects.....	18
Figure 14 - Remote Upstream Scheduler Data Model .....	18
Figure 15 - MAC Domain REQ-GNT Scheduler Data Model.....	20
Figure 16 - MAC Domain MAP Builder Data Model .....	21
Figure 17 - MAP Slot Base Type and Derived Types.....	22
Figure 18 - MAC Domain MAP-UCD Replication Data Model .....	23
Figure 19 - Upstream Scheduling Locality Optimization.....	24

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - REQ-GNT Delay Elements and the Impact of the MAP Interval and CIN Distance .....	8

# Introduction

The cable access network is undergoing a radical transformation from the traditional integrated CCAP architecture to a distributed access architecture (DAA), driven by the growing capacity crunches and the cost pressures to deliver gigabit broadband services. With DAA, cable operators are able to push fiber deeper and replace legacy fiber nodes with DAA devices, achieving higher capacity with both better signal quality and reduced service group sizes.

Depending on how the CCAP MAC and PHY functions are separated, there are two basic architectural options to DAA, Remote PHY (R-PHY) and Flexible MAC Architecture (FMA).

In the R-PHY architecture, the PHY element is removed from the CCAP core and added to the fiber node as a Remote PHY Device (RPD). The basic design philosophy is to put the least amount of hardware and software at the endpoints and keep the complexity centralized. It also allows operators to leverage existing CCAP functions as much as possible for a fast and seamless transition to DAA with both integrated PHY and the Remote PHY potentially connected to the same CCAP core.

The FMA, on the other hand, moves both the CCAP MAC and PHY elements to the node, either as an integrated Remote MAC-PHY Device (RMD) or a combination of Remote MAC Core (RMC) and RPD. Essentially, an FMA DAA device is a small scale CMTS without the routing and management functions. Compared to R-PHY, the DAA device requires significantly more hardware and software functions, and inevitably imposes design and deployment challenges when the device is constrained by power and cost. Moreover, given the complexity of the DOCSIS MAC layer, FMA requires a fairly comprehensive standard interface for the upper network layer and management applications to talk to the DAA devices.

From the latency point of view, one architectural difference between FMA and the R-PHY today is the location of the upstream (US) scheduler. With all the MAC layer functions centralized at the CCAP core, R-PHY has been using a centralized US scheduling scheme that requires the request (REQ) and grant (GNT) information to be exchanged across the Converged Interconnect Network (CIN). CIN delay has no impact on the US scheduling latency as long as it is not the dominating factor, which is the case when GNTs are carried in the de facto 2 millisecond (ms) MAPs, and the CIN distance is within the normal 100-mile (160 km) DOCSIS operational range assumed for I-CMTS deployment [1].

As the network keeps transitioning to DAA, there are, however, reported cases 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 new low latency applications, like cloud gaming and mobile xHaul [2], the DOCSIS REQ-GNT protocol is being tightened to shorter MAPs, such as 1 millisecond MAPs, on DOCSIS 3.1 OFDMA channels [3]. In such circumstances, the CIN delay could be exposed as a significant factor in the REQ-GNT latency equation.

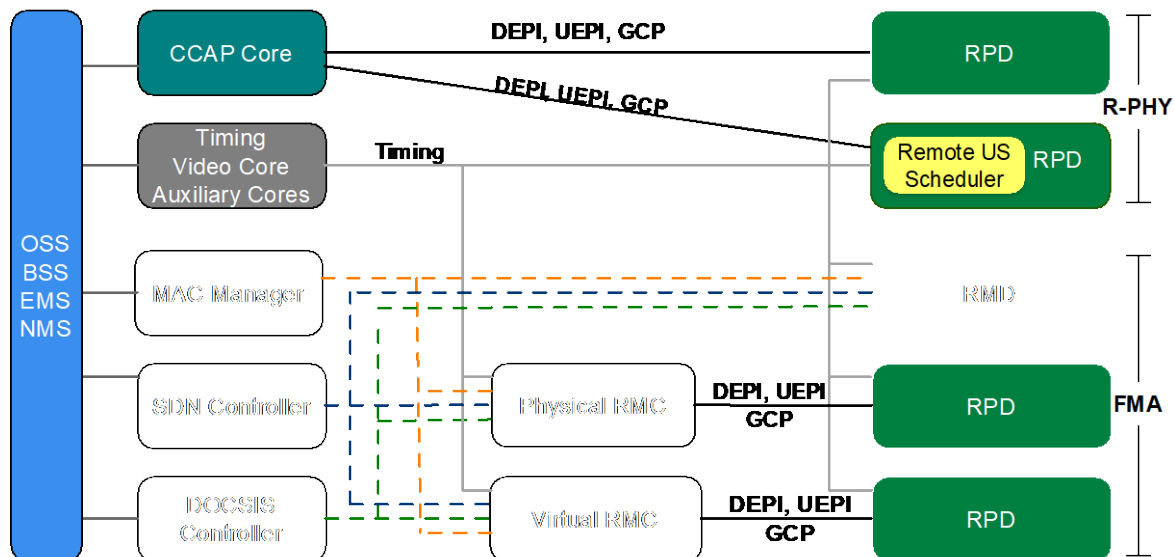
The reason why FMA is immune to CIN delay is because the US scheduler, the MAC element that handles REQ-GNT, is co-located with the PHY where the REQ is received; while in the R-PHY case, the US scheduler is at the core, separated from the RPD across the CIN. This realization leads to the question to be tackled in this paper: Is it possible to put a remote US scheduler at the RPD to help with the latency sensitive REQ-GNT processing?

The remote US scheduler idea was actually considered way back at the beginning of the R-PHY design and development and is mentioned as an option in the current R-PHY specification. However, since the initial R-PHY deployment goal was to replace I-CMTS, it was deferred as a future enhancement. Now the



time has come to move forward with the remote US scheduler design to provide the low-latency scheduling (LLS) needed for long-distance R-PHY deployment.

R-PHY with a remote US scheduler adds a new DAA scenario as shown in Figure 1. Latency-wise, it is equivalent to FMA, however, with much less cost and complexity. It offers FMA-lite functionalities with R-PHY's efficiency and simplicity.



**Figure 1 - DAA Options Including R-PHY with Remote US Scheduler**

Since the remote US scheduler is internal to the RPD, it can leverage the RPD hardware and software platform and the established forwarding, control and management plane interfaces such as R-UEPI, R-DEPI and GCP. The remote US scheduler APIs will be Yang data model-based and will be able to take advantage of the new control plane infrastructure proposed for R-PHY2.0 [4].

This paper is organized as follows. Section 1 explains the rationale for splitting the upstream scheduling between the core and the RPD with a remote RPD US scheduler. Section 2 defines the DOCSIS upstream service model and the remote upstream scheduler service categories. Section 3 examines the split upstream scheduling impact on R-UEPI, and Section 4 examines the impact on R-DEPI. Section 5 discusses the remote US scheduler APIs using the data modeling approach. Section 6 explains how scheduling locality optimization may be applied to improve latency and efficiency for a R-PHY system as a whole. Finally, the paper will be concluded by summarizing the study highlights.

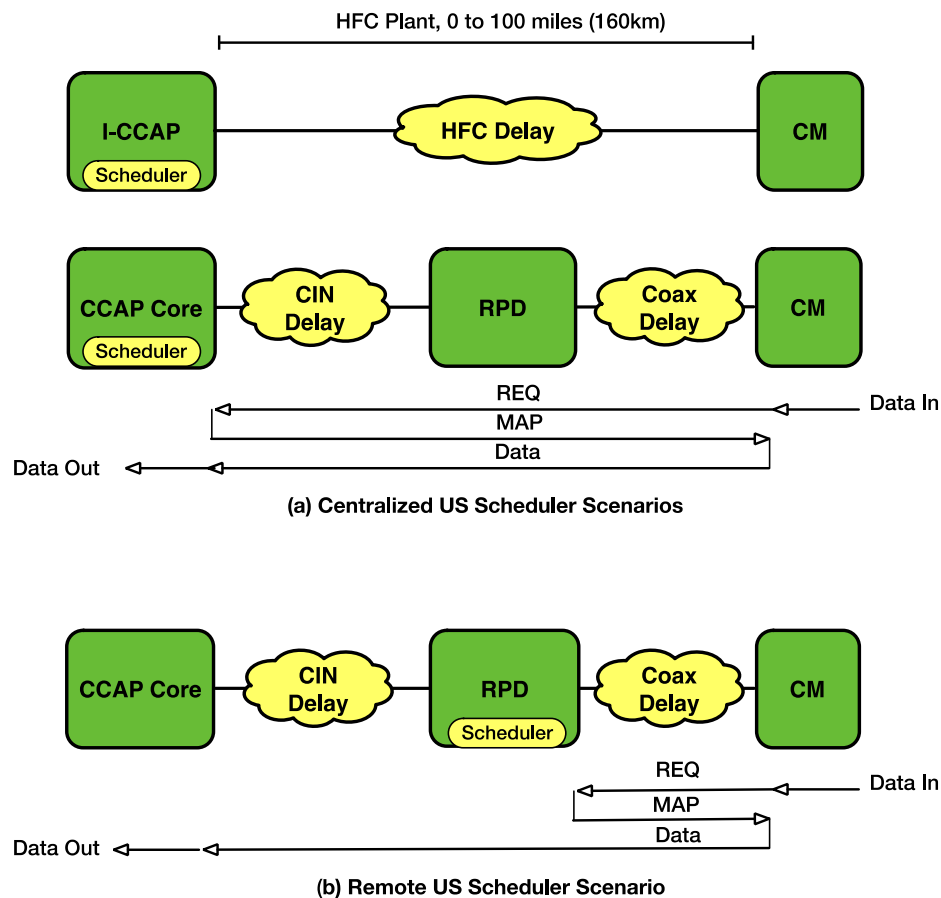
## Content

### 1. R-PHY US Scheduler Location Options

Since the beginning of the R-PHY architecture, there has been a technical debate as to where the US scheduler should be placed. Should it be in the CCAP core where the rest of the software is or should it be in the RPD with the US PHY? To answer this question, there are both business and technical reasons to consider when choosing one location over the other.

From a technical point of view, latency is the main consideration in comparing the two location options. In this perspective, R-PHY with a centralized US scheduler is equivalent to I-CCAP when operating at 2 ms MAPs over a 100-mile plant. R-PHY with a remote US scheduler is expected to provide better latency when operating at shorter MAPs and across a longer CIN distance.

Figure 2 depicts centralized vs. remote US scheduler deployment scenarios and the traverse paths of the REQ-GNT(MAP) messages.



**Figure 2 - Centralized vs. Remote US Scheduler Deployment Scenarios.**

### 1.1. The Case for a Centralized US Scheduler


The basis of the R-PHY architecture is to move the PHY and replace the analog optical link between the CCAP and the node with a digital link. Just with this initial step, cable operators would be able to get better SNR performance, pull the fiber deeper, rebuild the plant and cut a large N+M service group into much smaller ones. All these can be achieved by simply moving the PHY element out of the CCAP core, while keeping all MAC elements including the DOCSIS US scheduler centralized. This also allows operators to leverage the existing CCAP MAC functions to simultaneously support both integrated PHY and remote PHY for a seamless transition to DAA.

Besides the business reasons, the main technical reason for picking centralized US scheduling is based on whether the CIN delay is the dominating factor affecting the request-grant (REQ-GNT) latency.

DOCSIS uses a REQ-GNT protocol to arbitrate US channel access among CMs. When US data arrives at a CM, the CM sends a REQ to the CCAP US scheduler. The US scheduler arranges an individual transmission opportunity by encoding a GNT for the requesting CM in the bandwidth allocation map (MAP) message and broadcasts it to the listening CMs. For regular sized MAPs, from the CM point of view, at any time, there is always one MAP in use, and one MAP “on deck” that is about to be used. So the shortest possible latency for the CM to receive the GNT from the CCAP is every other MAP, which translates to two MAP intervals in time. Most of the time, the REQ needs to wait at the CMTS for the next MAP to be built, with an average wait time around half of the MAP interval. As a result, the total REQ-GNT protocol minimum delay with a 2 ms MAP interval is about 5 ms.

The REQ and GNT messages also take time to process and transport between the CMTS scheduler and the CMs. In this perspective, the REQ-GNT delay is the sum of the REQ propagation time over the coax plant and CIN (in the case of a centralized scheduler), REQ queueing time at the CMTS, CMTS MAP processing time, MAP propagation time over the CIN (in the case of a centralized scheduler) and coax plant, CM MAP processing, as well as the necessary US and DS PHY serialization / framing time. If the total REQ-GNT processing and transport delay, which includes the CIN delay, is less than or comparable to the REQ-GNT protocol delay, R-PHY centralized US scheduling will have no impact on US latency.

Figure 3 visualizes the composition of the REQ-GNT delay elements, and the impact of the MAP interval and the HFC/CIN distance. The latency values of the corresponding delay elements are listed in Table 1, calculated at different MAP intervals, 2 ms vs. 1 ms, and different HFC/CIN distances, five miles (short) vs. 100 miles (maximum I-CCAP operation range). Note that since the coax delay is negligible compared to the CIN delay in R-PHY, the CIN distance and the HFC distance are used interchangeably in this paper.



**Figure 3 - REQ-GNT Delay Elements and the Impact of the MAP Interval and CIN Distance**

As shown, the impact of the MAP interval is mainly reflected in the DOCSIS REQ-GNT protocol delay, which is 2.5 times of the MAP interval as explained earlier. The CIN distance, on the other hand, affects the round-trip CIN delay for transporting the REQ and MAP (GNT) messages. The maxima between the REQ-GNT protocol delay and REQ-GNT transport-processing delay determines the minimum REQ-GNT delay. At a 2 ms MAP interval, the REQ-GNT protocol delay is the dominant factor within the I-CCAP maximum operation range of 100 miles. In this case, R-PHY with centralized scheduling is equivalent to I-CCAP, and there is no latency benefit for moving the scheduler to the RPD.

**Table 1 - REQ-GNT Delay Elements and the Impact of the MAP Interval and CIN Distance**

LATENCY (ms)		TIME VARIANCE			
		2 ms MAP Interval		1 ms MAP Interval	
		5-mile HFC/CIN	100-mile HFC/CIN	5-mile HFC/CIN	100-mile HFC/CIN
<b>REQ-GNT Transport &amp; Processing Delay (Dt)</b>	REQ-GNT (MAP) CIN delay round trip	0.08	1.6	0.08	1.6
	Average REQ queuing delay	1	1	0.5	0.5
	CMTS MAP processing time	0.4	0.4	0.4	0.4
	CM MAP processing time	0.6	0.6	0.6	0.6
	PHY delays and system margin	1.4	1.4	1.4	1.4
	Total	3.48	5	2.98	4.5
<b>REQ-GNT Protocol Delay (Dp)</b>	2* MAP Interval +				
	Average REQ queueing delay	5	5	2.5	2.5
<b>MIN REQ-GNT Delay</b>	max (Dt, Dp)	5	5	2.98	4.5

## 1.2. The Case for a Remote US Scheduler

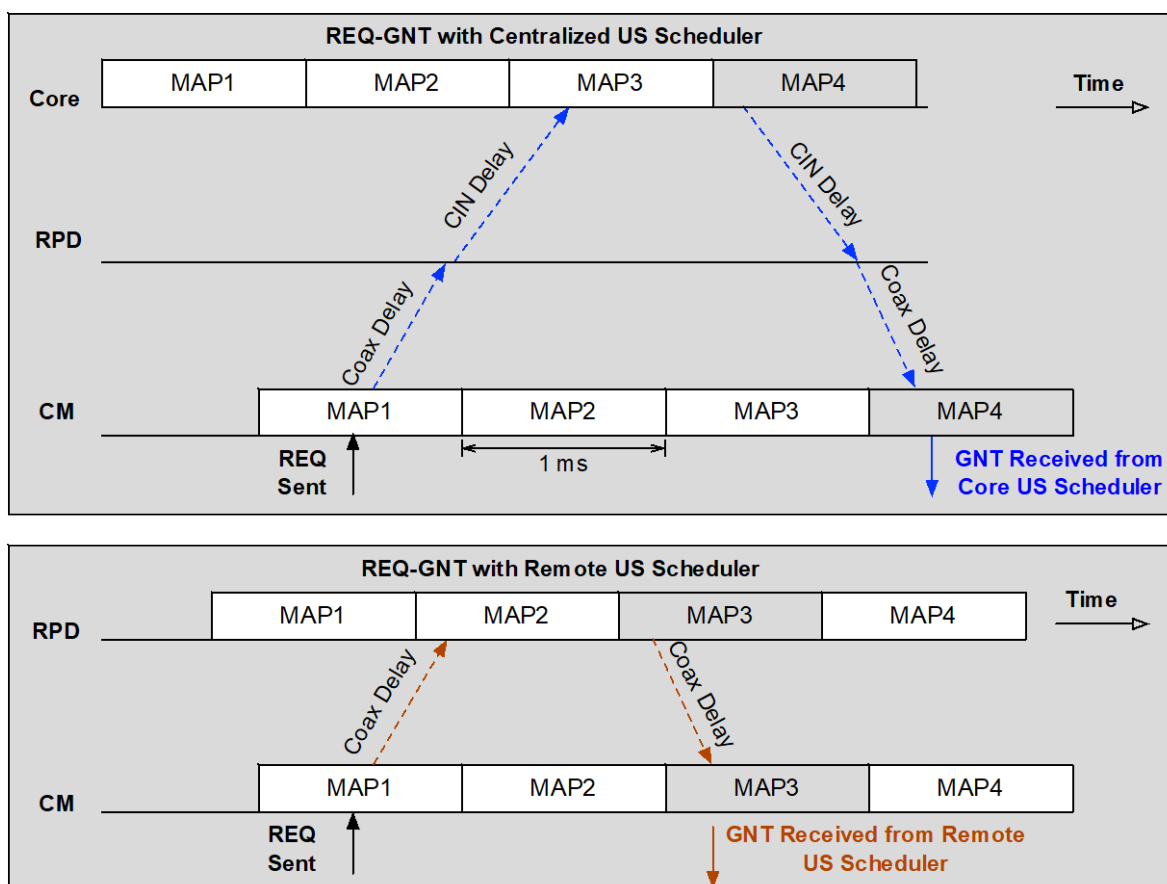
As cable network evolves, new use cases for remote US scheduling start to emerge, primarily driven by the need for long-distance R-PHY deployment and low-latency support over DOCSIS.

### 1.2.1. REQ-GNT Protocol Tightening

As part of the low-latency initiatives, the DOCSIS REQ-GNT protocol is being tightened with shorter MAPs (1 ms) and shorter CMTS MAP processing time (400  $\mu$ s) in order to support low latency over DOCSIS [3]. At the 1 ms MAP interval, the REQ-GNT protocol delay is cut by half, which makes the REQ-GNT processing and transport delay a likely dominating factor of the minimum REQ-GNT delay, as in the example shown in Figure 3. Removing CIN delay in this case does improve the latency performance especially for the longer CIN distance case.

It's also interesting to note from Figure 3 that for long-distance R-PHY deployment, reducing the MAP interval has little impact on the REQ-GNT delay unless the scheduling location is at the RPD. On the other hand, for a short-distance R-PHY deployment, the scheduler location does not matter much if the REQ-GNT processing delay is comparably large with respect to the CIN round trip delay.

Figure 4 illustrates the impact of the CIN delay on the REQ-GNT time when the MAP interval is comparable to the CIN delay. For the REQ sent during the MAP1 interval, the GNT issued by the remote US scheduler at the RPD is available to the CM in MAP3. In comparison, the GNT issued by the core US scheduler for the same REQ is available to the CM at MAP4, as the REQ missed MAP3 building time at the core US scheduler due to the CIN delay.



**Figure 4 - CIN Delay Impact on the REQ-GNT Latency at Short MAP Intervals.**

### **1.2.2. Long Distance R-PHY Deployment**

In a long-distance R-PHY deployment, the CIN distance is beyond the 100-mile range, making CIN delay a dominant factor in the REQ-GNT latency. Long-distance R-PHY deployment may be needed for hub site consolidation where the CCAP core is moved to a central headend or a regional data center. In addition to latency, longer CIN distance may also result in larger jitter due to the queue buildups in the network. Using centralized scheduling for the REQ-GNT handling obviously has challenges to meet the latency and jitter requirements for timing-sensitive applications such as voice and gaming.

On the other hand, adding a remote US scheduler at the RPD can effectively decouple the CIN from the REQ-GNT loop, offering a significant latency improvement for a long-distance R-PHY deployment.

### **1.2.3. CCAP Realtime Performance Acceleration**

DOCSIS US scheduling, characterized by its stringent timeliness, is a real-time task that presents design challenges in the virtualized or cloud CCAP core environment, as the software components with different levels of criticality are all running on the same compute platform. In other words, proper separation and isolation must be in place for the US scheduling task to meet the MAP timing requirements.

With the recent advance of DOCSIS 4.0 technologies, including FDX DOCSIS and extended spectrum, there is a pressing need to scale up the US scheduling capacity to match the upstream bandwidth capacity, which can be up to 50x more than what legacy DOCSIS can offer.

Distributing the timing-sensitive US scheduling tasks to the RPDs can effectively isolate the real-time processing and prevent the disruptions from the non-real time tasks running at the core. In this perspective, the RPD remote US scheduler essentially accelerates the CCAP real-time performance and achieves the scaling required for supporting DOCSIS 4.0 with a distributed scheduling model.

#### ***1.2.4. Easy to Add and Flexible to Adapt***

From an R-PHY architecture point of view, the remote US scheduler is a component internal to the RPD, and can therefore leverage the RPD hardware and software platform, forwarding, control and management interfaces, and the common configuration and performance monitoring infrastructure.

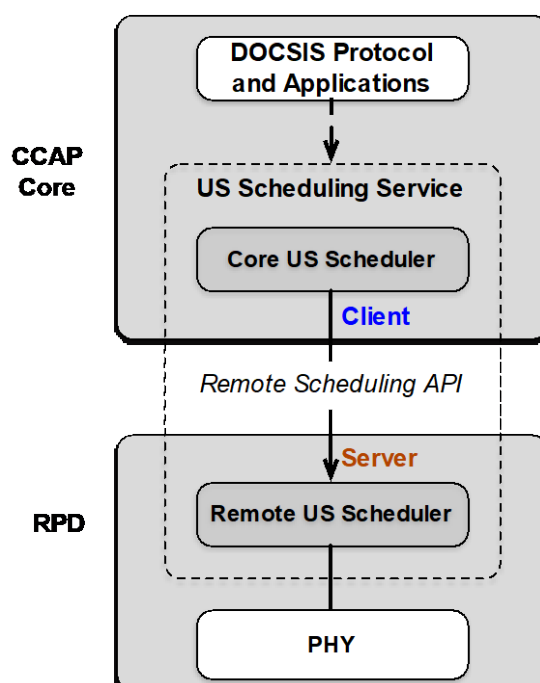
The R-UEPI and R-DEPI architecture established today can adapt to the remote US scheduler by changing the endpoint location of certain scheduler related pseudowires (PWs) on the MAC side. In other words, no change is needed on the US or DS PHY silicon.

US scheduling with the remote US schedulers is not a decentralized scheduling approach. Instead, it is a distributed scheduling scheme with centralized control performed by a core US scheduler, such that scheduling tasks can be load balanced vertically between the core US scheduler and the remote US scheduler at per-US service flow and / or per-channel basis. This architecture permits latency and efficiency optimizations as well as backward compatibility with legacy RPDs.

The RPD remote US scheduler is therefore a light-weight, promising solution to address R-PHY latency concerns and accelerate real-time performance for the virtual or cloud CCAP core environment.

## **2. Remote US Scheduling Services**

In a R-PHY system that has remote scheduling, there are actually two schedulers. There is the core scheduler that lives in the CCAP core, and the remote scheduler that lives in the RPD. Together, they form a client-server relationship and collectively provide upstream scheduling services to other MAC elements and applications as shown in **Figure 5**.



**Figure 5 - R-PHY Remote Scheduling Model**

The core US scheduler oversees the scheduling needs and manages remote scheduling services for latency sensitive scheduling tasks. It optimizes latency performance for the system as a whole by controlling which scheduling tasks to run remotely and when to trigger them. For the remaining MAC layer and upper layer applications, the core US scheduler provides a top-layer scheduling service, for example enabling an activated service flow, arranging US transmission opportunities for ranging and probing, or coordinating with profile management and proactive management applications with special grant allocations in MAPs.

Since US scheduling is per-US, it is technically possible to have some channels running with a remote scheduler and other channels running with a centralized scheduler. For example, legacy DOCSIS 3.0 ATDMA channels could remain with a fully centralized scheduler and the DOCSIS 3.1 OFDMA channels could run with a distributed scheduler.

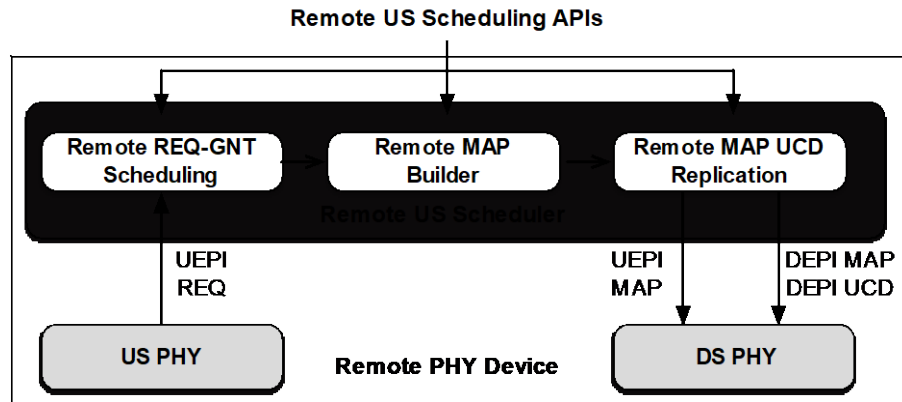
The northbound interface between the core US scheduler and the upper MAC and applications is internal to the core, therefore there is no need for standardization. The southbound interface facing the remote US schedulers is a candidate to be standardized via a set of remote US scheduling APIs. This will allow the CCAP core to interoperate with the remote US schedulers from different RPD vendors.

On the RPD side, the remote US scheduler implements the services declared by the remote US scheduling APIs. During the run time, the remote US scheduler decodes the incoming scheduling requests from the core US scheduler, executes scheduling actions and encodes responses.

The remote US scheduler will provide the following services, also shown in **Figure 6**, namely:

- **Remote REQ-GNT Service:** generates GNTs at the RPD where the REQs are received.
- **Remote MAP Builder Service:** encodes the scheduling decisions into per-US channel DOCSIS MAP messages.

- **Remote MAP UCD Replication:** replicates MAPs and UCDs and transmits them in order in DEPI and UEPI formats. For DEPI MAP and UCDs, the MAP and UCD replicas will be generated for each downstream channel designated to carry MAPs and UCDs for a given US channel.



**Figure 6 - Remote US Scheduling Service Model**

## 2.1. Remote REQ-GNT Service

The REQ-GNT service implements the reactive scheduling process where grants are given in response to the REQs received from the CMs. When REQ-GNT service is offered remotely at the RPD, it effectively cuts out the CIN round trip delay from the REQ-GNT transport time.

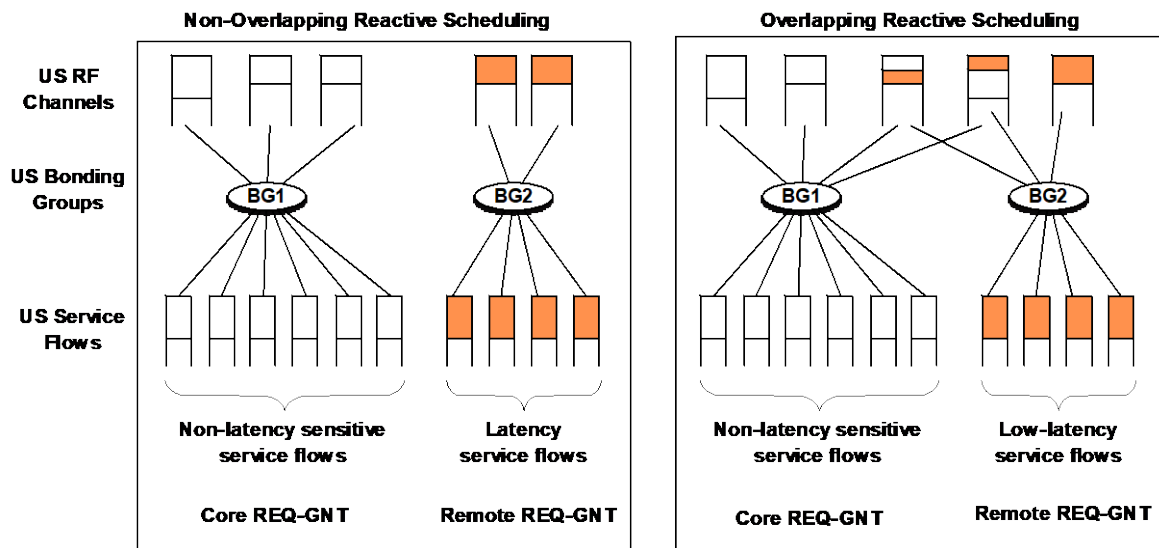
The scope of the REQ-GNT service can be defined on a per-US service flow basis. Once the DOCSIS control plane activates an US service flow, the core US scheduler can request the remote REQ-GNT service by calling the remote REQ-GNT service API. Through the API, the core US scheduler can pass in all the necessary parameters including active QoS parameters, US channel set, and SID to channel assignment associated with the US service flow.

An US service flow without the remote REQ-GNT enabled will remain served by the core US scheduler. With this split-scheduling arrangement, the CCAP core can adjust the remote scheduling workload based on service flows latency requirements and the RPD's capability and capacity in supporting the remote scheduling.

Depending on how the REQ-GNT scheduling tasks are partitioned, there are two REQ-GNT operational modes, overlapping and non-overlapping, as shown in **Figure 7**. In the overlapping mode, the US channel set for the core REQ-GNT service may overlap with the US channel set used by remote REQ-GNT service. In the non-overlapping mode, there is no channel overlapping between the two scheduling entities.

The non-overlapping REQ-GNT scheme is the basic operation mode which is sufficient to separate out the latency sensitive service flows vs. the latency tolerant service flows between the remote US scheduler and the core US scheduler. The overlapping REQ-GNT scheme is an advanced mode, which offers a bandwidth efficiency benefit by allowing the spectrum resource to be shared between the two scheduling entities. The remote scheduling APIs will be structured to allow the flexibility for supporting both REQ-GNT operation modes.





**Figure 7 - REQ-GNT Operation Modes**

## 2.2. Remote MAP Builder Service

The remote MAP builder provides the DOCSIS MAP encoding service to translate a scheduling decision into a MAP slot that represents an US transmission opportunity at specific time and frequency. It ensures the nominal MAP message interval and the MAP encoding consistency with the PHY layer configurations carried in the UCD messages.

The remote MAP builder serves the core US scheduler for the portion of the scheduling that is not handled by the remote REQ-GNT service, including the allocations of the ranging request opportunities, probing, OUDP burst transmission opportunities for DOCSIS 3.1 profile test, unsolicited grants for UGS, RTPS and PGS, or proactive data grants based on traffic predictions.

The remote MAP builder service also allows the client to request a MAP slot at a projected time with specifications for periodicity and jitter tolerance. This enables certain core applications like US symbol capture and sounding in FDX, that are latency tolerant however need to be triggered at a specific time in MAP.

Within the remote scheduler, the remote MAP builder is the next node in the service chain after the remote REQ-GNT process. It can therefore be used to serve the REQ-GNT scheduling module internally by converting a grant in bytes to a grant in minislots and encoding it as a MAP IE.

The remote MAP builder service can be enabled on a per-US channel basis. Given the remote REQ-GNT service is enabled on a per-US service flow basis, the remote MAP builder service needs to be enabled on all US channels in the bonding group associated with the US service flow, which in turn requires the remote MAP and UCD replication service described below.

## 2.3. Remote MAP and UCD Replication Service

After a DOCSIS MAP is built for an US channel, it typically needs to be replicated, as multiple DS channels may be assigned to carry the MAPs for the US channel. Similarly, the DOCSIS UCD message needs to be replicated for each of the MAP carrying DS channels, and sent in sequence with the MAPs to ensure MAC and PHY consistency upon a UCD change.

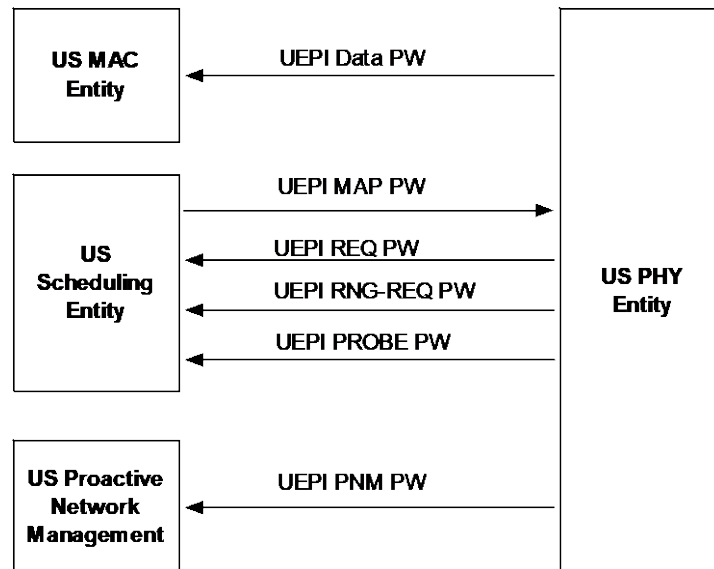
For any US channel enabled for remote MAP builder service, the MAPs and UCDs must be replicated remotely as the next step. However, remote MAP and UCD replication can be enabled as an independent service by itself even in the centralized scheduling case, which will help offload the CCAP processing if the replication is performed in software, and reduce network traffic load, especially as MAPs and UCDs need to be high priority across the CIN.

To facilitate remote MAP and UCD replication, the DS packet scheduler at the CCAP core must take into consideration the bandwidth consumed by the MAPs and UCDs on the replicated DS channels when shaping the DS traffic flows.

### 3. R-UEPI in Remote US Scheduling

The Remote Upstream External Physical Interface (R-UEPI) [5] consists of a set of L2TPv3 PWs connecting the US MAC and PHY in between the CCAP core and the RPD. The information exchanged between them includes various DOCSIS US bursts in the RPD to core direction, and DOCSIS MAPs in the core to RPD direction. In the R-UEPI architecture today, the centralized US scheduler at the core is a common end point for four types of UEPI PWs listed below:

- |            |   |
|------------|---|
| MAP PW     | From CCAP core to the RPD, containing the DOCSIS MAP messages. There is one MAP PW session for each US channel.                   |
| REQ PW     | From RPD to the CCAP core, containing DOCSIS REQ information extracted from the REQ burst or the piggybacked data burst.          |
| RNG-REQ PW | From RPD to the CCAP core, containing DOCSIS ranging-request message and the US PHY metrics measured from the ranging burst sent. |
| Probe PW   | From RPD to the CCAP core, containing the US PHY metrics measured from the probes.  |



**Figure 8 - UEPI Architecture with Centralized US Scheduling**

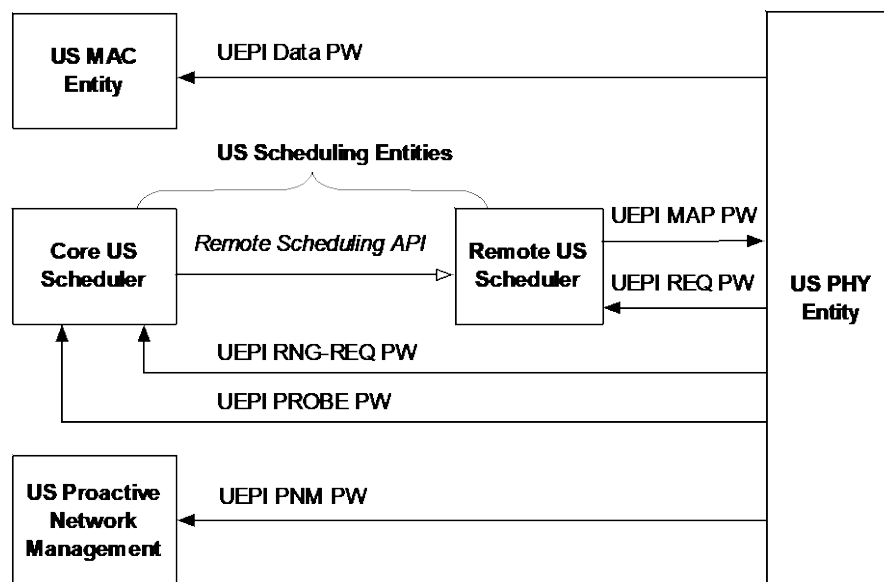
As part of the scheduling function moved to the RPD, the end points of the relevant UEPI PWs will naturally split between the core US scheduler and the remote US scheduler depending on the work load partition between the two scheduling entities.

**Figure 9** shows an R-UEPI arrangement option in a remote US scheduling scenario. The UEPI MAP PW and the REQ PW are shortened and terminated at the remote US scheduler co-located with the US R-PHY, the RNG-REQ PW and probe PW remain terminated at the core US scheduler. This arrangement is based on the required MAC processing and response time per DOCSIS protocol. For ranging and probing, the response time limit is 200 ms, well beyond the time scale of the CIN delay. The REQ-GNT time on the other hand needs to be as short as possible, in the 1~2 ms range, for supporting low-latency data services.

If a UEPI PW is terminated locally at the RPD, the UEPI control plane does not need to use L2TPv3 signaling. Instead, the local CPU can do the UEPI configuration through direct register access, same as in the embedded UEPI architecture.

From the US PHY point of view, there is no difference between an embedded PW or an external PW, as the UEPI framing and the forwarding plane setup will remain the same, therefore there is no impact to the US PHY silicon.

In summary, the native UEPI architecture can readily support the R-PHY architecture with remote US scheduling. It has the flexibility required for load balancing the US scheduling tasks to maximize the latency benefit of the remote scheduler and the core computation capacity for non-real time scheduling services.

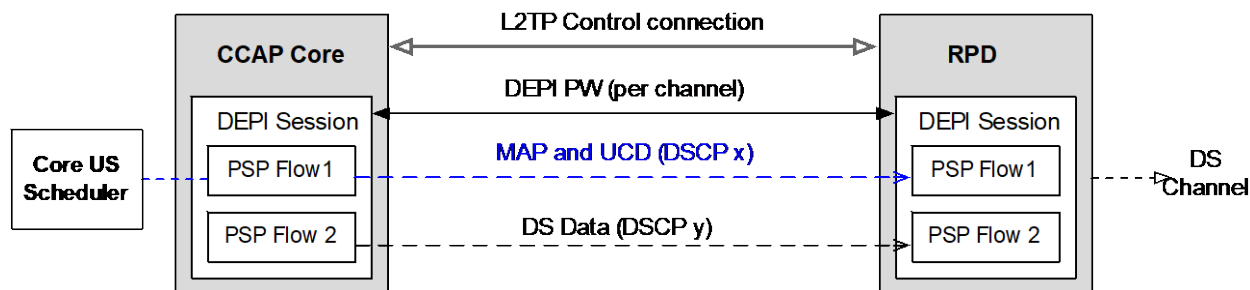


**Figure 9 - UEPI Architecture with Remote US Scheduling**

## 4. R-DEPI in Remote US Scheduling

The Remote Downstream External Physical Interface (R-DEPI) [6] is used to carry the DS DOCSIS data and signaling from the CCAP core to the RPD. It uses L2TPv3 PWs and separate packet stream protocol (PSP) flows for data traffic and MAP/UCDs.

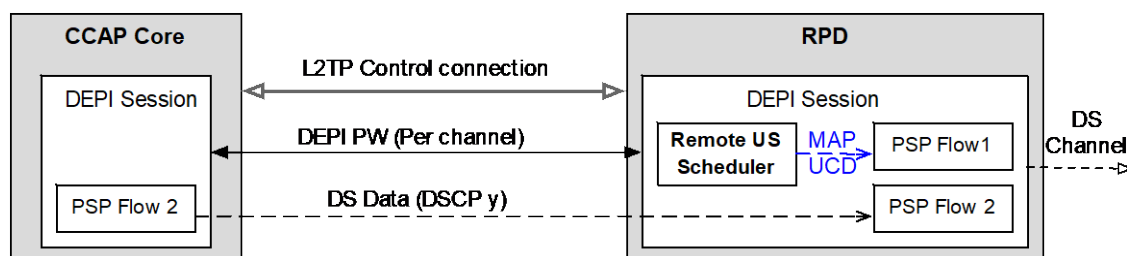
**Figure 10** shows a typical R-DEPI connection setup with centralized US scheduling. The CCAP core and the RPD form a pair of L2TP control connection endpoints that have a DEPI PW session for each channel. Each session has two PSP flows, PSP flow1 for DOCSIS MAPs and UCDs and PSP flow2 for carrying regular data traffic. PSP flow1 is typically encoded with the DSCP expedited forwarding code point for the MAPs and UCDs to be delivered at high priority across the CIN.



**Figure 10 - R-PHY DEPI connection with Centralized US Scheduling**

Since the sequence of the DEPI PSP segments is maintained per PSP flow, PSP flow1 can have a completely different forwarding path from the rest of the channel. This flexibility allows the DEPI MAPs and UCDs to be injected from a location totally different from the DS data PSP flow. **Figure 11** shows the DEPI connections in the remote US scheduling case, where DEPI MAPs and DEPI UCDs are inserted locally at the RPD. From the DS PHY point of view, there is no difference between the centralized US scheduling case and the remote US scheduling case, therefore there is no silicon impact to the DS PHY.

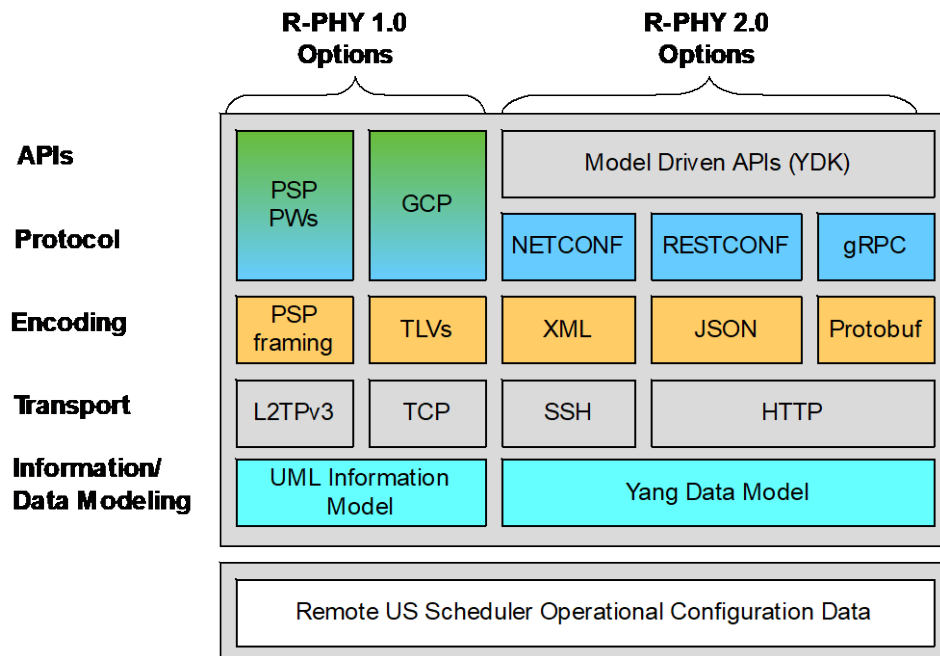
In summary, the remote US scheduling scheme can be readily supported in the R-DEPI architecture.



**Figure 11 - R-DEPI Connectivity with Remote US Scheduling**

## 5. Remote US Scheduling APIs

The scope of work for defining the remote US scheduler APIs falls into two categories, configuration and operational data modeling which determines the content of the API, and the mechanism to express and convey the APIs between two network entities including transport, encoding and protocol design/selections. **Figure 12** shows the API stack in today's R-PHY, R-PHY 1.0, and the new options for R-PHY 2.0 [4].



**Figure 12 - Remote US Scheduling Service API - Scope and Choices**

R-PHY 1.0 uses the UML information model and a messaging scheme that is built on DEPI/UEPI framing and GCP TLVs. APIs are described as rules and requirements declared in the R-PHY specifications published by CableLabs. This process is expected to be overhauled in R-PHY 2.0, where cloud-friendly APIs can be directly enforced based on the Yang data models. Tools like the Yang Development Kit (YDK) can be used to compile the Yang data model to provide APIs in several programming languages. These APIs provide precise definitions of the service contract, automatic data validations, and abstractions of the protocol, encoding and transport details. The data model-based API specifications are expected to significantly simplify the feature development by avoiding the conventional pitfalls caused by different interpretations of the written requirements.

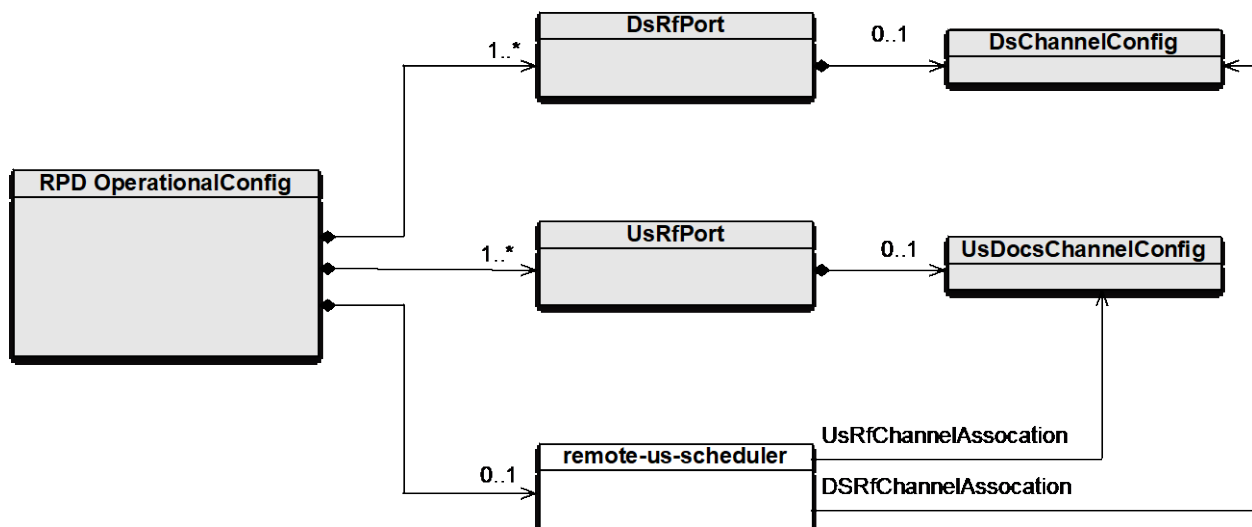
This section focuses on the data modeling portion of the APIs, in other words, what information needs to be exchanged between the CCAP core and the remote US scheduler for configuration management, service enabling and operational state collections. The API messaging mechanism will leverage the R-PHY 1.0 or R-PHY 2.0 messaging methodologies.

### 5.1. Top-Level RPD Operational Configuration Objects

The remote US scheduler module will be a new managed object visible from the top under the RPD operational configuration root module, as shown **Figure 13** in conventional UML format to show the relationship with other R-PHY top-level objects as specified in [7].

An RPD that is capable of remote US scheduling contains one *remote-us-scheduler* module, which holds the configuration and operational data objects accessible through the remote US scheduling APIs.

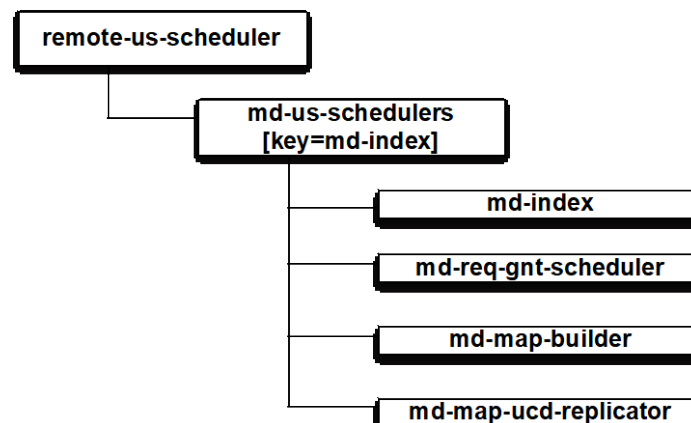
The *remote-us-scheduler* is associated with the existing RPD DS and US objects to access the RF channel configurations through reference needed for MAP building and MAP/UCD replications.



**Figure 13 - Remote US Scheduler Module in Relation with Top Level RPD Configuration Objects**

## 5.2. Remote Upstream Scheduler Data Object Tree

The *remote-us-scheduler* module includes all the configuration and operational objects for performing remote US scheduling at the RPD. **Figure 14** is a simplified Yang data model diagram showing the main branches of the *remote-us-scheduler* data tree. As shown, the *remote-us-scheduler* contains a list of per MAC domain US scheduling instances. Each MAC domain US scheduling instance collects the configuration and operational data elements for all the remote US scheduling service categories.



**Figure 14 - Remote Upstream Scheduler Data Model**

The main elements included in the *md-us-scheduler* are listed below:

<b>md-index</b>	Contains the MAC domain index that identifies a MAC domain from the list
<b>md-req-gnt-scheduler</b>	Contains the information elements for managing the remote US REQ-GNT services for the US service flows in the given MAC domain.

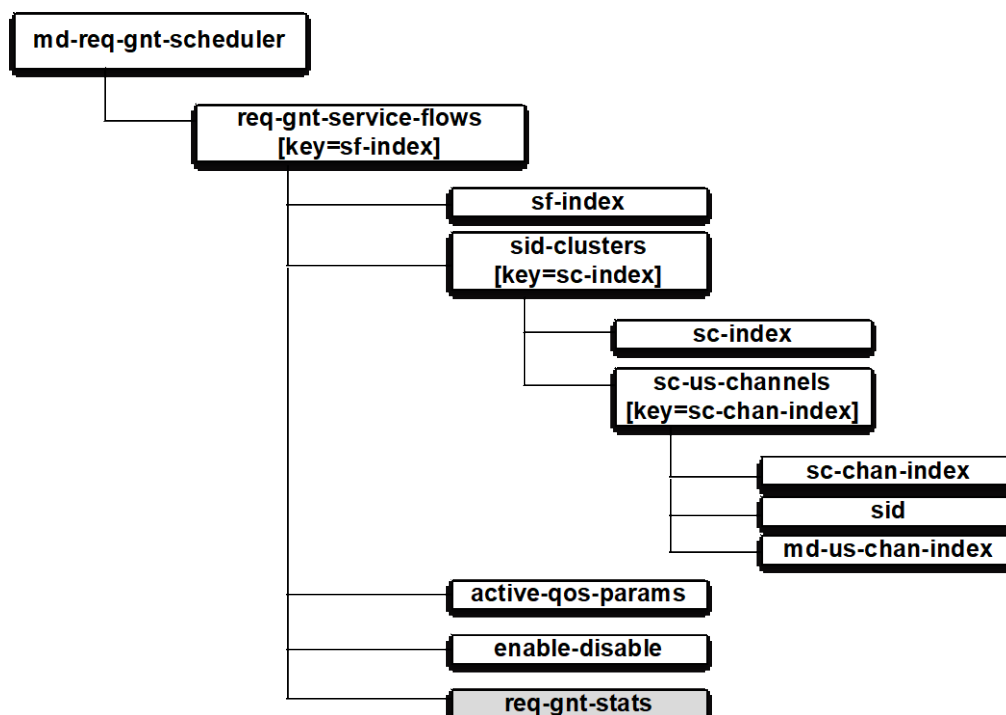
<b>md-map-builder</b>	Contains the information elements for managing the remote MAP-builder services for the US channels in the given MAC domain.
<b>md-map-ucd-replicator</b>	Contains the information elements for managing the remote MAP and UCD replication services for the given MAC domain.

### 5.2.1. MAC Domain REQ-GNT Scheduler

The MAC domain REQ-GNT scheduler, *md-req-gnt-scheduler*, is the container that organizes the configuration and operational data for a list of US service flows enabled for the remote REQ-GNT scheduling service, as shown in **Figure 15**. In this diagram, the configuration data elements (read-write) are shown as white boxes, and the operational data elements (read-only) are shown as shaded boxes.

Each US service flow element contains the following items:

<b>sf-index</b>	Contains the service flow index that identifies a service flow from the list
<b>sid-clusters</b>	Contains the SID cluster configurations that associate the SIDs carried in the REQ with proper US channel resources
<b>active-qos-prams</b>	Contains the service flow active QoS parameter configurations, for example traffic priority, maximum sustained rate, etc.
<b>enable-disable</b>	Contains the action to enable or disable the remote REQ-GNT service for the given service flow
<b>req-gnt-statistics</b>	Contains the operational data of the REQ-GNT statistics, for example, the total number of bytes requested and the number of bytes granted on the given service flow

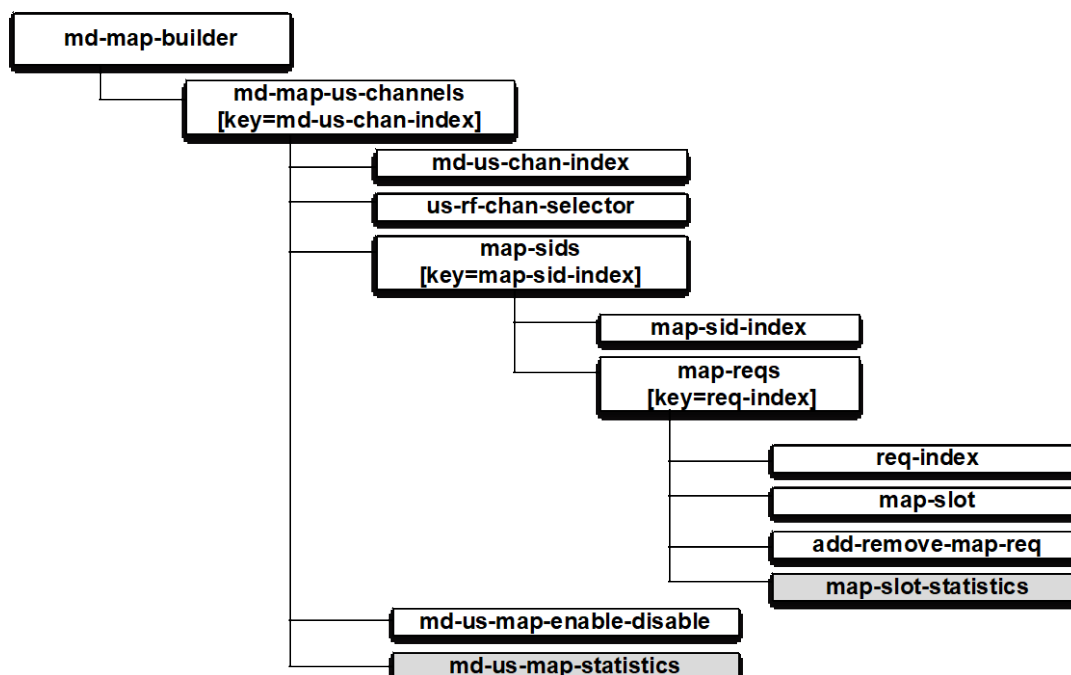


**Figure 15 - MAC Domain REQ-GNT Scheduler Data Model**

### **5.2.2. MAC Domain MAP Builder**

The MAC domain MAP builder, *md-map-builder*, is the container that organizes the configuration and operational data for a list of US channels intended for the remote MAP-builder service, as shown in **Figure 16**.



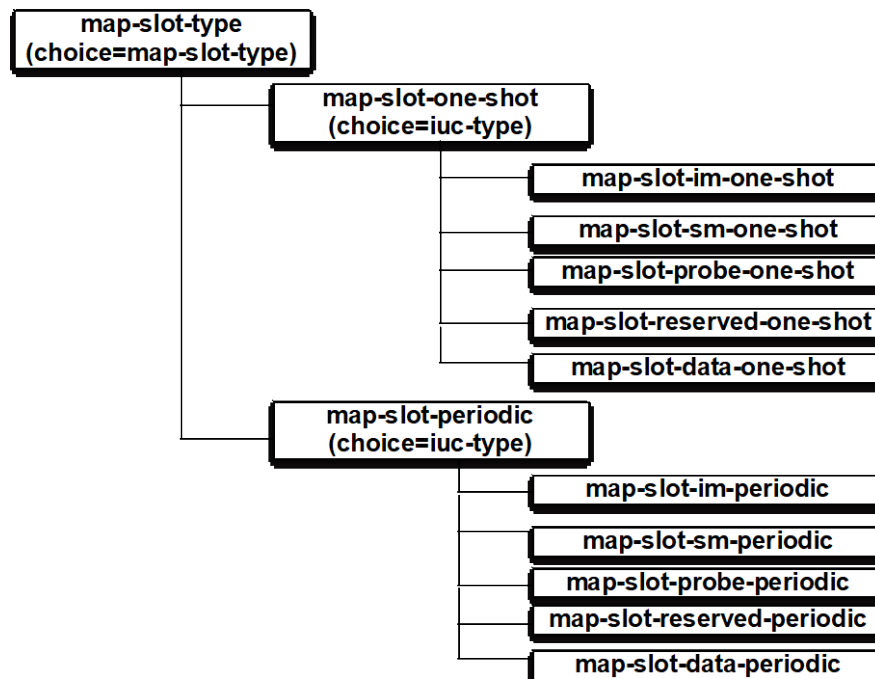


**Figure 16 - MAC Domain MAP Builder Data Model**

Each MAC domain US channel element contains the following items:

<b>md-us-chan-index</b>	Contains the US DOCSIS channel index that identifies an US channel from the per-MAC domain channel list.
<b>us-rf-chan-selector</b>	Contains the RF channel identifier, as defined in [6], associated with the MAC domain US channel.
<b>map-sids</b>	Contains a list of MAP-SID elements that require transmission opportunities in MAPs. Each MAP-SID element contains a list of requests for transmission opportunities, <i>map-slots</i> .
<b>md-us-chan-enable-disable</b>	Contains the action to enable or disable the MAP-builder service on the given US channel.
<b>md-us-map-statistics</b>	Contains the MAP builder statistics, for example, the percentage of minislots allocated for each map-slot type.

The MAP slot request, *map-slot-req*, contains a scheduling decision that needs to be translated into a slot in MAP, or a transmission opportunity from the CM's point of view. The CCAP core uses this API to schedule transmission opportunities that are not handled by the remote US scheduler, such as ranging, probing, profile testing, spectrum capture, unsolicited or proactive data granting, etc. The MAP slot request can be either one-shot for certain ad hoc scheduling needs, or periodic at regular intervals. **Figure 17** shows the types of MAP slots that the CCAP core can request for MAP-build services.



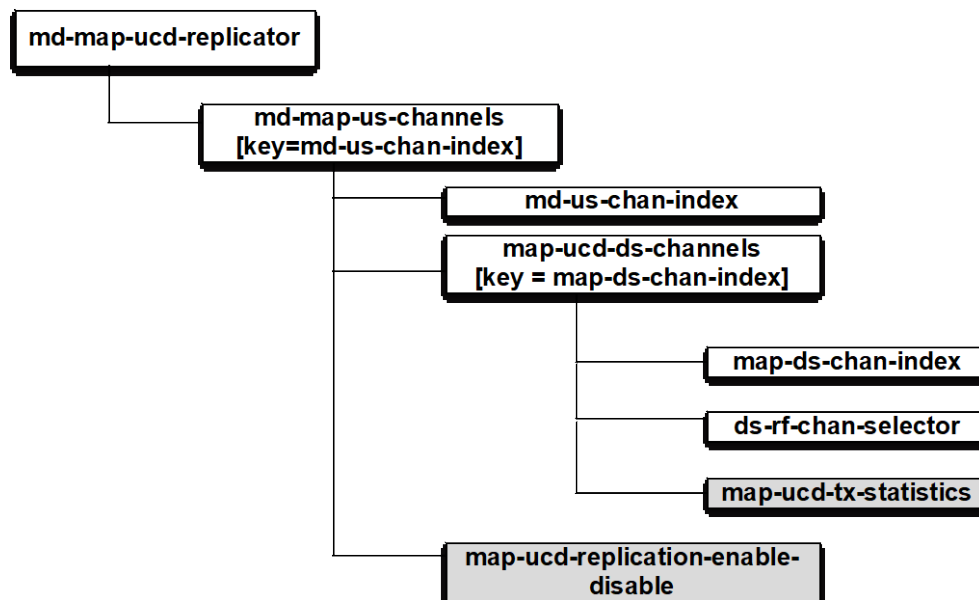
**Figure 17 - MAP Slot Base Type and Derived Types**

### **5.2.3. MAC Domain MAP-UCD Replicator**

The MAC domain MAP-UCD replicator, *md-map-ucd-replicator*, contains the US channel to downstream channel binding information for the MAP and UCD replication function, as show in **Figure 18**.

Each MAC US channel element contains the following items:

<b>md-us-chan-index</b>	Contains the US DOCSIS channel index that identifies an US channel from the per-MAC domain channel list.
<b>map-ucd-ds-channels</b>	Contains the list of DS channels assigned to carry the MAPs and UCDs for the associated US channel. Each DS channel element contains the association to the DS RF channel identifier, and the MAP UCD transmission statistics.
<b>map-ucd-replication-enable-disable</b>	Contains the action to enable or disable the MAP and UCD replication service for the given US channel.

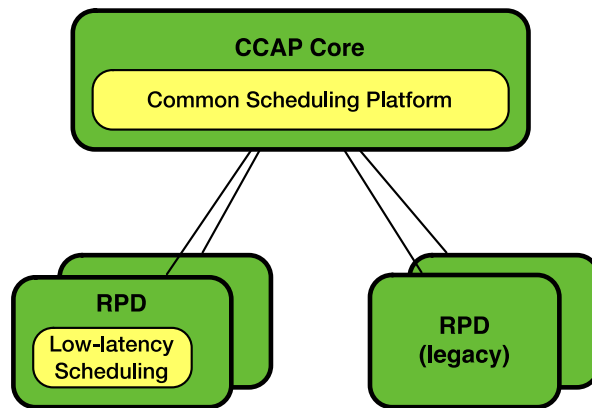


**Figure 18 - MAC Domain MAP-UCD Replication Data Model**

## 6. Locality Optimization with Remote US Scheduling

The essence of split US scheduling with the remote US scheduler at the RPD is to enable system-wide locality optimization. By co-locating the reactive granting portion of the scheduling with the PHY where the REQ is received, latency is improved by avoiding the CIN delay in the REQ-GNT process. By keeping the latency-tolerant and computation-intensive part of the scheduling centralized at the core, efficiency is maximized by positioning the core as the common computation platform accessible to all RPDs, as shown in **Figure 19**.

This scheduling model is different from the decentralized scheduling model as used in FMA that has no centralized control from the CCAP core. The core US scheduler in the R-PHY case provides a unique value for global locality optimization that takes into consideration the per-service flow latency requirement, RPD capabilities / constraints, and the CCAP core real-time processing capacity. In this perspective, the ability to do remote scheduling with centralized control gives R-PHY the architecture advantages to achieve low-latency and high efficiency, and remain backwards compatible with legacy R-PHY deployments.



**Figure 19 - Upstream Scheduling Locality Optimization**

## Conclusion

Cable networks have been going through a radical transformation, changing from bandwidth-limited and latency-tolerant networks to a high-capacity, low-latency, multi-service edge access network. Adapting to the change by enabling low-latency US scheduling in R-PHY is a step taken to accelerate this transformation and prepare cable networks for the future.

R-PHY low-latency US scheduling involves moving the latency-sensitive scheduling tasks such as REQ-GNT handling to the RPD, while keeping the latency-tolerant scheduling tasks centralized to retain the centralized MAC advantages. For the US scheduling service, the core and RPD form a client-server relationship, where the RPD remote US scheduler provides services to REQ-GNT low-latency service flows, builds MAPs for both core and the remote schedulers, and replicates MAP UCDs to the proper DS channels. Such services can be precisely defined using data model-based APIs, which can be autogenerated based on published Yang data models.

The RPD remote US scheduler can be built on top of an existing R-RPY platform, which contains the basic MAC and PHY building blocks and the glue logic. The addition of the remote US scheduler has no impact on the US PHY or DS PHY silicon and can be readily supported by the R-UEPI and R-DEPI architecture, as the only change needed is the endpoint location of the UEPI and DEPI PWs on the MAC side.

The addition of the remote US scheduler to the R-PHY US scheduling scheme enables a distributed scheduling model where the core can optimize the scheduling locations and conduct the vertical load balancing between the core and the RPD. This scheduling model is unique to the R-PHY architecture, being able to achieve system-wide optimization in both latency and efficiency, and simultaneously maintaining backwards compatibility with legacy RPDs.

## Acknowledgements

The authors would like to thank our Cisco colleagues Brian Bresnahan and Wenkai Zhu for their valuable contributions to the paper.

# Abbreviations

CM	cable modem
CCAP	converged cable access platform
CIN	converged interconnect network
DAA	distributed access architecture
DS	downstream
FMA	flexible MAC architecture
GNT	DOCSIS bandwidth grant
LLS	low latency scheduling
REQ	DOCSIS bandwidth request
R-PHY	remote PHY
US	upstream

## Bibliography & References

- [1] John Chapman, Gerry White, Hang Jin; *Impact of CCAP to CM Distance in a Remote PHY Architecture*; 2015 Spring Technical Forum Proceedings
- [2] John Chapman, Jennifer Andreoli-Fang; *Mobile Backhaul over DOCSIS*; 2018 SCTE Technical Forum Proceedings
- [3] “CM-SP-MULPIv3.1-I18-190422: MAC and Upper Layer Protocols Interface Specification”, CableLabs, 2019
- [4] Pawel Sowinski, Andy Smith, Tong Liu etc; *Remote PHY 2.0*; 2019 SCTE Technical Forum Proceedings
- [5] CM-SP-R-UEPI-I05-170111: “Remote Upstream External PHY Interface Specification”, CableLabs.
- [6] CM-SP-R-DEPI-I06-170111: “Remote Downstream External PHY Interface Specification”, CableLabs.
- [7] CM-SP-R-PHY-I10-180509: “Remote PHY Specification”, CableLabs

# **Delivering the Highest IP Video Quality Efficiently While Improving Customer Experience**

A Technical Paper prepared for SCTE•ISBE by

**Garey Hassler**

Distinguished Engineer Software Architect

Comcast

1515 Wynkoop St. Ste. 200, Denver, CO 80202

720-502-3717

Garey\_Hassler@comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
IP Video Service Resilience.....	3
1. Redundant and Resilient Linear IP Video Service.....	3
1.1. Synchronized Transcoding and Packaging.....	3
1.2. Hypertext Transfer Protocol Version 2.0 (HTTP/2) .....	4
1.3. Content Delivery Network .....	5
1.4. Video Quality and Automated Origin Process.....	6
2. Proof-of-Concepts.....	7
2.1. HTTP/2 and Session Management .....	7
2.1.1. Dispelling the Single Viewer Misconception.....	9
2.2. Video Quality.....	9
Conclusion .....	10
Abbreviations.....	11
Bibliography & References .....	11

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Simple HTTP Service.....	4
Figure 2 – HTTP Service with CDN.....	5
Figure 3 – HTTP Fan-Out Example.....	5
Figure 4 – HTTP/2 and Session Management Test Environment.....	7
Figure 5 – Video Quality and Automated Origin Processing POC.....	10

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Average Range of 50 Test Runs of 5 Minute Duration .....	8

# Introduction

Television audiences expect the best video quality their devices can render and to have it available for viewing in seconds. Years of video quality branding, like high-definition (HD) and 4K, and other consumer educational efforts, have fashioned discernible viewers capable of distinguishing sharpness, brightness and resolution and equating this to quality. Furthermore, with decades of broadcast consumption, viewers have been observing diminishing tune times, or time-to-first-frame (TTFF), creating an expectation, which can be challenging to achieve in a redundant IP video service.

Multichannel video program distributors (MVPDs) can dramatically improve their ability to achieve viewers' expectations by enhancing a video origin to support Hypertext Transfer Protocol Version 2.0 (HTTP/2) [1], synchronizing video transcoding and packaging, applying video quality measuring, and automating origin selection.

A proof-of-concept (POC) was constructed to independently measure the video quality of a single video source, relative to the synchronized transcoding and packaging of two geographically dispersed sites. Applying automated decision logic on the quality measurements to select an origin for distribution on a per-fragment interval ensures the highest quality is distributed to viewers. Utilizing (HTTP) version 2.0 PUSH method to immediately distribute the IP content through a content delivery network (CDN), reduces network traversal from a round-trip-time to an end-to-end latency.

Evaluations of the POC were conducted in simultaneous test executions compared with existing client/server production components, in a non-isolated network, thus any network congestion or impairments applied equally to the test measurements. Preliminary test results over multiple executions indicate a 45% – 56% improvement in TTFF, while delivering the highest quality video fragment to all test players.

## IP Video Service Resilience

### 1. Redundant and Resilient Linear IP Video Service

Delivering a linear IP video service redundantly and resiliently is a challenge for MVPDs. With the introduction of duplicate versions of a channel and quality monitoring tools, new latencies can impact the service. The following subsections look at some of the drivers leading to the proof-of-concepts covered in this document.

#### 1.1. Synchronized Transcoding and Packaging

With the development of coordinated processing by transcode vendors, it's possible to configure a pair of transcoders into a master/slave relationship, allowing a MVPD to improve the reliability and resilience of an IP video service. When transcoding from two origins which are synchronized and operating at two independent datacenters, devices are capable of seamlessly retrieving and playing the video content distributed from either video origin. This feature significantly reduces the likelihood a customer will experience the dreaded buffering symbol



when there is a network impairment or hardware disruption. However, this introduces new complexities and questions for an MVPD to consider:


- What methodology should the distributor employ when operating the origin: even distribution across channel lineup, load balancing by viewer, primary versus secondary model, or some other approach entirely?
- How does a distributor ensure customers receive the best available video quality when there is a video provider source or network issue occurring at only one origin?
- How does a distributor prevent the introduction of additional latencies into the delivery of the video?

Serving the broadest set of devices is every MVPD's desire. To facility this objective, a standard from SCTE affectionately called Common Intermediate Format [2] is the output produced by the packaging system. The media files utilized are a transport stream and this format is maintained until processing occurs at a Just-In-Time-Packager, which transforms the manifest and media files into a format requested by a client.

## 1.2. Hypertext Transfer Protocol Version 2.0 (HTTP/2)

The advanced features provided by the HTTP/2 standard afford some significant benefits for the distribution of video content. The standard defines an upgrade mechanism for client devices that support both HTTP/1.1 and HTTP/2 to propose use and support of these advances in requests it initiates. Furthermore, the additional features allow compatible servers and clients to support a new feature called Server Push, which essentially allows a server to preemptively deliver associated files to the client based upon the previously initiated request. Common examples of the types of additional files might be images or Cascading Style Sheets (CSS) that are associated with the content requested by the client.

By expanding the interpretation of a request for a linear video channel from a single request for a fragment<sup>1</sup> of some common duration into a broader definition to one of requesting to consume a channel, it becomes possible to theorize and conceptualize a single request for all future variants of a video manifest and media fragments. In its simplest form of a client and server, shown below, a client would initiate a request to “tune” to a channel along with an indication it supports upgrading to the HTTP/2 standard. The server will respond by accepting to upgrade, notifying the client of a future content via a Server Push using the PUSH\_PROMISE of a new stream to be initiated by the Server, and then deliver the requested video manifest file.



**Figure 1 – Simple HTTP Service**

---

<sup>1</sup> A small discrete video file typically of a common size. In this POC, fragment durations are 2 seconds in length.

In this configuration, a Linear IP Video Origin Server can repeatedly initiate a new stream with the client, send a new PUSH\_PROMISE for a subsequent stream and deliver the most recent updated changes to the manifest or media file as they are produced on its normal cadence of processing. The server is capable of detecting or inferring when the client no longer needs the media through the simple termination of the Transmission Control Protocol (TCP) communication link. Alternatively, the server can force a reestablishment of the process at any time by closing the TCP connection with the client.


### 1.3. Content Delivery Network

For an MVPD providing thousands of streams servicing millions or even tens of millions of video customers, the network is far more complicated than the one shown in Figure 1 – Simple HTTP Service. They will frequently include multiple layers of a Content Delivery Network (CDN) to support the caching that reduces latencies and improves resiliency when delivering video over large areas. The diagram below illustrates an example of a mid-level or mid-tier of a CDN cache and an edge cache, the latter often being in close proximity to the consumer or client device.



**Figure 2 – HTTP Service with CDN**

In a typical video provider network there will be multiple mid-tier caches and a greater number of edge caches distributed throughout the service providers network to serve the millions of client devices. Delivery of the video across these growing number of network hops is a common network form or pattern called a “fan-out distribution.”



**Figure 3 – HTTP Fan-Out Example**

The introduction of a CDN results in the need to propagate the HTTP/2 through one or more intermediary systems, disrupting the server and clients' ability to infer when either system no longer needs a video channel. With the introduction of a Session Management module into the CDN service, the “proxying” of the TCP communication link closures and terminations can be propagated through the entire network. The Session Management processing is responsible for maintaining a mapping or association of upstream links (communication links towards the server) and with communication on its downstream links (those communication links that originate from the client side of the diagram.) Once all of the downstream links in a component for a channel terminate, the Session Management closes its upstream link. With this concept, propagation of the inference of channel viewing can be conveyed through the video distribution network without any new messaging beyond standard TCP messages.

#### **1.4. Video Quality and Automated Origin Process**

As mentioned previously, a video distributor wants to be able to provide the best available video quality equally to all of its subscribers. To achieve this goal, the distributor monitors the video quality it is receiving from a source provider as well as the quality being produced by the transforms conducted by transcoding and packaging the media. When augmenting this distribution with redundant transcode and packaging, there are greater opportunities for the quality to vary between those two independent video origins. The MVPD can use a variety of vendors or open source tools to perform the quality measurement, but how does it go about ensuring all viewers are receiving the best quality at a particular moment in time?

The introduction of a new component dubbed in this document as the Video Quality Agent (VQA) can apply processing and decision logic based upon the scores of each transcode and packaging origin pair of a channel. Based upon some threshold of deviation in quality between the two origins, the VQA controls which origin is the source of content for all viewers, through notification to one origin to serve the content, and the other or secondary origin to redirect content to the origin currently serving the channel. This control can consider multiple factors in its decision processing and can be as granular as a single fragment of media.

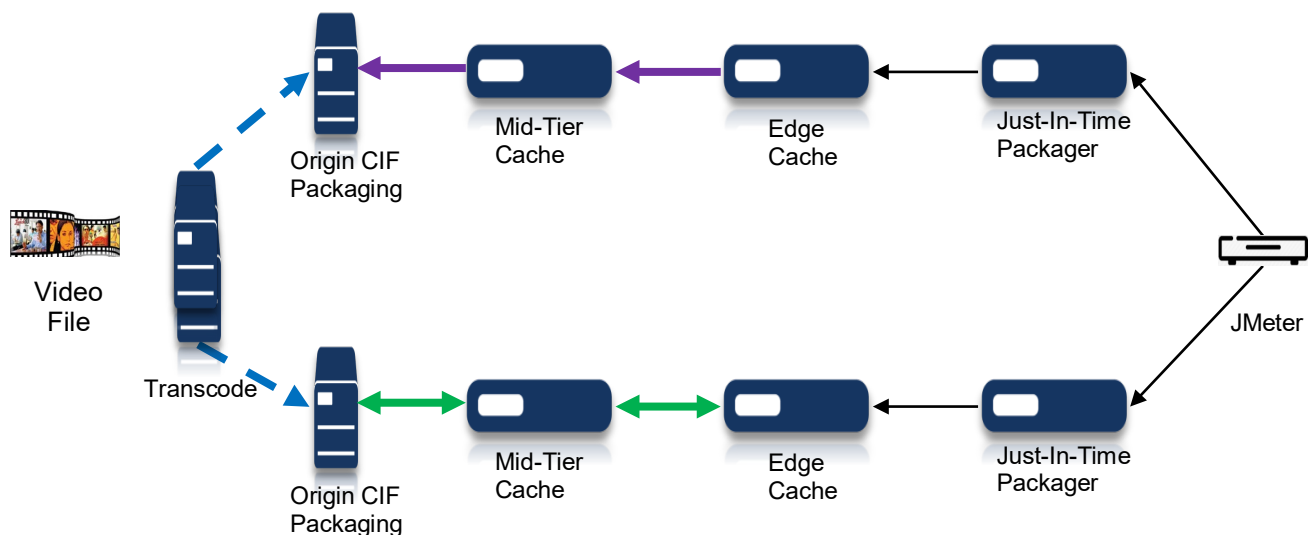
In the situation where a distributor wants to operate in an even distribution of requests to two redundant origins, such as a simple round-robin approach from Client devices, the redirecting of traffic to a single origin introduces additional latencies and increased network load for some clients. When combined with the HTTP/2 Server Push functionality and the expanded Session Management functionality with the two new features, this mechanism can propagate the highest quality in an Automated Origin process. To facilitate the indication to downstream systems that a media file is duplicative from multiple origins, the active Origin Server includes a new custom HTTP header, such as X-Video-Origins, indicating it is the authoritative origin for both redundant servers. The new HTTP header contains the host information of the peer Origin Server of the paired servers and is provided by the Video Quality Agent or through configuration. When the media file is “pushed” or published to the downstream systems connected in the fan-out, the Session Manager of the system creates a new link or index mapping that references the cached file as an item to serve for any request for either origin, thus a request for the secondary origin can be served with the file sent by the primary origin. In a similar manner, the cache can publish

to the next downstream systems that have upgraded to use the HTTP/2 mechanism for delivery. This pattern is repeatable in each system's Session Management processing.

## 2. Proof-of-Concepts

### 2.1. HTTP/2 and Session Management

Evaluating the effectiveness of this conceptual set of hypotheses was conducted using software derived from the preexisting production software systems, emulating the HTTP/2 and Session Management functionality. The diagram below illustrates the test environment created to conduct the POC experiments.



**Figure 4 – HTTP/2 and Session Management Test Environment**

Beginning from the left-hand side of the diagram, an HD mezzanine quality video file is the input into a transcoder, which is processed into 5 MPEG-4 audio and video profiles of two seconds in duration with an average bandwidth of 3.5 Mbps for the highest resolution, 2.0 Mbps for the middle resolution, and 0.5 Mbps for the lowest resolution profile. The output of the transcoder is a multicast UDP packets sent to both Origin CIF Packaging servers, as indicated by the dotted blue lines. The upper Origin serves the media files using the HTTP/1.1 client/server methodology over the systems, indicated by the purple lines, and the lower Origin delivers the media files to the Mid-Tier and Edge-Cache, depicted by the green lines, with arrows on both ends, in reference to communications beginning by a client and then propagating downstream, using HTTP/2 and employing the Session Management functionality. The POC limits the use of HTTP/2 to the Origin, Mid-Tier and Edge Caches for testing for two reasons: 1) to limit the software updates prior to obtaining demonstrative evidence of actual improvements and 2) to validate that the optimization could occur, without the need for all systems to implement simultaneously, before benefits could be realized by the service and the MVPD.

In order to reduce the complexity of the Session Management algorithm, separate communication links were employed for the manifests from those of the fragment files. This

simplified the code development and management of the synchronization, in cases when publishing multiple fragment files prior to publishing the manifest to the next downstream systems.

To minimize external variables, such as network latencies skewing the measurements across each test run or between the two paths, the systems were all deployed in the same datacenter, on virtual machines with the same resource configurations and similar network traversal paths. To further mitigate any outside bandwidth impairments, data collection occurred simultaneously over both methods. Each test execution was conducted by a single JMeter instance, operating within two mutually exclusive threads with the same script to allow for common clock and time collection as they emulated a client's behavior.

Test results of 50 executions of the environment, each 5 minutes in duration on the three different video resolution levels are illustrated below. Table 1 – Average Range of 50 Test Runs of 5 Minute Duration shows the summary of the span of average times over all of the test runs. The test collected the complete time necessary to request and receive a complete copy of the manifest and the next two-second fragment, beginning with a simulated “tune” to the channel. In addition, the average improvement of each test run was calculated with the range of results shown in the final column.

**Table 1 – Average Range of 50 Test Runs of 5 Minute Duration**

<b>Test Case Complexity</b>	<b>Avg. HTTP/1.1 Delivery Range</b>	<b>Avg. HTTP/2.2 Delivery Range</b>	<b>Improvement Range</b>
Highest 3.5 Mbps Profile	13.316393 – 17.230584 (ms)	7.595753 – 10.039682 (ms)	45% - 56%
Middle 2.0 Mbps Profile	10.332659 – 14.64361 (ms)	7.121569 – 8.585287 (ms)	34% - 41%
Lowest 0.5 Mbps Profile	9.518596 – 11.713008 (ms)	6.899321 – 8.29262 (ms)	28% - 35%


In most circumstances, when a viewer tunes to a channel, another viewer or service, like a cloud recording service, will have already requested the media files. This is valid for all except for the very first system needing a file. A technique was defined (though not evaluated in the POC) to improve a client's time-to-first-frame by enhancing the operations of the Origin Server, upon any initial communication link, to request a manifest from a downstream system. This request indicates the support of the upgrade process. Furthermore, it can conceivably be considered to be attempting to consume the corresponding media. The concept is to automatically, preemptively transmit a predefined number of the most recently created media files produced by the server to the requesting system, in anticipation of subsequent requests for the associated media files, upon a client's processing of the manifest file returned in the initial request. Even with modern browser improvements that initiate multiple concurrent requests to fill their video buffers, so as to expedite the retrieval of multiple fragments, the reduced latency and playback initiation occurred in a manner that was perceptively sooner when using the HTTP/2 operations.

### **2.1.1. *Dispelling the Single Viewer Misconception***

Upon initial reflection, most people are under a misconception the improvements demonstrated by the POC don't justify the additional overhead expense of the new processing, because only the first client actually retrieves the files from the origin. However, this does not take into consideration all of the devices actually waiting upon the cache to obtain a copy of the content to service subsequent requests. Most CDNs and caching software today will actually hold or queue up multiple requests received while an outstanding request for the same file is actively being requested from the origin. The feature has different names in CDNs. As examples, Apache Traffic Server calls it "reader-while-writer" and Varnish calls it 'coalesced'. Regardless of the name, all requests received by the CDN, while awaiting the file from the upstream system, experience some latency and delay. In a widely distributed network with a large viewership, it is quite common for tens of thousands of requests to experience some level of queueing in various systems, as the content traverses the fan-out. This is alleviated with the HTTP/2 method, as the manifest is the final file delivered. And, because the Origin initiates the operation, it is only bound by its' computational ability and network transfer capacities to deliver to the next downstream systems.

## **2.2. Video Quality**

A discrete POC was created to evaluate the implications of the Video Quality and Automated Origin Process when employing the delivery of content using the HTTP/2 methodology. The Video Quality Agent (VQA) evaluates the quality of the transcoding and packaging systems, determining a quality score for the combined transforms. This scoring can be performed on a referenced basis, through comparison with the original video file or source input, or non-referenced basis, using no comparison to the input. When the quality produced in one origin path deviates by some margin, perhaps 5%, the VQA directs an Origin Server change. The VQA coordinates and controls all HTTP/2 publishing by directing the Origin CIF Packaging Server with the best quality to enable publishing, and directs the lower quality or scoring Origin to disable publishing.



**Figure 5 – Video Quality and Automated Origin Processing POC**

The VQA is responsible for communicating the host information of the peer Origin CIF Packaging Server for inclusion in the X-Video-Origin. When the video quality varies beyond one of the thresholds, the VQA performs the control logic for the channel's origin switch. From the client's perspective, it continues to request content from the same origin and is completely unaware of the transition between origins.

Testing was conducted to disrupt the video by introducing packet loss which reduced the qualities created at one of the transcoders. The tests were conducted 20 times each, with varying the levels of impairment, in an attempt to cause a large enough issue or disruption to be visually impacting during playback. In all tests, the playback was uninterrupted and continued without any awareness or perception from the viewer.

## Conclusion

The methods described in this document demonstrate the improvements possible for a Linear IP Video Service. Given the range of the test runs, it is safe to infer that the majority of the timing reduction in obtaining a media file from an origin is the result of the file traversing from the origin to a client without the requirement to be initiated by a preceding request. In other words, by transforming the delivery model from one of request-and-response to a client-initiated file delivery, delivery of files can be estimated to be equivalent to the end-to-end latency of the service, instead of the round-trip-time. The techniques outlined leveraging synchronized origins with video quality measurements and controls ensure all viewers receive the highest quality video available, with minimal additional complexity and no additional network communication.

## Abbreviations

CDN	Content Delivery Network
CIF	Common Intermediate Format
CSS	Cascading Style Sheets
HD	High Definition
HTTP	Hypertext Transfer Protocol
ISBE	International Society of Broadband Experts
MVPD	Multichannel Video Programming Distributor
POC	Proof-of-Concept
SCTE	Society of Cable Telecommunications Engineers
TCP	Transmission Control Protocol
TTF	Time-to-First-Frame
VQ	Video Quality
VQA	Video Quality Agent

## Bibliography & References

[1] IETF RFC 7540 2015, *Hypertext Transfer Protocol Version 2 (HTTP/2)*

[2] ANSI/SCTE 214-4 2018, *MPEG DASH for IP-Based Cable Services Part 4: SCTE Common Intermediate Format (CIF/TS) Manifest for ATS Streams*



# Remote PHY 2.0

## The Next Steps For Remote PHY Technology

A Technical Paper prepared for SCTE•ISBE by Pawel Sowinski, Andy Smith and Tong Liu.

**Pawel Sowinski**  
Principal Engineer  
Cisco Systems Inc.  
psowinsk@cisco.com

**Andy Smith**  
Principal Architect  
Cisco Systems Inc.  
andsmit@cisco.com

**Tong Liu**  
Principal Engineer  
Cisco Systems Inc.  
tonliu@cisco.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
1. The Opportunity .....	3
2. Introducing Remote PHY 2.0 .....	3
3. Cloud Friendly Control Plane .....	4
3.1. What is R-PHY Control Protocol? .....	4
3.2. Control Protocol in Remote PHY 1.0 .....	5
3.2.1. GCP Protocol Stack .....	5
3.2.2. Transactions .....	6
3.2.3. Extensibility .....	6
3.2.4. Performance Requirements .....	7
3.3. R-PHY 1.0 Control Protocol Status .....	7
3.4. Remote PHY 2.0 Control Protocol .....	8
3.4.1. Why YANG? .....	9
3.4.2. Proposed Methodology .....	9
3.5. 1.0 to 2.0 Transition .....	11
3.6. Enabling Automation .....	11
4. Model Driven Telemetry .....	12
5. R-PHY with Remote Upstream Scheduler .....	13
6. Data Plane Improvements .....	15
7. Other Functional Improvements .....	17
Conclusion .....	18
Acknowledgments .....	19
Abbreviations .....	19
Bibliography & References .....	20

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – R-PHY Control with multiple RPD and CCAP Cores .....	5
Figure 2 – RPD 1.0 control protocol stack .....	6
Figure 3 – RPD 2.0 control protocol stacks .....	10
Figure 4 – RPD 1.0 and 2.0 control protocol stacks .....	11
Figure 5 – Model Driven Telemetry with R-PHY 2.0 Protocol Stack .....	12
Figure 6 – Centralized vs. remote R-PHY upstream scheduling options for R-PHY1.0 and R-PHY 2.0 .....	15
Figure 7 – MPLS Transport for R-PHY .....	17

# Introduction

In the period between the initiation of the Remote PHY program and the deployment of phase one significant changes have taken place in cable operator services, and network deployments and also in the technologies that have become available to build and operate systems such as Remote PHY. It is appropriate to consider the next generation requirements R-PHY will need to solve and what are the best tools to use to do this.

As the first phase of Remote PHY is moving towards technological maturity, the time is right to plan the next steps in its evolution. The paper proposes a strategy for enhancing the Remote PHY architecture by examining a set of its real and perceived issues, and by suggesting a how to effectively tackle them.

The paper focuses on the following technical issues:

- **Cloud APIs and Automation:** How to transition the R-PHY control plane towards mature, open model driven network management protocols such as NETCONF, RESTCONF or gNMI.
- **Manageability:** A recommendation for model driven telemetry for Remote PHY.
- **Latency:** How to eliminate latency issues resulting from the physical distances between R-PHY system components by incorporating a DOCSIS upstream bandwidth scheduler into the RPD.
- **Operation in Multi-Service Networks:** The application of modern network standards, such MPLS and Segment Routing to R-PHY data plane transport.

The paper presents an in-depth technical analysis and discusses the utility and economic value of the proposed enhancements. The paper demonstrates that the proposed functional advancements, taken together represent the next generation of R-PHY architecture, Remote PHY 2.0.

The paper does not explain the details of R-PHY architecture. The paper is written with the assumption that the reader has at minimum a rudimentary familiarity with Remote PHY. The necessary background information can be found in [RPHYTR] and [RPHY] .

## Content

### 1. The Opportunity

Remote PHY technology has finally entered the phase of wide scale deployments. At the time of this paper's writing, several cable operators are providing commercial service based on R-PHY technology to hundreds of thousands of subscribers. The multi-year Remote PHY standardization efforts led by CableLabs are drawing to a conclusion. Soon, the R-PHY project at CableLabs will enter the maintenance phase. The working group's focus will shift towards fixing specification bugs rather than the definition of new functionality.

The confluence of these events creates a perfect opportunity to take step back, take a critical review of the R-PHY technology, assess its weaknesses and gaps and devise a strategy to best address these issues. This paper presents a menu of options for several selected new R-PHY features. Our intent is to initiate a conversation within the industry about the future direction of R-PHY technology. Therefore, the list of R-PHY 2.0 features discussed within this paper is open to further additions and changes.

### 2. Introducing Remote PHY 2.0

In this paper we refer to the existing R-PHY technology and specifications as R-PHY 1.0 or simply 1.0.

“Remote PHY 2.0” is nothing more than a convenient name chosen as a common label applied to the set of new R-PHY architecture options proposed in this paper. We don’t claim that any of these features cannot be added to existing specifications and products without such a label. We believe however, that there are tangible benefits and a convincing argument to be made for separating these options from current R-PHY technology and packaging them under a new version label. The primary concern is the ability of the existing products to support these new options. The following factors also need to be considered:

- The proposed features are interdependent. For example, the proposed model driven telemetry relies on RPD supporting the data driven control plane.
- The proposed technical solutions do not constitute incremental development. They offer replacement for currently utilized techniques and may not provide backward compatibility.
- It is beneficial to logically separate these options because of the large scope of changes to the involved software infrastructure.

### **3. Cloud Friendly Control Plane**

In this section we propose a strategy to replace the main control protocol deployed in R-PHY 1.0 architecture. First, we describe the existing R-PHY control protocol and analyze its strengths and weaknesses. Later we detail the approach to upgrade the control protocol and how to minimize the transition impact on the existing R-PHY system. Finally, we explain the technical and business benefits of the proposed transition.

#### **3.1. What is R-PHY Control Protocol?**

In a Remote PHY Architecture, the integrated Converged Cable Access Platform (CCAP) is separated into two distinct components. The first component is the CCAP Core. The second component is the Remote PHY Device (RPD). The CCAP Core inherits all I-CCAP functions except for the PHY layer which is implemented in the RPD. The CCAP Core and the RPD communicate over a permanent IP connection.

The relationship between the CCAP Core and the RPD resembles a master-slave communication model. The direction of control is from the CCAP Core to the RPD. The CCAP Core remotely controls the functions of the RPD through a protocol which we refer to as the R-PHY Control Protocol.

The R-PHY control protocol incorporates all elements of the FCAPS (Fault, Configuration, Administration, Performance, Security) management framework. In this context the CCAP Core acts as the Network Management System and the RPD acts as the Managed System. There is however a number of important differences in requirements for a typical FCAPS operation and for a R-PHY control protocol with the Core and RPD having a much tighter coupling than a typical FCAPS manager and client. In many instances, the CCAP Core and the RPD operate with a common set of configuration parameters and state information. Whenever the operator, or internal processes in the CCAP Core impose changes to the values of these parameters or state variables, the control protocol needs to coordinate them between the systems, sometimes with tight real-time constraints. For example, when the configuration of a downstream profile changes on an OFDM channel, the change needs to be enacted in both systems by a detailed procedure prescribed by the control protocol.

The R-PHY architecture incorporates a great deal of flexibility in how the CCAP can be functionally decomposed into a set of independent CCAP Cores. For this reason, each RPD is required to provide service to multiple (from one up to 10) CCAP Cores. Serving multiple masters is hard. The R-PHY

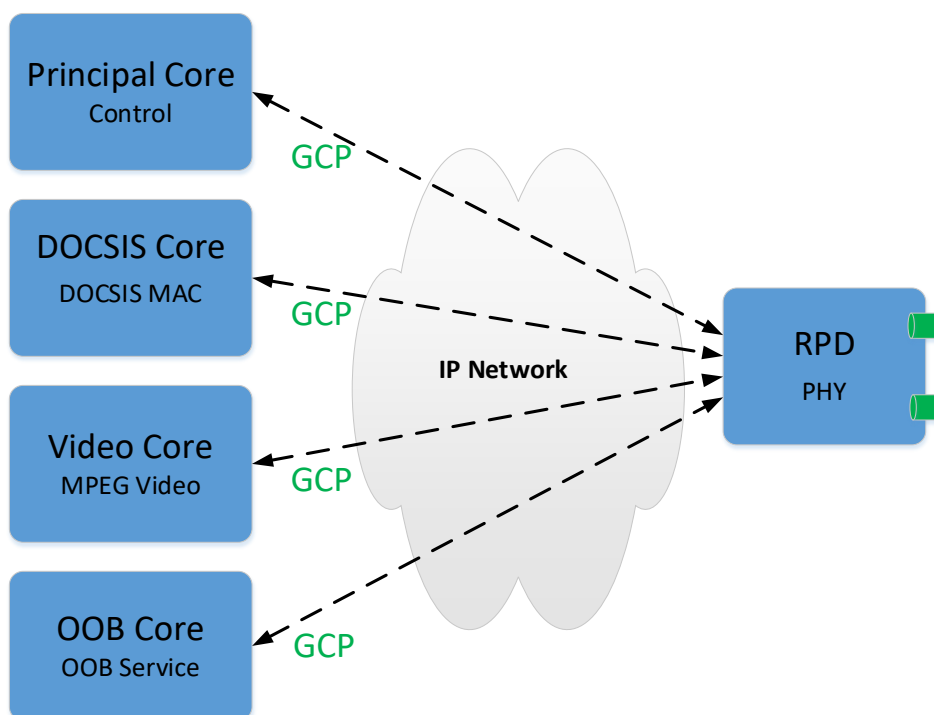
control protocol solves this problem by subdividing the set of managed resources into isolated information silos. Each silo is controlled by a single Core.

For example, a selected CCAP Core, referred to as the Principal CCAP Core, is designated to provide the central management functions such as the initial configuration of the RPD-CCAP Core pairing, the division of managed resource between the Cores, the RPD general configuration, the fault handling, the control over device software upgrades, etc. The Principal Core does not handle any video or data signals. Other CCAP Cores can provide individual CCAP data services, e.g. DOCSIS, or SCTE 55-2 out-of-band service, and only manage the RPD resources dedicated to these services.

### 3.2. Control Protocol in Remote PHY 1.0

In R-PHY 1.0, virtually all aspects of the master-slave relationship are managed with a protocol commonly referred to as Generic Control Plane/R-PHY Control Protocol (GCP/RCP). R-PHY also relies on several other protocols for narrower purposes, such as Layer 2 Tunneling Protocol Version 3 (L2TPv3) control protocol. In this section we focus solely on the GCP/RCP protocol.

Figure 1 shows an example of GCP connections between an RPD and a set of four CCAP Cores.



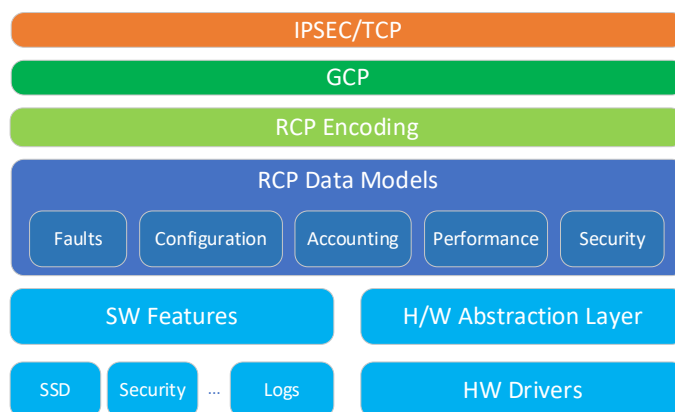
**Figure 1 – R-PHY Control with multiple RPD and CCAP Cores.**

In the example, as mentioned earlier, the Principal Core provides general management of the RPD. The DOCSIS, Video and OOB Cores control distinct sets of RPD resources associated with the services they respectively provide.

#### 3.2.1. GCP Protocol Stack

Figure 2 presents a typical control protocol stack of an RPD. The GCP protocol relies on TCP for reliable transport and on the IPsec suite for security protection. The formats of messages and their exchange rules

are specified in [GCP] . The upper layers implemented by the RCP are documented in [RPHY] and in [RPHYOSS].



**Figure 2 – RPD 1.0 control protocol stack.**

The RCP protocol design was intended to mirror the DOCSIS protocol deployed for communication between the CMTS and Cable Modems. RCP operates as an abstraction layer over the foundation provided by the GCP protocol. Just like the DOCSIS protocol, RCP carries data in hierarchically organized Type-Length-Value (TLV) tuples.

Unlike DOCSIS, RCP defines rules for structured access to an RPD data model represented by a hierarchy of application specific object-TLVs. Object-TLVs form a tree in which each node has a type and either a value or a set of child nodes. The detailed specifications of objects-TLVs serve two roles, the formal definitions of the elements of the information models and how to encode the information in exchanged messages. The definition of RCP object-TLVs includes semantics and syntax, default values, attribute units as well as range and access constraints. The specifications render selected object classes as UML diagrams to provide an informal, visual representation of the information models of object-TLVs constructs. An extensive set of distinct objects describes RPD capabilities. Lastly, RCP also directly incorporates selected DOCSIS messages for configuration of specific DOCSIS channels' operational parameters.

The RCP provides the CCAP Core with the ability to remotely manage properties of modelled objects such as RF channels, RF ports, performance variables, etc., maintained by the RPD. The relative complexity of R-PHY information models, in our opinion, reaches at best modest levels.

### **3.2.2. Transactions**

Just like most network management protocols, RCP allows the CCAP Core to perform CRUD (create, read, update, and delete) operations on object classes and individual attributes. RCP configuration operations are transactional within sets of grouped objects. The protocol also provides a means for the RPD to send asynchronous notifications to CCAP Cores to inform them about defined events, e.g. when the state of a modelled attribute changes or to report errors.

### **3.2.3. Extensibility**

The RCP mirrors the DOCSIS protocol extensibility mechanism through vendor-proprietary extensions at the top of the TLV hierarchy. Some implementations also support proprietary extensions at lower levels in the hierarchy.

### **3.2.4. Performance Requirements**

The volume of traffic exchanged over the GCP/RCP connection is highest during the initial RPD configuration. The size of a configuration set for a typical RPD is modest. It can be measured in 10s of kilobytes. After the initialization, the volume of GCP traffic varies; it depends primarily on the level of status information retrieved from the RPD by the CCAP Core(s). The specifications do not impose limitations on the traffic volume. We can assume that the volume of GCP traffic is not an issue we need to be concerned with.

The RCP operates with few real-time constraints. The majority of RCP protocol interactions require a response within one second. The most stringent real-time requirements are imposed by procedures for dynamic updates to DOCSIS channel parameters such as DOCSIS OFDM profiles or DOCSIS upstream channel parameters. During these procedures, the GCP/RCP transport carries selected DOCSIS messages from the DOCSIS Core to the RPD. The procedures require coordination with parallel procedures conducted between the CMTS Core and Cable Modems. For example, during upstream channel parameters change, the RPD is mandated to process received UCD message in less than 50 msec.

To summarize, the RCP defines abstract information models and a set of protocol rules for CRUD operations on objects from these models. The RCP closely follows the operational principles of network management protocols and is subject to few performance or stringent real-time requirements.

### **3.3. R-PHY 1.0 Control Protocol Status**

At the functional level, few complaints can be made about the R-PHY 1.0 control protocol. The protocol has demonstrated sufficient flexibility to enable effective replication of I-CCAP features in the distributed R-PHY environment without compromising functionality or performance. RCP 1.0 also offers very compact encoding.

Remote PHY 1.0 specifications certainly meet the criteria of an open standard. The specifications have been developed by a working group open to any willing participants. CableLabs processes make the specifications available, royalty free for download to the public. R-PHY 1.0 specifications, including the control protocol have been developed with the goal of ensuring multivendor interoperability. Successful deployments of R-PHY systems with interoperable components from several vendors have proven that the R-PHY 1.0 control protocol has successfully achieved this goal.

The R-PHY 1.0 control protocol is deployed just in one, relatively narrow application field, the Remote PHY architecture. While the development of the R-PHY 1.0 control protocol for a system does not represent a high technological barrier, the number of existing implementations is limited to a handful of equipment vendors and cable operators. The ecosystem of applications and support tools available to test GCP/RCP is scarce or non-existent.

The R-PHY 1.0 control protocol was designed to implement the RPD control functions efficiently and to meet certain DOCSIS real-time requirements. The reliance on a purposely developed and narrowly deployed protocol as well as the inherited real-time constraints made the R-PHY 1.0 control protocol unique.

This uniqueness is also the Achilles heel of the R-PHY 1.0 control protocol because it translates into a set of business issues such as the difficulty in testing and validation, the slower adoption curve and in the end into higher OPEX for the operators.

Without a doubt, the availability and the maturity of the development and test ecosystems has been a contributing factor to some of the interoperability problems encountered during initial deployments. The root causes of interoperability problems can be notoriously complex. This is especially true in a multivendor environment, where vendors cooperate developing on ever-changing specifications but also fiercely compete against each other in the marketplace. The causal analysis of the interoperability issues lies beyond the scope of this paper.

Despite the protocol design based on open CableLabs specifications, the developer community has currently no Open Source code project at its disposal to stop reinventing software each time they need to build a new product.

Another issue is the maturity of the 1.0 protocol definition. We are referring here to the 1.0 control protocol rules, not the information models on which the protocol operates. Design issues are uncovered, new protocol rules are added, existing rules are modified or clarified in each release of the R-PHY specifications. Even though the rate at which the protocol changes are introduced is consistently decreasing, these processes are likely to continue for some time.

Finally, we need to consider the significant changes in the cable operator infrastructure environment in which the R-PHY is implemented. The R-PHY 1.0 protocol was developed with requirements of coherent integration with the physical CCAP infrastructure. As operators transition towards a cloud-based infrastructure, the requirements for the protocol shift as well. The R-PHY control protocol needs to be reimagined to better conform to the new cloud environment.

### 3.4. Remote PHY 2.0 Control Protocol

In this section we describe the proposal for the control protocol for R-PHY 2.0. We explain the goals, the methodology, discuss the feasibility and demonstrate how the proposal addresses the issues with R-PHY 1.0 control protocol identified in the preceding section. Finally, we discuss the options for a transition from 1.0 to 2.0 control protocol.

We propose to transition the R-PHY control protocol away from GCP/RCP towards modern, YANG model driven protocols. YANG, which stands for Yet Another Next Generation, is simply a better choice, especially for cloud APIs. This is not a new idea even in the domain of Distributed CCAP Architectures. The first example is CCAP Config. CableLabs specifications have been publishing YANG based APIs for CCAP configuration for almost a decade. Most recently, the Flexible MAC Architecture (FMA) group has embraced a similar approach using YANG for the management APIs of the FMA MAC Network Element.

YANG is a data-modeling language used to describe network device configuration and operational data developed by the Internet Engineering Task Force (IETF). YANG models the hierarchical organization of data as a tree in which each node has a name and either a value or a set of child nodes. YANG provides clear and concise descriptions of the nodes and of the interaction between them. Details about YANG can be found in **Error! Reference source not found.** and **Error! Reference source not found.**

Moving to a YANG models will address the R-PHY 1.0 control plane issues identified previously, enable cloud friendly tools and automation, and improve system manageability, testability and multivendor interoperability.



### **3.4.1. Why YANG?**

Over the past decade, YANG became a universally adopted standard for modeling of APIs for management of physical and virtual network elements. YANG emerged as the default choice for network management APIs when automation, agility, and scaling are the key requirements. The proliferation of YANG based programmatic interfaces extends into the Internet of Things (IoT) space, and outside of networking, into fields such as medical, vehicular and even aeronautical technology.

The IETF has been developing standards with YANG as the data modeling language for all elements of FCAPS framework. To date, within the IETF, about a hundred of YANG modules have been adopted for standard track and hundreds more are circulating as drafts.

The wide availability of mature development ecosystems, including Open Source code libraries, toolchains and applications are the key factors driving YANG's adoption. The development organizations can pick from dozens of Open Source rooted tools and commercial systems specifically developed to help with YANG model creation, validation and testing. The functionality supported by tools includes conversion from other modeling methods and even automatic API source code generation in modern programming languages.

### **3.4.2. Proposed Methodology**

#### **3.4.2.1. Data Model Translation**

The first step towards R-PHY 2.0 control protocol is the formalization of RPD data models in YANG. The RPD data models can be translated from the current representation in RCP object-TLVs to equivalent YANG modules. Such a translation, in our opinion, is feasible, if not completely straightforward or even somehow mechanical in nature. The existing hierarchy of RCP object-TLVs, and their constraints can be directly replicated into YANG. The translation does not need to result in a perfect mirroring of the existing RCP models. Where optimizations are appropriate, the formalization process can include a desired level of refactoring. Selected YANG modules developed by the CableLabs FMA working group could be adopted for reuse in R-PHY 2.0. The product of the translation will be a set of YANG modules representing the same managed objects of an RPD as those embedded in RCP 1.0.

#### **3.4.2.2. Relaxation of Real Time Requirements**

To eliminate the most stringent real-time requirements on the control protocol, we propose to remove DOCSIS MAC Management messages from the R-PHY 2.0 control protocol. The UCD, OCD and DPD messages are sent from the CMTS to Cable Modems in-band, as packets embedded in downstream data streams. In R-PHY 1.0 these packets simply pass transparently through the RPD. [MULPI] specifies precise procedures by which DOCSIS CMs operate on these messages. An R-PHY 2.0 compatible RPD can snoop the messages from the data plane stream sent to CMs and participate in the channel change procedures just like Cable Modems.

#### **3.4.2.3. Protocol Selection**

The next step is the selection of the protocol over which the CCAP Core and the RPDs exchange information. Three protocols have emerged as the favorite choices for operation with YANG defined APIs. These protocols are NETCONF, RESTCONF and gNMI.

- The Network Configuration Protocol (NETCONF) defines a mechanism for manipulating configuration data and for retrieving operational data. NETCONF carries configuration and

operational data encoded in XML over a reliable transport. NETCONF's definition can be found in **Error! Reference source not found..**

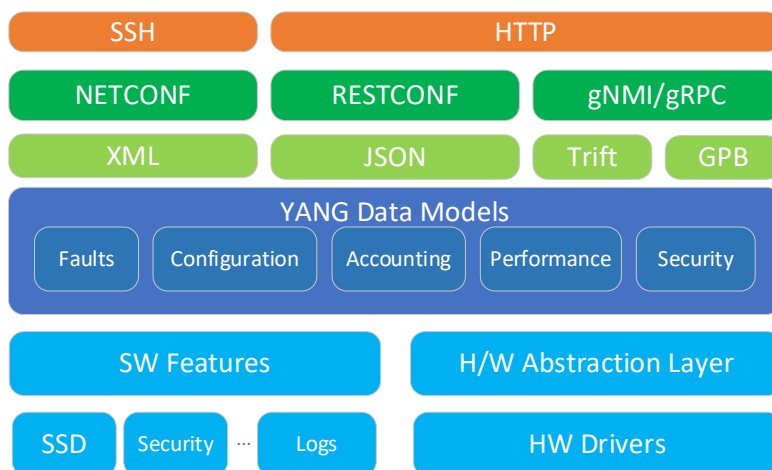
- RESTCONF stands for Representational State Transfer Configuration Protocol. RESTCONF is a REST-like protocol which relies on HTTP protocol and methodology. Request and response data can be represented in XML or JSON format. RESTCONF is described in **Error! Reference source not found..**
- gNMI is a network management protocol developed primarily by Google. gNMI provides the mechanisms to manage the configuration of network devices, and also to retrieve operational data. gNMI typically relies on Thrift or Google Protocol Buffers (GPB) for data encoding and serialization. gNMI is specified in [GNMI-SPEC].

It's outside of the scope of our paper to analyze the technical differences between these protocols. A good example of such analysis between NETCONF and RESTCONF can be found in [CLAISE]. We will however examine the common properties of these protocols to consider their suitability as the replacement for GCP/RCP 1.0.

All of these protocols have been adopted by major network equipment providers and gained strong industry support in both physical and virtualized applications. A rich ecosystem of tooling and test equipment benefits from participation by a wide vendor community. Open source code libraries for the client and the server side are available to accelerate development, lower costs and ensure seamless interoperability.

We don't claim that there are not any differences between these protocols. On the contrary, the design of each protocol allows it to perform certain tasks better while having disadvantages in other areas. They have been developed to solve different issues with a different set of goals. However, the following properties are common to all the protocols. Each protocol supports a superset of the primitives offered by RCP 1.0. All protocols offer security protection features that can be seamlessly integrated into the CableLabs public key infrastructure security defined for R-PHY 1.0. Mirroring a capability of other networking devices, a R-PHY 2.0 compliant RPD could even have the flexibility to simultaneously operate all three protocols on top of common YANG models.

The multi-protocol RPD stack is shown on Figure 3.

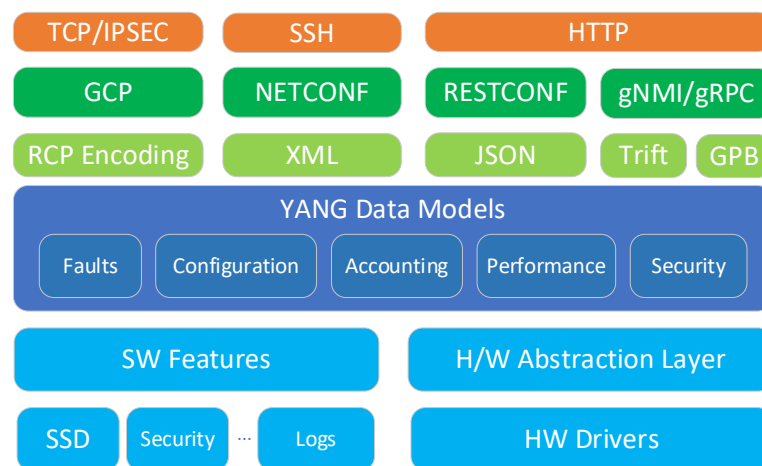


**Figure 3 – RPD 2.0 control protocol stacks.**

The choice of which protocol is enabled on RPDs in any particular deployment could be left to operators and driven by the requirements of the CCAP Cores and other OSS infrastructure systems deployed in their networks. Another choice is to narrow the RPD mandate to just one, carefully selected protocol, to maintain the low development costs and reduce the complexity of the RPD.

### 3.5. 1.0 to 2.0 Transition

By applying the translation methodology explained in the previous sections, the R-PHY specifications can establish a precise correspondence between object classes and individual attributes represented in YANG and in RCP object-TLVs. For example, a definition of each YANG leaf attribute could include a cross-reference to an RCP object-TLV. With such a mapping, a 2.0 RPD can support dual 1.0 and 2.0 control protocols. The protocol stack of such an RPD is shown on Figure 4.



**Figure 4 – RPD 1.0 and 2.0 control protocol stacks.**

It is worth noting that an RPD with a dual protocol stack may not only allow backward compatible operation with R-PHY 1.0 but can also inter-operate with 1.0 and 2.0 CCAP Cores simultaneously. The operators can utilize this feature as a part of their transitional strategy to R-PHY 2.0. It offers the flexibility to selectively and incrementally upgrade CCAP Cores to version 2.0 protocols. Further, certain CCAP Cores supporting narrow functionality, e.g. SCTE 55-2 OOB, may continue to operate in R-PHY 2.0 environment at 1.0 level until they reach end-of-life.

### 3.6. Enabling Automation

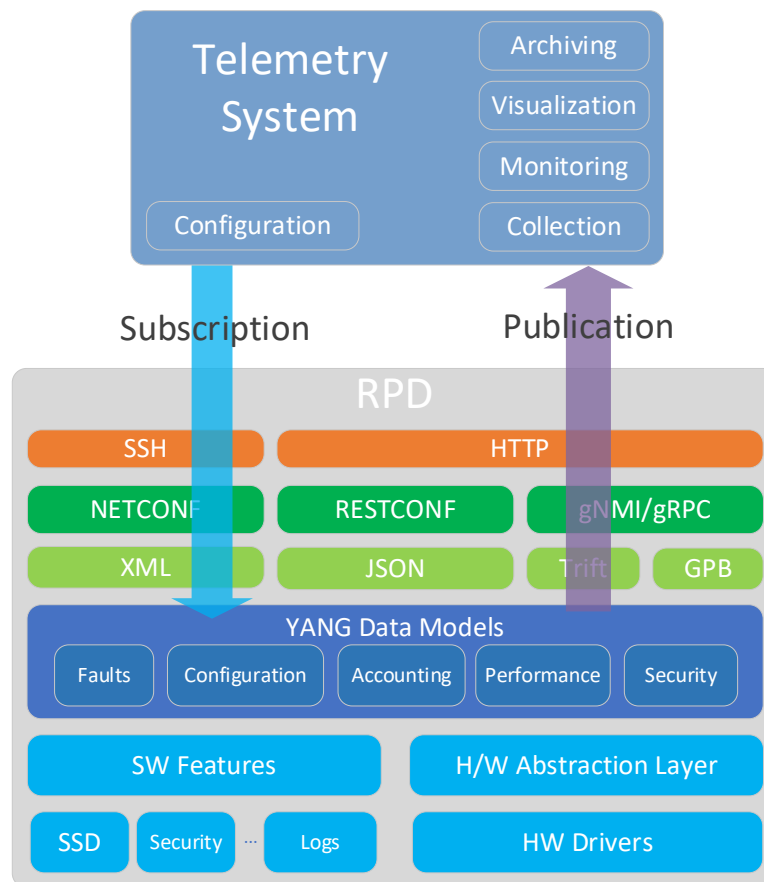
Studies show that networks are evolving faster today than they have in the previous decades while their OPEX and CAPEX are being continually reduced. The key evolutionary drivers are automation and virtualization. An R-PHY control protocol transition to depend on YANG models and widely deployed, standard-based protocols will align the R-PHY with modern cloud-native technologies. It will also help in addressing all of the RCP 1.0 issues explained earlier in the paper. Few automation tools support GCP/RCP protocol. Many existing cloud automation tools are available, and their APIs are YANG based. Thus, transitioning to a 2.0 control protocol will be the necessary step to more easily integrate with cloud-native CCAP Core systems and automated OSS systems. The results will be the enablement of automation, the acceleration of the network evolution and significant reduction of the total cost of ownership for cable operators.

## 4. Model Driven Telemetry

Model-Driven Telemetry (MDT) is a modern technique for monitoring in which operational data is streamed from network devices continuously using a push model. Applications can subscribe to selected elements of YANG data models over a standards protocol such as NETCONF, RESTCONF or gNMI/gRPC. Model driven streaming telemetry allows monitored data to be pushed from the monitored device, e.g. the RPD, to an external collector at a higher frequency than polling, as well as to push data only when a change is recorded. A periodic collection method, when a device pushes data at a defined interval, is better suited to monitoring of frequently changing metrics, e.g. data plane statistical counters. An on-change collection method is a better fit for monitoring infrequently changing data such as state objects, faults or error counters. Through a combination of these methods, MDT provides a highly flexible, efficient communication process for automatic near real-time access to operational data.

In order to stream data from the device the application, or the collector, establishes a subscription to a data set which can be any subset of a device's YANG model. A subscription is a contract between a subscription service and a collector that defines the data set to be pushed and the collection methods. Subscription allows clients to subscribe to modeled data. The device pushes the data to the collector as per agreed contract.

Figure 5 shows how MDT could be integrated with the proposed R-PHY 2.0 protocol stack.



**Figure 5 – Model Driven Telemetry with R-PHY 2.0 Protocol Stack.**

The expanding popularity of MDT is driven by many factors, including its simplicity, the reliance on open standards and YANG models, the broad availability of commercial solutions and Open Source software for all elements of the development ecosystems. For example, well-known Open Source components such as the Apache Kafka messaging bus and the ELK stack (Elasticsearch, Logstash, and Kibana), can be used to build a reliable MDT infrastructure and automated systems for processing and visualization of received information.

The cable industry is already quite familiar with model driven telemetry. MDT originated as a data collection technology for cloud-based infrastructure, but it is also implemented on modern CCAP hardware-based platforms. Several cable operators already deploy MDT collection and monitoring systems in their networks. Including MTD in the R-PHY 2.0 feature set will provide a simple, yet extremely powerful technology for pushing useful metrics from where they are generated to where they are consumed, fitting well into the operators' publish/subscribe (PUB/SUB) model. MTD will become a foundation for modern monitoring of the real-time health of the RPD population as well as of the services it provides.

Finally, we examine a simple example which shows how MTD technology could be applied to monitor a vital metric of the health of the R-PHY data plane. In R-PHY, the user data is transported over L2TPv3 pseudowires. In each consecutively transmitted packet, the L2TPv3 transmitter increments a sequence number embedded the packet header. By examining the continuity of the sequence numbers, the receiver can detect when the network drops packets in-between the transmitter and the receiver. In such a case the receiver (i.e. the RPD) increments a statistical counter of sequence errors for the corresponding pseudowire. Any change to the values of the sequence error counters provides an immediate indication of a potential issue with the health or the performance of the network. The RPD telemetry agent could be configured to monitor changes to the values of the sequence error counters and stream the counters' values to the MTD collector whenever they change. An application within the MTD collection system could then in real-time analyze the received data and alert a network manager about on-going network problems.

## **5. R-PHY with Remote Upstream Scheduler**

The R-PHY remote upstream scheduler is an architectural option that moves the real-time DOCSIS upstream scheduling function together with the PHY element to the RPD. It is suitable for providing low latency DOCSIS transport over long distance R-PHY deployments.

The location of the upstream scheduler has been part of the R-PHY architecture consideration since the very beginning of the R-PHY development. The optimum location choice depends on both business and technical reasons.

For R-PHY 1.0, the primary goal is to enable DCA by replacing the analog optical link between the CCAP and the Node with a digital link. Just with this initial step, cable operators would be able to get better SNR performance, pull the fiber deeper, rebuild the plant and cut a large service group into much smaller ones. All these can be achieved by simply moving the PHY element out of the CCAP Core, while keeping all MAC elements including the DOCSIS upstream scheduler centralized. This also allows the operators to leverage the existing CCAP MAC functions to simultaneously support both integrated PHY and Remote PHY for a quick and smooth transition to DCA.

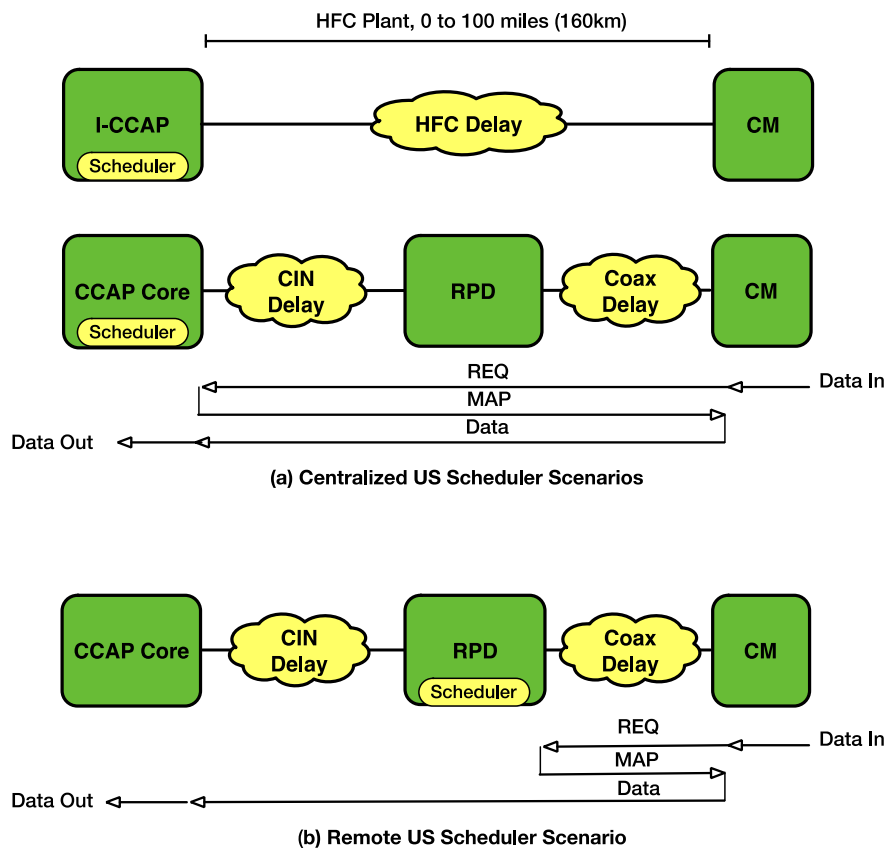
The main technical reason for applying centralized upstream scheduling to R-PHY 1.0 is because with a 2 msec MAP interval, the CIN delay is not a dominant latency factor within the 100-mile I-CCAP HFC

reference range. In this case R-PHY 1.0 is equivalent to an I-CCAP in terms of the upstream request-grant (REQ-GNT) latency.

As the network keeps transitioning to DCA, there are, however, reported cases where the CIN is stretched beyond the 100-mile mark, for reasons such as hub-site 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 REQ-GNT protocol is being tightened to shorter MAP intervals, such as 1 millisecond on DOCSIS 3.1 OFDMA channels. In such circumstances, the CIN delay could be exposed as a significant factor affecting the REQ-GNT latency. This problem can be solved by simply moving the REQ-GNT handling to the RPD, which will effectively cut the CIN out of the REQ-GNT loop as shown in the Figure 6 below.

The remote upstream scheduler will be proposed as a R-PHY 2.0 feature for low latency support. The specification will focus on the remote upstream scheduling interface definition to allow the CCAP Core to work with remote upstream schedulers from different RPD vendors. The remote upstream scheduling management interface will be a data model driven, taking advantage of the new R-PHY 2.0 control plane infrastructure.

The overall remote upstream scheduling definition will provide R-PHY 2.0 with the flexibility for load balancing the upstream scheduling tasks between the CCAP Core and the RPD, enabling backward compatibility with 1.0 RPDs, and ultimately the ability to maximize both the centralized and distributed computation resource to achieve low-latency at high system efficiency.



**Figure 6 – Centralized vs. remote R-PHY upstream scheduling options for R-PHY1.0 and R-PHY 2.0.**

Additional details about the Remote Scheduler can be found in [RemoteScheduler].

## 6. Data Plane Improvements

Since the early days of M-CMTS, the MAC and PHY split has been enabled through the application of L2TPv3 tunneling technology. The (DEPI) tunneling scheme utilized L2TPv3 over IPv4 or IPv6 as a simple, lightweight and standards-based encapsulation to enable scalable connectivity between the CMTS and the edge QAM device.

L2TPv3 was utilized again with the advent of Remote PHY, to enable a MAC and PHY split. DEPI was used on the downstream and the upstream direction (called Upstream External PHY Interface or UEPI) was added. In the Remote PHY architecture, the L2TPv3 UEPI tunnels are unicast (point to point) from the RPD to the CCAP Core. In the downstream direction, the L2TPv3 DEPI tunnels are either unicast from the Core to a single RPD or multicast to multiple RPDs. Multicast DEPI tunnels permit an efficient allocation of CCAP Core resources across many DOCSIS service groups and are ideal when adapting existing centralized hardware CCAP devices to a Remote PHY deployment. The network that DEPI and UEPI tunnels transit is called the CIN (Converged Interconnect Network).

While the existing L2TPv3 DEPI and UEPI tunneling schemes have served both M-CMTS and Remote PHY well, as cable networks and CCAP software evolve it might be prudent to re-examine the tunneling architecture of a MAC/PHY split in cable. The current architecture leaves a few things to be desired:

- It is difficult to traffic engineer IP tunnels without resorting to another, additional encapsulation. Specifically, traffic engineering refers to redistributing traffic loads across different paths.
- ECMP (Equal Cost Multi Path) load balancing in the CIN can be a challenge, primarily due to the lack of decipherable entropy in the payload of the packet. ECMP is a hop-by-hop algorithmic load balancing mechanism that depends on sufficient input to the algorithm, known as ‘entropy’, to make an efficient decision. Because DEPI and UEPI tunnels carry encrypted DOCSIS traffic and have a common set of IP addresses, a standard router has limited visibility in to how to best make a load balancing decision.
- A high caliber network architecture is required. It is incumbent that the CIN be engineered akin to circuits, not paths. This is a packet transport network, not an internet routing network. Remote PHY architecture requires that packets arrive in order, no high priority packets be dropped and symmetric latency (round trip times) be maintained.
- L2TPv3, despite its versatility, is still a niche tunneling protocol in the industry. Service provider networks have generally embraced Multi-Protocol Label Switching (MPLS) as the tunneling technology of choice. If you look at the protocol diagrams for R-PHY 1.0, it allowed for an expansion to include MPLS.

The state of cable access architecture is evolving quickly to encompass software, cloud native technology and multi-modal access methods. Software implementations of CMTS infrastructure permit a horizontal scaling of DOCSIS resources, which eliminates the need for IP multicast in DEPI tunnels for DOCSIS. Many cable operators would like to build out one CIN or Ethernet aggregation network, of which DOCSIS technology is but one method for last mile connectivity. Also, service provider control and data planes, automation methods, telemetry retrieval and system programmability have all evolved significantly since M-CMTS and Remote PHY were first proposed.

This confluence of events and technology means it’s a good time to revisit the CIN architecture and the way in which the Remote PHY system interconnects the CCAP Core with RPDs. This paper explores two options. In both cases, MPLS technology will play a much larger role in the transport of Packet Streaming Protocol (PSP). Option 1 is to simply encapsulate the existing DEPI and UEPI unicast into an MPLS LSP (Label Switch Path). Option 2 is to eliminate the L2TPv3 tunneling layer altogether, and tunnel PSP natively over MPLS.

### **Option 1: PSP over L2TPv3 over MPLS**

In this option, the existing Remote PHY DEPI and UEPI tunnels are further encapsulated into MPLS, typically by the first hop router in the CIN. Once in an MPLS LSP, Remote PHY traffic can be merged and integrated into a multi-purpose CIN.


This option has the benefit of maintaining backward compatibility with existing CCAP and RPDs, including the L2TPv3 control channel. While this is the easiest option from a Remote PHY architecture change perspective, and it pushes any MPLS integration work to a CIN engineering exercise, it is suboptimal from a life-of-a-packet perspective. It is generally unadvisable to deploy networks with multiple layers of encapsulation, especially when the final encapsulation (in this case MPLS) is capable of carrying the ultimate payload (in this case Remote PHY PSP). L2TPv3 in this case is simply redundant and adds unnecessary complexity, MTU size, and administrative complications into the architecture.



## Option 2: PSP over MPLS

In this option, R-PHY PSP is natively transported over MPLS. This is a unified tunneling approach, in that one tunneling mechanism transports not only Remote PHY traffic but also any other traffic the CIN may be called upon to transport. MPLS provides a payload agnostic method of transporting any type of data across a packet switched infrastructure, and Remote PHY will be one of any number of services.

Benefit	Drawback
RPD traffic is under policy control with XTC (PCE). TE is a function of network management.	RPD and CCAP participate in control plane – must engineer correctly for scale and resiliency
MPLS enables a multi-purpose network	MPLS skillset development needed
MPLS enables scale and automation, alignment with progressive SP network directions	More software running on the RPD and CCAP (control plane only – ISIS, BGP, PCEP)
CIN is free of any remote phy state (labels only)	Software and standards changes for Remote PHY architecture



**Figure 7 – MPLS Transport for R-PHY.**

It is worth noting that the RPD has legacy DOCSIS traffic that does not use PSP as well as video traffic that is based on an MPEG-TS. All this legacy traffic could also be placed natively on an MPLS infrastructure or over-laid with old encapsulation over new encapsulation.

Key to the use of MPLS in this architecture is creating the equivalent of a circuit for PSP packet transport. In Remote PHY it is critical that between the CCAP Core and the RPD, no packets should be dropped, no packets have asymmetric latency, and no packets arrive out of order. Traffic Engineering (TE) permits this type of network to be built by specifying path or path constraints that MPLS encapsulated traffic must follow. Modern approaches to TE, such as Segment Routing, permit a lightweight and scalable approach to delivering the type of network Remote PHY performs best in.

## 7. Other Functional Improvements

In addition to the four architectural options outlined above, Remote PHY 2.0 as described can provide a foundation for a much broader set of functional improvements. Here we list several such options without describing them in detail.

- **Extended Spectrum DOCSIS (ESD)**, a part of DOCSIS 4.0. The process of standardization of ESD has only just begun at CableLabs but it will undoubtedly require changes to R-PHY specifications.

- **NetFlow** agent in the RPD. NetFlow is a valuable tool which can help in debugging end-to-end data plane issues.
- **Broadband Digital Forward and Broadband Digital Return.** Building on Moore's law progress these techniques push existing R-PHY features, such as Narrowband Digital Return (NDR) and Narrowband Digital Forward (NDF) into digitization of wider spectrum blocks. DCA effectively replaces the legacy technologies.
- **Advanced Power Management** of the RPD system and its RF module may allow for significant reduction of the power consumption of the HFC network.

R-PHY 2.0 can be also used to more completely specify those data unit formats which are kept as vendor proprietary in R-PHY 1.0 specifications.

## Conclusion

R-PHY 1.0 provides a valuable addition to the toolset available to operators as they continue to extend service offerings and provide ever increasing bandwidth while reducing capital and operational costs. The paper has described a number of issues faced by the 1.0 version which may limit the utility of the R-PHY system going forward.

The paper describes a menu of potential architectural improvements for Remote PHY technology and describes four of these in some detail. Individually, each one of the proposed options solves a different problem and offers valuable business and operational benefits to the cable operators. The paper demonstrates that each one of these options is worthy of the consideration by the cable operators in its own right and that when combined they can provide value which is greater than the sum of the parts. Taken together, these technical improvements constitute a new generation, Remote PHY 2.0.

# Acknowledgments

The authors would like to sincerely thank John Chapman and Gerry White for their valuable comments and contributions to the paper.

## Abbreviations

BGP	Border Gateway Protocol
CAPEX	Capital Expenditure
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
DCA	Distributed CCAP Architectures
DEPI	Downstream External PHY Interface
DOCSIS	Data over Cable System Interface Specification
ECMP	Equal-Cost Multi-Path
ESD	Extended Spectrum DOCSIS
FCAPS	Faults, Configuration, Accounting (or Administration), Performance, Security
FMA	Flexible MAC Architecture
HFC	Hybrid Fiber Coax
I-CCAP	Integrated CCAP
IoT	Internet of Things
ISYS	Intermediate System to Intermediate System
GCP	Generic Control Plane
GPB	Google Protocol Buffers
L2TP	Layer 2 Transport Protocol
L2TPv3	Layer 2 Transport Protocol version 3
LSP	Label Switching Path
MDT	Model Driven Telemetry
MPLS	Multi-Protocol Label Switching
msec	millisecond
NETCONF	Network Configuration Protocol
OPEX	Operational Expenditure
PCE	Path Computation Element
PSP	Packet Streaming Protocol
RCP	R-PHY Control Protocol
RPD	Remote PHY Device
RR	Route Reflector
R-PHY	Remote PHY
SCTE	Society of Cable Telecommunications Engineers
SR	Segment Routing
SSH	Secure Shell
TE	Traffic Engineering
TLS	Transport Layer Security
UEPI	Upstream External PHY Interface
XTC	XR Traffic Controller
YANG	Yet Another Next Generation

# Bibliography & References

- [CLAISE] <https://www.claise.be/2017/10/netconf-versus-restconf-capability-comparisons-for-data-model-driven-management-2/>, Benoit Claise
- [GCP] DOCSIS DCA-MHAv2, Generic Control Plane Specification  
CM-SP-GCP-I04-180509, CableLabs
- [GNMI-SPEC] "gRPC Network Management Interface (gNMI) v0.6.0", OpenConfig operator  
working group  
<https://github.com/openconfig/reference/blob/master/rpc/gnmi/gnmi-specification.md>
- [MULPI] DOCSIS 3.1 MAC and Upper Layer Protocols Interface  
Specification, CM-SP-MULPIv3.1-I18-190422, CableLabs
- [NetProg] Network Programmability with YANG, Benoit Claise, Joe Clarke, Jan Lindblad,  
ISBN-13: 978-0135180396
- [RemoteScheduler] SCTE-Expo 2019, R-PHY with Remote Scheduler, Tong Liu
- [RFC3931] <https://tools.ietf.org/html/rfc3931>
- [RFC6020] <https://tools.ietf.org/html/rfc6020>
- [RFC6241] <https://tools.ietf.org/html/rfc6241>
- [RFC7950] <https://tools.ietf.org/html/rfc7950>
- [RFC8040] <https://tools.ietf.org/html/rfc8040>
- [RPHY] DOCSIS DCA-MHAv2, Remote PHY Specification, CM-SP-R-PHY-I12-  
190307, CableLabs
- [RPHYOSS] DOCSIS DCA-MHAv2, Remote PHY OSS Interface Specification,  
CM-SP-R-OSSI-I12-190510, CableLabs
- [RPHYTR] DOCSIS DCA - MHAv2, Modular Headend Architecture v2 Technical Report,  
CM-TR-MHAv2-V01-150615, CableLabs
- [R-DEPI] DOCSIS DCA-MHAv2, Remote Downstream External PHY Interface  
Specification, CM-SP-R-DEPI-I12-190307, CableLabs
- [R-UEPI] DOCSIS DCA-MHAv2, Remote Upstream External PHY Interface  
Specification, CM-SP-R-UEPI-I10-190307, CableLabs

# Capitalizing On The Evolved Communications Experience

A Technical Paper prepared for SCTE•ISBE by

**Kjell Johansson**  
Director, Technology Transformation  
Ericsson Inc  
6300 Legacy Drive, Plano, Texas, 75024  
+1 404-936-3889  
kjell.m.johansson@ericsson.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	4
Content .....	5
1. Consumer communication trends and use cases.....	5
1.1. Companion Devices .....	5
1.2. Personal Assistant / Smart Speaker.....	6
1.3. Audio / Video Doorbells .....	6
1.4. Improved communications at venues .....	7
1.5. Fixed Wireless Access (FWA) with consumer Voice over IP (VoIP).....	8
2. Business communication trends and use cases.....	8
2.1. Enhanced voice call and WebRTC enrichment.....	8
2.2. VoLTE for UC - Integration of VoLTE with Enterprise UC .....	9
2.3. Voice for Internet Of Things (IOT).....	11
2.3.1. Introduction.....	11
2.3.2. Connected Cars .....	11
2.3.3. Fleet Management.....	12
2.3.4. Elevators and Alarm Panels .....	12
2.4. Augmented Reality (AR).....	13
2.5. Smart Speaker for business.....	14
3. Specific Technology Dependencies.....	14
3.1. Consumer Communication Dependencies .....	15
3.2. Business Communication Dependencies .....	16
4. Common technology dependencies.....	17
4.1. Enablers Explained.....	17
4.1.1. NFV-I Orchestration and Cloud Execution Environment (CEE) .....	17
4.1.2. Automation and Analytics.....	18
4.1.3. Service Enablement.....	19
4.1.4. Container Distribution.....	19
4.1.5. Dynamic Orchestration.....	19
4.1.6. Cloud SDN.....	20
4.1.7. Distributed Cloud .....	20
4.2. Service provider scale, ambition and approach .....	20
4.3. Advanced Use cases common enabler dependencies.....	22
Conclusion .....	22
Abbreviations.....	23
Bibliography & References .....	25

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Connected wearables shipments by device category (World 2017–2023).....	5
Figure 2 - NNI-based VoLTE for UC Source: Ericsson .....	10
Figure 3 - Network architecture for remote manufacturing .....	13
Figure 4 - Dependencies on feature functionalities for the consumer trends and use cases .....	15
Figure 5 - Dependencies on feature functionalities for the consumer trends and use cases .....	16
Figure 6 - Reference stack blueprint .....	17
Figure 7 - Segmenting Service Providers across common enablers .....	21
Figure 8 - Common enablers across advanced use cases.....	22

# Introduction

Almost 20 years after the first communication service providers began talking about “all IP” as a vision for network evolution, we have reached a stage where a vast majority of services are being delivered according to an architecture where IP is truly used end-to-end, including in the end-user devices, and at scale.

About 10 years ago, Android was introduced, helping to reduce the number of mainstream platforms in popular devices – most notably smart phones, tablets and TVs. Android quickly became an application development ecosystem, and a very open one, easy to integrate into various types of networks.

About 5 years ago, the industry took a few first vital steps towards communication services delivered from private clouds, which can be defined as using Network Function Virtualization, to deliver software only applications onto an independently provided cloud platform. These are just the first steps on a journey that will bring greatly enhanced automation, elasticity and flexibility in deployment, based on improved analytics and agility, as well as a much-needed performance improvement.

The upcoming introduction of 5G for wireless, will bring enterprises and consumers onto a common and shared all-IP network, providing both an internet and communication service experience that is much enhanced. This will be possible because of a highly distributed cloud, orchestrated to provide both enterprises and consumer segments and use cases, with the look and feel of running on their own (dedicated) network.

In this paper, we will be examining a set of communication use cases and trends which all can enhance and differentiate the end user experience, improve revenues, and reduce churn. Each of these have dependencies on one or more enablers from the four technology waves above, in order to be able to scale and multiply. Our objective is to drill down on these technical dependencies, with the objective to “show the path” to realizing them.

We have elected not to cover regulatory opportunities in this document, mainly because user experience is at the center of the theme, and user experience for mission critical and government services such as Wireless Priority Service (WPS), are driven by different factors than for consumers and non-government enterprises. It is understood that the regulatory opportunities right now are indeed very exciting and are certainly an increasingly important part of what the communication services industry needs to address moving forward. Perhaps, deserving of its own document.

It should be noted that this document is a technology inventory of trends, use cases, and dependencies. Not all of them apply to any one reader. Each vendor and each service provider will need to decide which enablers are relevant, and should also consider additional trends, use cases and enablers, applied to their specific situation.



# Content

## 1. Consumer communication trends and use cases

### 1.1. Companion Devices

During most of this decade, smart phones and home phones, have been joined by other popular device types such as tablets and wearables, and smart-home devices, into what is a growing ecosystem for communication and entertainment, providing enhanced features to the existing services. The new communication features are

- Multi-Device, which is the ability to consume the service, on multiple devices including companion devices
- Multi-Persona, which is the ability to use multiple identities such as phone numbers, on a single device, for example one work identity and one personal identity
- Multi-User, which is the ability to share a common group identity, e.g. a family number, in addition to the personal identities of each group member

These capabilities greatly enhance the end user experience of the communication service and can in many cases increase usage of the service. Perhaps more importantly, service providers can find ways to differentiate, i.e. create a “stickiness”, and thereby reduce the risk of churn.

Million	2017	2018	2019	2020	2021	2022	2023
Fitness & activity trackers	58.0	65.0	71.5	78.7	82.6	86.7	88.9
Smartwatches	29.5	45.5	56.9	71.1	85.3	102.4	117.7
Smart glasses & HMDs	0.2	1.5	3.2	5.1	7.7	10.0	11.9
Medical devices & mPERS	1.3	1.8	3.0	4.9	5.4	6.1	6.9
Others	2.0	3.0	4.5	6.5	9.0	11.0	13.0
<b>Total</b>	<b>91.0</b>	<b>116.8</b>	<b>139.1</b>	<b>166.2</b>	<b>189.9</b>	<b>216.2</b>	<b>238.5</b>
Cellular attach rate	6.7 %	13.9 %	17.2 %	20.4 %	22.8 %	25.8 %	28.4 %
Cellular device shipments	6.1	16.2	24.0	34.0	43.2	55.8	67.7

**Figure 1 - Connected wearables shipments by device category (World 2017–2023).**

Source: Connected Wearables, M2M Research Series 2019, Berg Insights

This trend will continue, driven by for example a projected strong growth in wearables shipped globally (figure 1), in particular wearables that are suitable for voice, like smart watches.

## **1.2. Personal Assistant / Smart Speaker**

After its introduction just a few years ago, the smart speaker or personal assistant has been tremendously successful, particularly in North America. This device represents a shift to smart home, and a completely new type of user experience / life style. This device type integrates to the smart device ecosystem in the house, but it also very quickly became integrated to the smart phone and the home phone, for making and receiving calls, sending text messages, etc.

Smart Speakers come with their own phone number, so that you can communicate with telephony devices on public networks but is otherwise lacking in integration to the communication services providers.

Some communication service providers and smart speaker providers are now taking steps to further integrate their services, like in the example of Vodafone and Amazon.

[\(Reference 1\)](#)

This type of tighter integration provides several benefits:

- it removes the need for a separate new phone number, associated with the smart speaker
- other devices in the home can become companions to the smart speaker, and these can for example have simultaneous ringing with the smart speaker when the family number is called
- it can increase usage of all the services in the smart home
- it can reduce churn

As smart homes and smart speakers continue to evolve, we will certainly find new use cases for integration to communication services. Revenue connected to smart speakers is expected to grow at an average of 2B USD per year annually, and reach 12 B USD annually in 2023

[\(Reference 2\)](#)

It will be important to find added value for both the communication service provider and the smart home device vendors, so not all use cases will be sustainable.

## **1.3. Audio / Video Doorbells**

Another recent and emerging communication device in the North American smart home is the connected audio / video door bell, with a similar growth projection as the smart speaker.

[\(Reference 3\)](#)

The communication aspects however, are quite different. Although many of these doorbells will not be by the front door, the primary value of this category of device is security of the home owner and their family. Also, the identity of at least one of the participants in a typical “doorbell call” is potentially unknown and inherently not trusted.

Perhaps therefore there really has not been any notable integration between the doorbell and the communication service providers. The doorbell provider typically provides a smart phone application, and the communication is handled over a pure internet channel between the smart phone app and the doorbell, remote controlled by the home owner.

However, in the broader context of a smart home, a market need for communication integration to this type of device could very well develop over the next few years, driven by for example the following needs;

- **Video Quality.** Video quality is one of the most important characteristics of the doorbell as a security device, and the improvement of video quality in the doorbells we expect to see in the next several years, may drive the need for a higher video Grade Of Service (GOS), which in turn will drive the need for traffic prioritization, which communication service providers specialize in.
- **Security itself.** If the communication service provider also provides the security system and the video doorbell, then there are several other enhancements that can be made based on the trust relationship between the home owner and the service provider. An example of this is the possibility of a “grey list” of people with conditional access to the home, through for example biometrics, or digital credentials (bar codes).
- Another interesting ability which has already been applied, is to integrate the doorbell with a smart speaker (see section 1.2), which was showcased at Consumer Electronics Show this year ([Reference 4](#)). This type of integration / companionship could create additional value for communication services.
- Emergency services, which can be needed at the outdoor camera spot, by anyone, is a third potential use case that can be added. The communication provider is many times already providing this for the address of the doorbell, and it can be triggered by a button on the doorbell, or from the home owner smart phone.

Overall, the video doorbell, while not yet communication service provider integrated, deserves keeping an eye on, and behind the smart speaker it is perhaps the second most interesting point of integration, into the smart home, from a communication service perspective, looking forward.

## **1.4. Improved communications at venues**

Many service providers, including cable operators, are today providing broadband, Wi-Fi and cellular coverage at venues, which have traditionally suffered from poor connectivity and bandwidth. For this section, these service providers will be called “wireless” service providers. This is done as a business-to-business arrangement with the landlord/venue.

Whether the context is sports, music or other entertainment events, there is a growing trend of making use of this improved coverage, by opening a channel to the audience to improve their experience of the event. This can include social networking, voice, chat, group messaging, real time statistics, replays, back-stage contexts, alternative camera angles, different monetization strategies can be explored, including advertising.

Depending on who the venue, and content partners are, two business models are feasible or can be combined:

- **Business-to-Business-to-Consumer (B2B2C)**  
The wireless service provider participates in a joint model with the venue/content owner, to facilitate a complete set of services to the audience.
- **Business-to-Consumer (B2C)**

The wireless service provider can provide their communication, mobility and internet connectivity services directly to the audience member. This is sometimes referred to as a retail service.

B2C enriched / advanced messaging – an add on to Rich Communication Suite (RCS) messaging, can be leveraged by the wireless service provider, to provide an advertising channel to the audience member. Note that this is sometimes referred to as RCS Business Messaging.

For more on RCS Business Messaging see ([Reference 5](#))

## **1.5. Fixed Wireless Access (FWA) with consumer Voice over IP (VoIP)**

As the industry is rolling out FWA – a technology on the way to 5G introduction in mobile networks, there are also initiatives to provide this service using unlicensed spectrum in the 3.5 MHz frequency band, so called Citizens Broadband Radio Service (CBRS).

So far, the cable service providers have not engaged on this opportunity, but some of the wireline operators are exploring it. The reason is that there is no business case for replacing existing cable access, and it is more feasible to compete against Regional Bell Operating Companies (RBOCs), using the traditional cable access.

However, as the technology evolves, and as the industry spectrum situation becomes clearer across North America, FWA may again become feasible as an alternative to cable access, within territory of each cable operator.

If FWA becomes viable, an immediate question and concern will be whether the service provider should offer home phone service. Key questions and concerns will be:

- Voice and Internet should be supported in the same chipset, in the FWA modem, to improve business case over current situation
- Should voice be standardized, and aligned with the VoLTE service “without mobility”, i.e. equivalent to Multi Media Telephony (MMTel) standards, and therefore be less aligned towards the voice service provided by the current Multiple System Operator (MSO) covered homes?
- What should be the value-added voice related services provided as part of the Quad Play?

## **2. Business communication trends and use cases**

### **2.1. Enhanced voice call and WebRTC enrichment**

Several service providers are looking to provide their business customers in the small and medium segment, with ways of enhancing existing voice services, which is being provided to the business by the service providers. These existing voice services being enhanced can be

- a. mobile subscriptions for the employees, with 2G/3G Circuit Switched (CS) based voice, or Voice Over Long Term Evolution (VoLTE) for capable devices
- b. Unified Communication (UC) service with support for VoIP on desk phones, computers, and potentially also to those same employee VoLTE enabled mobile devices (see section 2.2)

- c. Customer support, for terminating traffic to the business from any device that is capable of Web RTC (Web based Real Time Communication).

The enhancements can be of two types:

1. For service a and b above, the service provider can provide pre-call, during-call, and post-call enrichment of the call, which includes a broad suite of messaging, photo and file sharing, group features. This generally requires an upgrade or a re-configuration of the core network, to support the necessary triggering functions, and an application server upgrade, or re-configuration to support the specific use cases.

The Global System Mobile (GSM) Association (GSMA), has defined these services, based on RCS 2.0 ([Reference 6](#)), which is broadly supported among service providers.

2. For service c, co-browsing, can be introduced to enrich the voice call end user customer experience, across mobile devices, tablets and laptops. Co-browsing means that the customer support person and the customer can view the same web page, including the cursor/pointer, during the call. If needed, a co-edit capability can be added for e.g. jointly filling in forms.

For a quick demo of co-browsing, see ([Reference 7](#))

As part of Hypertext Markup Language revision 5 (HTML5), the WebRTC standard provides real time communication between the customer support agent and the customer, including the use of a data channel, which can be used for the co-browsing, and other enrichment to the voice/video call. For more on this, please see ([Reference 8](#)), and the enabler section 3 (Technology Enablers).

There is a lot to gain for a business by providing “in context” communication whether it is for the employees, or for the customers of the business. Both RCS enriched calling and the customer service with co-browsing solutions provides sustained relevance, can increase usage of existing services, and improve the ability to retain the business customer.

## **2.2. VoLTE for UC - Integration of VoLTE with Enterprise UC**

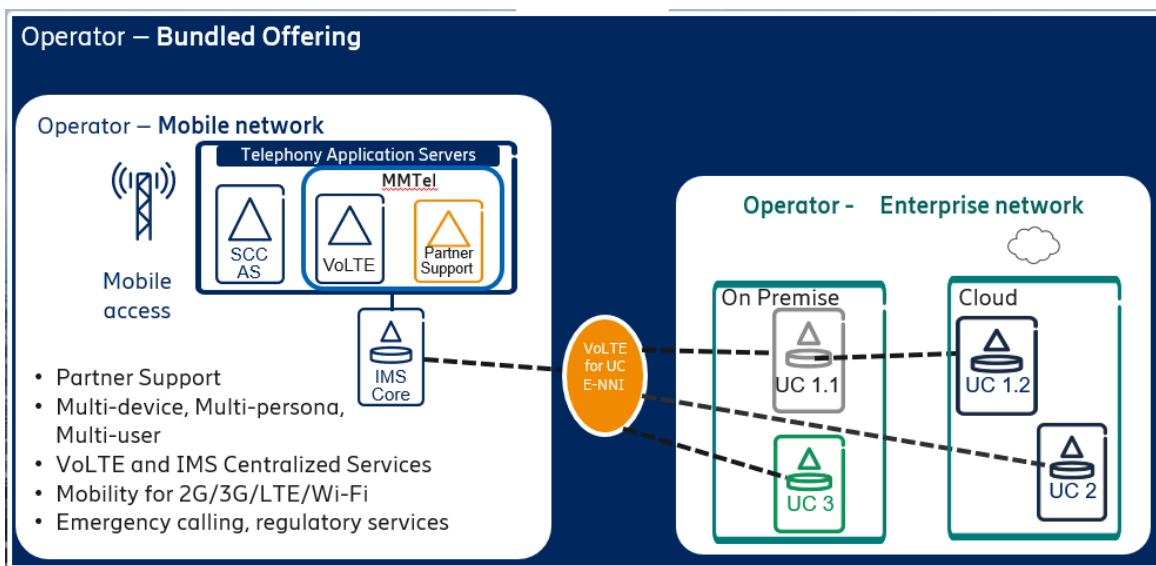
For over a decade, service providers have been looking for a cost-efficient way of adding mobility to existing enterprise voice services, especially in the small to medium segment of businesses, where it is possible to bundle mobile subscriptions for the employees, into the overall business with the business.

This has proven difficult for the following main reasons:

- 2G/3G networks were lacking the standards to provide an integration to Business Voice and were limited to the Primary Rate interfaces between a Mobile Switching Center (MSC) and on-premise Private Branch eXchange (PBX), which did not support making the PBX service available on the mobile. A few “point solutions” did exist, most of which were not successful, mostly due to a poor end user experience.
- The IP Multimedia Subsystem (IMS)-standards had more support for hybrid services, but were aiming to support new services, as “add-ons” or extensions to a base IMS application, through so called “chaining” over the IMS Service Control (ISC) interface, between the application server and the IMS-core.

- Business VoIP, and VoLTE have not been standardized to complement each other, but rather have large overlaps in functionality. Subsequently, when you use ISC-chaining, interactions get complicate.
- ISC-chaining does not support having multiple Enterprise UC solutions integrated with the same IMS-core
- Some of the UC solutions are only available as cloud services (private, or public), and it is difficult to connect to those over the ISC-interface.

There are now new ways to possibly solve these issues, based on separating the two services into two IMS networks (could potentially still be served by one IMS), and handle this the same way we do when we interoperate between two VoIP networks, with a Network-to-Network Interface (NNI), and a Network Session Border Controller (N-SBC). The UC solution does not have to be IMS-based but can be based on the predecessor architecture – soft switching but needs to support Session Initiated Protocol (SIP).



**Figure 2 - NNI-based VoLTE for UC Source: Ericsson**

Initial trials and a few production deployments have shown very promising results. In the service provider, there are often two different groups operating the UC service and the VoLTE service.

To summarize the market potential of VoLTE for UC:

- Provide a value added “UC on the road”, on the mobile device, as a potential up-sell
- Reduce dependencies and complexities by splitting the solution into two domains, each with its own SIP-network, and thereby creating a “separation of concerns”.
- Allow independent evolution of each service
- Solve any remaining interactions by negotiating / iterating a “new” NNI or Application Programming Interface (API)
- Can potentially also be used to collaborate between service providers where one provides business services and one VoLTE services

It is important for service providers to grow both their enterprise business and consumer business separately, and to find ways to do so in an efficient way.

As the service provider network expands and starts to serve a larger number of constituents or “tenants”, such as enterprises, Mobile Virtual Network Operators (MVNOs), wholesale partners, the more modular and distinct each piece of the network (network slice) is, the easier it will to continue scaling the business.

## **2.3. Voice for Internet Of Things (IOT)**

### **2.3.1. Introduction**

The Internet Of Things represents a tremendous growth segment for several different types of service providers, including communication service providers. As an industry, we have all been engaged for several years, we have solutions that are live, and more to follow in the next 12 months.

In parallel, the industry business blueprint is still under construction, as various actors are trying to jointly build an ecosystem, determine our individual roles, how to be successful within that role, and more importantly how to stretch that success to the use cases and business.

In this section we will discuss some common themes and experiences of this first generation IOT, what services providers have done in terms of communication services. In the next couple of sections, we will investigate particularly interesting use cases with Voice for IOT Enterprise, and how service providers can sustain and grow “share of overall value”.

We will intentionally skip the 2G and 3G based IOT-solutions, which at the time were labeled “Machine-to-Machine” (M2M). These were almost entirely sensor based and have evolved into what is now called “narrowband IOT”. These devices are not voice capable and there is typically no need for voice at all.

Please also note that “smart home” developments are already covered in section 1.2 and 1.3 of this document, and are considered consumer communication services, so will not be covered further here.

### **2.3.2. Connected Cars**

In this industry, the large global car Other Equipment Manufacturers (OEMs) are at the top of the value chain, and are looking for a standard set of capabilities – the very basic one being telematics / engine data channel - which can be deployed globally to support their network of factories and dealerships, including the following expected from communication service providers

- Concierge voice service for the vehicle and anyone in it  
The car is equipped with microphone, speaker and a concierge button, which can be pushed to get any help such as road side assistance or help with directions. This is typically a subscribed service.
- Emergency Service (eCall)  
An eCall can be initiated either by pushing a separate emergency button, or through the concierge service. It can also be initiated automatically, if the car is able to detect a crash. In either case, the call will be routed to a local Public Safety Answering Point (PSAP).
- Terminating call screening  
Callback of any kind to the car is only allowed from the PSAP and the Concierge service

Voice traffic is very low to the vehicles, and besides the three features above, there are a lot of features on traditional mobile soft switches, and IMS-cores that are not used.

The solution must be scaled down from a feature and capacity point of view and scaled up from a global reach and roaming point of view.

OEM (Other Equipment Manufacturer), MNO (Mobile Network Operators – providing Radio Access Network (RAN) in the different global regions), the Mobile Virtual Network Operator (MVNO) – the Service Provider, and perhaps a third-party host for some shared cloud-based capabilities (multi-tenant) must come together to make this service work everywhere.

Part of the requirements are regulatory, which means they represent an initial investment needed to get into the business.

In addition to the communication services, the Mobile Virtual Network Operator (MVNO) can take part in a retail service, where backseat passengers can purchase on-demand content from the MVNO. This is not a communication service, but a potential upside in the business case.

Worth mentioning is that there are alternative business models for the basic connected car above, where the driver builds their own ecosystem based on the open interface available as a standard on vehicles manufactured in the IOT-era. The service providers can potentially also address this space, but in vastly different (e.g. lower cost) ways.

### ***2.3.3. Fleet Management***

If we take a step northbound in the value chain, communication services providers can participate in Fleet Management, which is the scheduling, dispatch, service delivery, and payment part of small/medium businesses fleets of vehicles, which is a business towards the Enterprise owning the fleet, and not the car OEM. The Service Providers may provide all or just a few of the following:

- Connectivity and communication kits for each vehicle
- telematics
- geo-location, navigation
- fleet data collection
- dispatch / Push-To-Talk (PTT) – communication
- business messaging - communication
- machine learning

### ***2.3.4. Elevators and Alarm Panels***

From a functional point of view, this category of devices operates very similar to the connected car. A speaker and a microphone are present in the elevator or at an alarm panel by an un-manned door / gate. In an emergency, a single button call to a security officer / central location is made. These devices are typically of the type CATEGORY M1 (CAT-M1).

For a service provider point of view, this service is quite local to the geography where they are already present. Traffic is even lower from these devices than from the vehicles, so for any existing service provider voice infrastructure that is IP-based (IMS or soft switch), it may make sense to reuse it, if the devices supports the access network (Wi-Fi, GSM, CDMA 1x, LTE).

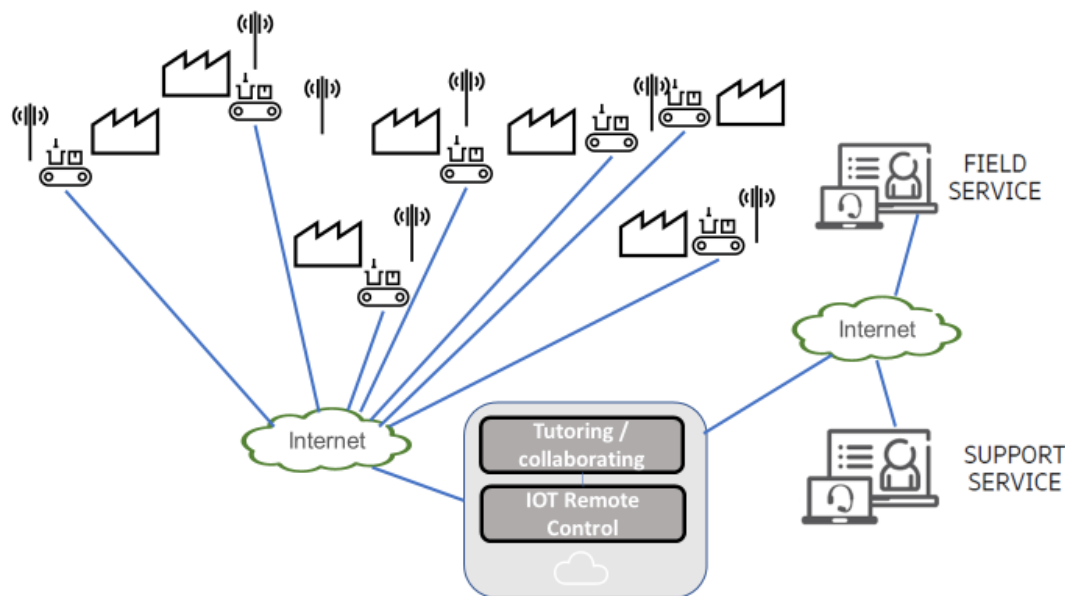


For anything else, that is non-existent, a common example is Subscriber Identity Module (SIM)-management, a centralized 3pp managed, multi-tenant solution should be considered to lower Total Cost of Ownership (TCO), especially if you look at this communication use case, as an isolated business case.

## 2.4. Augmented Reality (AR)

Augmented Reality, sometimes also referred to as Mixed Reality (MR), is expected to grow to a 14B USD business by 2022, as estimated by Forbes ([Reference 9](#)). Many of the use cases include “in-context” communication (see also section 2.1).

A good example of such solutions under trial, is “remote manufacturing”.



**Figure 3 - Network architecture for remote manufacturing**

Source: Ericsson

If an enterprise has factories in several geographic locations and has expensive and complex manufacturing machines in those locations, it is a common need to have a centralized a smaller group of experts (support service), to assist teaching and servicing field technicians on site in the factories.

### Tutoring use case

During an audio / video session between a support engineer and a field technician, a support engineer is to be able to capture in real time, the status of a machine at a factory, using the service “IOT remote control”, provided “as-a-service” by the service provider. The support engineer can then use the tutoring and collaboration service, to show the field technician the finding, by sharing the capture, and pointing to certain readings etc.

### Remote control use case

The support engineer can also use the same services, to actively control the machine from remote, while the field technician can observe, both the machine “in reality” and the remote control being done, as it happens.

As we move towards 5G, new low latency services will be possible, based on a highly distributed cloud capability. For this use case, performance will improve in terms of latency, as it will be possible to push / distribute the cloud data center closer radio network at each factory.

Given that the factories and the machines are a sunk investment, the service provider here, has more “share of value” as this capability helps lower the overall TCO, as training relatively many remote technicians from a central location using relatively few experts.

## **2.5. Smart Speaker for business**

We already covered smart speakers in section 1.2, but this device is also making its way into the enterprise. There are already enterprise versions of the smart speakers available, and integration to popular enterprise office software, including security and privacy systems is ongoing.

As we have already examined, the smart speakers have a natural fit with “in-context communication”, including UC, as well as customer support scenarios. Other use cases will surely be based on conference rooms and other collaboration spaces.

Service providers will need to continue making sure that the smart speaker can be a companion to their offered services, and conversely smart speaker vendors need to make sure that other office devices can be companions of the smart speaker.

## **3. Specific Technology Dependencies**

In this section, we will describe the specific enablers we need to realize the use cases and trends in sections 1 and 2. We will focus on the enablers that have not already been described in sections 1 and 2.

In section 4, we will describe common enablers, based on a conceptual blue print, to outline common needs across these use cases, that also depend on service provider scale, ambition and approach.

### 3.1. Consumer Communication Dependencies

## Specific functional technology dependencies Consumer

Dependency Trend/ Use Case	EPC ePDG, TWA G & AAA- server	IMS-extensions for Multi-X	Secure Entitlement	IMS-extensions new video codecs	B2B2C APIs B2C APIs	RCS Business Messaging	Camera feed enablement	VoLTE + Internet in one chipset
Companion Device	✓	✓	✓					
Personal Assistant	✓	✓	✓					
Audio / Video Doorbells	✓	✓	✓	✓				
Improved communications at venues					✓	✓	✓	
FWA with Consumer VoIP								✓

**Figure 4 - Dependencies on feature functionalities for the consumer trends and use cases**

Source: Ericsson

Wi-Fi-calling, which is used by smartphones and traditional companion devices like tablets, requires additional capabilities in the Evolved Packet Core (EPC):

- Enhanced Packet Data Gateway (ePDG). This is a secure termination node for IP Security (IPsec) tunnels established by the User Equipment (UE), supporting untrusted Wi-Fi.
- Trusted Wireless Access Gateway (TWAG). This is a secure termination node for protocols like GPRS Tunneling Protocol (GTP), for trusted Wi-Fi.
- AAA-server. Provides Authentication Authorization and Accounting and interworks with the cellular Home Subscriber Server (HSS), for access to cellular services when on Wi-Fi.

The IMS-extensions for Multi-X, supports Multi Device, Multi Persona and Multi User, as defined in section 1.

Secure Entitlement Server (SES) is a solution that interworks with the UEs, to get provision them in the Core Network for the companion device services, which are different in each service provider.

Whenever we add higher quality video (in doorbells), we may need to add support for that codecs, in the core network, so that you can have interoperability with devices that do not support the new format.

The venue uses cases require the different stakeholders such as the landlord, a sports league, and the service provider, to jointly provide the services. This in turn requires Application Programmable Interfaces (APIs), between the partners to settle billing, usage, exchange content etc.

## 3.2. Business Communication Dependencies

### Specific functional technology dependencies Business

Dependency Trend/ Use Case	IMS-based or IN-based in- call triggers	RCS 2.0 + enriched calling	In-context Web RTC IMS enablers	UC Partner extension on IMS	VoLTE for UC NNI- solution	Global SIM- management and connectivity	IMS/MSC support for Concierge + eCall	Fleet Manageme nt App Suite	Chipset support for VoLTE
Enhanced VoLTE call and Co- browsing	✓	✓	✓						
VoLTE for UC				✓	✓				
Voice for IOT Connected cars						✓	✓		
Voice for IOT Fleet Management							✓	✓	
Voice for IOT Elevators and Alarm Panels									✓
Augmented / Mixed Reality	✓		✓			✓			✓
Smart Speaker for Business (TBD)									

**Figure 5 - Dependencies on feature functionalities for the consumer trends and use cases**

Source: Ericsson

The Intelligent Network (IN) and IMS triggers are needed to trigger the enhanced service from a VoLTE call that has already been initiated.

There is a need for an operator owned RCS server, compliant to RCS 2.0 + enriched calling, to enrich a VoLTE call with messaging, carrying different media types such as hyperlink sharing and photo sharing.

The in-calling WebRTC with data channel, can be triggered from a WebRTC Gateway, which is often combined with a Session Border Controller (SBC), part of IMS, or it can be triggered from an ongoing call.

To implement VoLTE for UC, a UC-partner extension is needed as part of IMS, plus also a set of rules specified in an existing or new Network SBC (N-SBC), to implement a Network-to Network Interface (NNI), between the UC-network, and the VoLTE network.

For Connected Cars, if you as a service provider are aiming to provide global support for a car manufacturer, you will need a global SIM-management system, and connectivity to the various MNO RANs in the different geographies. The Service Provider will be an MVNO and will connect to the MNO RAN, using the VoLTE Roaming architecture for Home Routing.

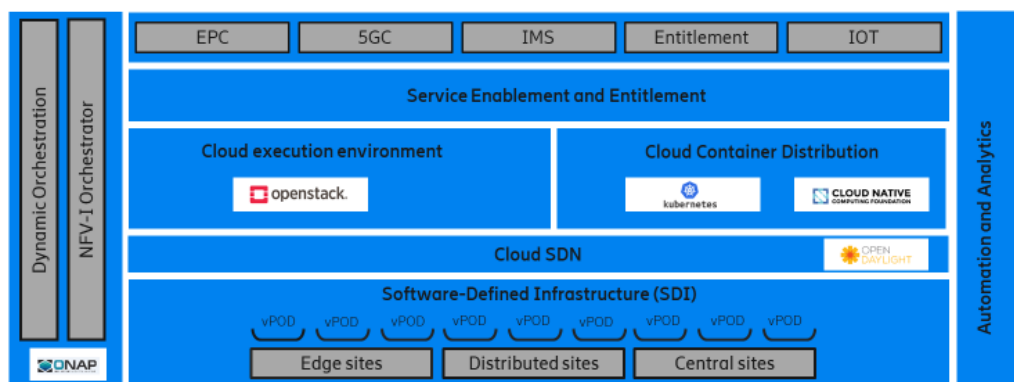
Note that you can also choose to play only locally (e.g. US only), for a car manufacturer, depending on the spectrum situation, you can may also use the MVNO architecture, or you can be the MNO. You can also be MNO in some areas and MVNO (with MNO partner) in other areas.

## 4. Common technology dependencies

In this section, we will walk through the technology enablers that reside in the cloud platform. Many of them are “non-functional” meaning that they are not providing any applications and they are not enabling any voice calls or internet sessions. However, they manage, optimize, secure and organize the cloud.

These enablers support multiple use cases, but how many of them you need, also depend on your ambition level when building your service provider cloud.

### Reference Stack



**Figure 6 - Reference stack blueprint**

Source: Ericsson

To guide this section, figure 6 represents a conceptual blueprint that can support the core network aspects that we are focused on in this document, showing the non-functional enablers combined in to conceptual “reference stack”.

We will use this to figure to explain these dependencies.

### 4.1. Enablers Explained

#### 4.1.1. NFV-I Orchestration and Cloud Execution Environment (CEE)

**Cloud Execution Environment (CEE):** CEE is an Infrastructure-as-a-Service (IaaS) solution. It is typically based on OpenStack, with additional features that expand its flexibility of use and meet the needs of communication service providers.

Openstack is an open source platform and has a wide participation in the service provider and IT space. Code is contributed by vendors to the service providers, and service providers themselves. The focus of the communication service providers has been on performance enhancement for our industry. This helps operators to avoid vendor lock in, benefit from the standard open interface and leverage on community fast evolvement.

The CEE is a software layer that manages the physical resources, which in the blueprint is called Software Defined Infrastructure (SDI). The resource managed are:

- Compute (CPUs)
- Memory
- Storage
- Networking / Switching

Openstack supports a set of APIs that can be used by the service provider to manage the cloud and the applications on it. An application / Virtual Network Function (VNF) can be on-boarded onto the cloud, instantiated (equivalent of installed in the physical world), and configured, and launched. Capacity can be expanded, or reduced, and the VNF can be upgraded, and decommissioned. There are also management APIs that provide performance statistics.

Network Function Virtualization Infrastructure (NFV-I) Orchestration entails performing these functions on a per application or VNF basis, continuously, and across several pods, as shown in figure 6. One data centre can have several pods. Each is an instance of the cloud, but the orchestrator is used to combine PODs into a network level cloud. The cloud itself can spread across many sites. Each VNF typically comes with an Element Management Server (EMS), which in standards is called a specialized VNF manager (S-VNFM).

This element knows the context of alarms and performance counters for its VNF or set of VNFs. Northbound, the VNFM interfaces the Orchestrator and takes part in digital “work flows”, that consists of instantiating, expanding, or contracting the VNFs, across the physical resources in each pod.

#### **4.1.2. Automation and Analytics**

The Orchestrator is providing automation of some basic tasks, but by using additional tooling we can automate more of the operations. By applying analytics of the performance of the cloud, we can also create Closed Loop Automation. A very good example of this is “self healing”. This is a process whereby the cloud can detect a piece of hardware as mal-functioning, is able to quarantine it, identify the applications / VNFs impacted (currently suffering from capacity degradation), restore the lost capacity on a new (healthy) server, and expand the affected VNFs to restore the capacity.

The other type of automation that can be supported has to do with Continuous Integration / Continuous Deliver (CICD). Service providers are adding tooling to receive software releases from the VNF vendors in a purely digital way, by connecting the vendor CICD systems with the Service provider Systems. By doing this, new features can be delivering quicker. CICD also provides ability to run predefined test cases, so will an opportunity to verify that the system performs well, before taking the new software into service.

### **4.1.3. Service Enablement**

This is the ecosystem support for the cloud, and this layer is responsible for exposing capabilities of underlying layers, towards application developers, using APIs. There will typically be several different service enablement components in this layer. Billing is perhaps the most obvious component. IOT may have one exposure component, and IN, and IMS have another one, and Entitlement can be a fourth component. Application providers from different industries will be able to put together the end user experience and complete it.

A good example of this type of application is developers of “in car camera feed (audio/video), for racing. APIs will need to be exposed from the enablement layer, to allow the application provider to tap into this feed from the platform. For optimal performance, this application should execute on premise at the race track, so should be distributed there.

### **4.1.4. Container Distribution**

Most of the enablers described in this section uses a virtual abstraction that can run on physical hardware, called Virtual Machines (VMs). A rich ecosystem has formed around these. As we approach 5G, the industry is looking at containers, as the next evolution steps. The most famous container execution environment is called Kubernetes.

A container is more efficient than virtual machines, because it does not allow the applications / VNFs to bring their own Operating System (OS), labeled “guest OS”, and load it on an underlying “host OS”. This provides a lot of flexibility as you can put a variety of different software with different OSes on the same cloud, but has drawbacks when it comes to performance, since the guest OS loaded on each VM will use extra CPU-resources etc. In a container environment, there is no Guest OS, and all applications must agree to run on one and the same OS, the container OS; frequently Linux.

We will see clouds introduce support for 5G Applications using Containers, but several different combinations will exist.

Initially, both VM and Container will be supported as well as Containers on VMs. Later, as container distributions and orchestration mature, VMs may disappear. The IT industry has been using containers for some time, so experiences from that can be reused.

### **4.1.5. Dynamic Orchestration**

Dynamic Orchestration is an extension of Cloud Orchestration, and can be delivered by the same product, or as separate products. Dynamic means adaptive, so able to partake in closed loop automation.

Whenever we need to orchestrate not only single VNF but complete groups of VNFs, then we talk about a network slice – something that is treated like a separate network, for instance the same common VoLTE IMS-core could serve multiple VoLTE for UC network slices (see section 2)

- a) One for each UC vendor
- b) One for each enterprise with the same UC vendor
- c) One for each UC vendor and enterprise

Another important area for dynamic orchestration is multi-site / multi data-center deployments. This is key for the distributed cloud that is envisioned for 5G. We must be able to instantiate and orchestrate a

virtualized and containerized network with sites distributed closer to the end users, to achieve the performance envisioned. To do that we need two main ingredients.

- i) ability to orchestrate a network slice or group of VNFs across multiple sites
- ii) ability to dynamically allocate and configure the transport network between the sites, using Software Defined Networking (SDN). This means that SDN is moving from just supporting each cloud data center or site, to supporting also the transport network.

SDN means that physical (e.g. fibre) capacity exists for example to a football stadium in over-capacity, but whenever the stadium is not used, then software capacity, including layer 3 virtual routing, that is needed to use the allocated physical capacity, is moved somewhere else.

Dynamic orchestration hence implies digital coordination with SDN-controllers, deployed as part of the SDN-layer in a larger data-center, must be part of the orchestration end-to-end. That way, when the football game is about to start, adequate capacity to support it can be moved there, as part of dynamic orchestration.

#### **4.1.6. Cloud SDN**

Cloud SDN simply implies both in-data center Software Defined Networking, and the site external SDN transport capability described in section 4.1.5, by means of an SDN-controller.

#### **4.1.7. Distributed Cloud**

Distributed Cloud means supporting a new type of site, further distributed than ever before, which the cloud can “spread” to, and where applications with extremely stringent requirements for low latency, can be orchestrated and configured dynamically. Sometimes the term Edge Compute is used for this as well, as described in section 4.15 and 4.1.6.

### **4.2. Service provider scale, ambition and approach**

In this section, we will be reviewing some ambition levels among service providers, and how they map to the non-functional common enablers.



# Common non-functional technology dependencies

## Scale, ambition and approach

dependency use case, scale, ambition or approach	NFV-I Orchestration + cloud execution env.	Automation and Analytics	Service Enablement	Container Distribution	Dynamic Orchestration	Cloud SDN	Distributed Cloud
Appliance operator or use case (network-in-a-box)	✓	✓					
Multi tenant operator, with cloud operations	✓	✓			✓		
5GC Multi tenant operator, with cloud operations	✓	✓		✓	✓	✓	✓

**Figure 7 - Segmenting Service Providers across common enablers**

Source: Ericsson

There is a high demand in the market for what is sometimes called Appliance, which means that it is a single tenant and single use deployment, which is static and hence does not need any dynamic orchestration.

Smaller service providers are often looking mainly for a smaller and more economic form factor and scale

But even larger operators are often interested in the ability to single out a use case, an enterprise, or a situation, and use an appliance, while we are waiting for advanced dynamic orchestration and network slicing, which is still a few years out.

Even after dynamic orchestration is out, there may be cases such as international deployments where the service provider is not present, or government critical infrastructure which simply cannot be shared, where appliance will still have a place.

Service Providers that choose the appliance model, still need cloud / NFV-I Orchestration and will also still need automation and analytics, but may not make use of the other enablers, as depicted in figure 7.

Cloud Operations, means running a separate organization just for the cloud platform described earlier in this section, taking full life cycle responsibility for that cloud platform, the vendors, and the open source software, and running it as a business separate from the application business. The cloud is distributed over multiple data centers, which brings the need for dynamic orchestration.

Those same service providers that are running cloud operations, are expected to start distributing further in the network in 5G time frame. No doubt, there will be some new service provider entering as well. In 5G timeframe, you will need a distributed cloud, so you will need almost all the enablers described earlier in section 4.

### 4.3. Advanced Use cases common enabler dependencies

In this section, we will be reviewing the more advanced use cases, and their dependency on common enablers.

## Common non-functional technology dependencies advanced communications use cases

use case, scale, ambition or approach	NFV-I Orchestration + cloud execution env.	Automation and Analytics	Service Enablement	Container Distribution	Dynamic Orchestration	Cloud SDN	Distributed Cloud
4G AR/MR	✓	✓	✓				
5G AR/MR	✓	✓	✓	✓	✓	✓	✓
Improved Communication at venues in 5G	✓	✓	✓	✓	✓	✓	✓
Enhanced Voice Call and Co-browsing	✓	✓	✓		✓		

**Figure 8 - Common enablers across advanced use cases**

Source: Ericsson

The use cases in this section, which have all been covered in section 1-3, are advanced mainly because they have dependency on low latency and therefore dependent on a distributed cloud and dynamic orchestration.

In 4G, some of these capabilities are not available.

Also noteworthy is that Service Enablement is needed for all of them, which means these are “ecosystem dependent”. Without an ability to expose APIs towards developers, a service provider will not be able to be successful in scaling and growing these businesses.

## Conclusion

For sure, making a one-on-one voice call just to say hello, is a service in decline. Like many other “mainstream” services, it has been replaced with several interesting niches. Everyone is now watching separate TV-shows so there is no “common ground” at the proverbial “water cooler” anymore.

In some of the “voice niches”, voice communication has an increase in value and an increase in usage, across age-segments.

It is up to each communication service provider to seek those out, and to connect them with other user experience values that they add, in their current business.

Some of these are obvious – content and media, broadband service, mobility. Some are more subtle, but just as important – trust and privacy, security, emergency services.

In North America, just the increase in connectivity options and devices drive a lot more opportunities across consumer and enterprise segments.

A few factors stand out, when reviewing the use cases and technology enablers in this document, in terms of how to address growth based on user experience going forward:

1. **Eco-system is vital.** This is another important strength in the cable industry. It has always been an eco-system, and collaborative. APIs will need to be opened up, for the service provider to be relevant in some of the key ecosystems like the smart home, the smart business, and the smart venue
2. **Platform matters.** The platform is very different now. It is open source, it is a separate business, and it is multi-vendor, but it is fundamental in seeking out the right niches, and enabling several very different businesses, to choose a service provider, for instantiating their own network slice, on the multi-tenant cloud.  
**Key platform Decisions:**
  - Appliance or cloud operations?
  - Single tenant, or multi-tenant with dynamic orchestration?
  - Distributed cloud or not?
3. **The biggest competitor to one-on-one voice, is text messaging,** another service you can bring with you...now evolved to RCS Business Messaging, with an opportunity to introduce and advertising revenue stream. How can that be captured? How can the cable ecosystem add value?
4. **Communication is now in-context, which exponentially increases its value.** Remote Manufacturing is dependent on the voice/video communication, to provide tutoring and remote control. A concierge call, or eCall from a connected car has very high value, and gamers are preferring a separate voice call over in-game voice, which also survives switching from one game to another...so a very long call.... how can we open those doors?

Let's pick our battles, and let's not go at it alone, and let's evolve how we think about communication services.

## Abbreviations

2G	2 <sup>nd</sup> Generation wireless
3G	3 <sup>rd</sup> Generation wireless
5G	5 <sup>th</sup> Generation wireless
AAA	Authentication Authorization and Accounting server
API	Application Programming Interface
B2B2C	Business-To-Business-To-Consumer
B2C	Business-To-Consumer

CAT-M1	CATegory M1
CBRS	Citizens Broadband Radio Service
CDMA	Code Division Multiple Access
CDMA 1x	1x Evolution-Data Optimized is next gen CDMA with improved data speed
CEE	Cloud Execution Environment
CICD	Continuous Integration Continuous Delivery
CPU	Central Processing Unit
CS	Circuit Switched
eCall	Emergency call from a connected vehicle (speaker and microphone)
ePDG	Enhanced Packet Data Gateway
EMS	Element Management Server
EPC	Evolved Packet Core
FWA	Fixed Wireless Access
GoS	Grade of Service
GSM	Global System for Mobile
GSMA	GSM Association
HMD	Head Mounted Display
HSS	Home Subscriber Server
HTML5	Hypertext Markup Language revision 5
IMS	IP Multimedia Subsystem
ISC	IMS Service Control interface
IOT	Internet Of Things
IP	Internet Protocol
IPsec	IP security
IT	Information Technology
LTE	Long Term Evolution
M2M	Machine-To-Machine
MMTel	Multi Media Telephony
mPERS	Mobile Personal Emergency Response Device
MNO	Mobile Network Operator
MVNO	Mobile Virtual Network Operator
MSC	Mobile Switching Center
MSO	Multiple System Operator
NFV	Network Function Virtualization
NFV-I	NFV Infrastructure
NNI	Network-to-Network Interface
N-SBC	Network-to-network SBC
OEM	Other Equipment Manufacturer
PBX	Private Branch eXchange
PSAP	Public Safety Answering Point
PTT	Push-To-Talk
RAN	Radio Access Network
RBOC	Regional Bell Operating Company
RCS	Rich Communication Suite
SBC	Session Border Controller
SDI	Software Defined Infrastructure
SDN	Software Defined Networking
SES	Secure Entitlement Server
SIM	Subscriber Identity Module
SIP	Session Initiated Protocol
S-VNFM	Specific VNF-Manager
TCO	Total Cost of Ownership

TWAG	Trusted Wireless Access Gateway
UC	Unified Communication
UE	User Equipment
USD	US Dollar
VoIP	Voice over Internet Protocol
VoLTE	Voice over LTE
VNF	Virtual Network Function
VNF-M	VNF-Manager
VM	Virtual Machine
WebRTC	Web based Real Time Communication
Wi-Fi	Wireless Fidelity
WPS	Wireless Priority Service

## Bibliography & References

- *Connected Wearables*, M2M Research Series 2019, Berg Insights
- Vodafone Press Release – Smart Speaker ([Reference 1](#))
- MARKETSANDMARKETS.com Smart Speaker Market ([Reference 2](#))
- marketwatch.com – Smart Doorbell market ([Reference 3](#))
- digitaltrends.com – best video doorbells at CES 2019 ([Reference 4](#))
- GSM-Association – RCS Business Messaging ([Reference 5](#))
- GSM-Association – RCS 2.0 Enriched Calling Specification ([Reference 6](#))
- Youtube demo – Telus Cloud Contact Center ([Reference 7](#))
- Html5rocks.com – Tutorial on WebRTC with data channel ([Reference 8](#))
- Forbes.com – Augmented Reality Market Projection ([Reference 9](#))
- *The future of the SIM*, potential market and technology implications for the mobile ecosystem, GSM-Association (available on [gsma.org](http://gsma.org))

# Operational Transformation Using GIS

A Technical Paper prepared for SCTE•ISBE by

**Derek Rieckmann**  
Senior GIS Manager  
Midco  
3901 N Louise Ave Sioux Falls, SD  
605-274-2977  
Derek.rieckmann@midco.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Content .....	3
1. Web-based GIS .....	3
2. Coax Data Conversion and Billing Integration .....	4
3. Plant Analysis and Visualization .....	5
Conclusion .....	6
Abbreviations.....	7

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Proportional Symbols of Data Usage by Address .....	6

# Introduction

Through digital transformation of legacy mapping to a more spatially intelligent option, new methodologies have been implemented across the enterprise around activities like plant maintenance, extensions, and troubleshooting. Midco continues to enhance its geospatial architecture by integrating with other systems and, depending upon the data within the system, to drive processes through the rest of the organization. Geographic Information Systems (GIS) are no longer just a system of record, but also becoming one of engagement and insight. Rather than simply reading data, it is the authoritative source for many datasets.

This paper focuses on how Midco is able to leverage GIS technology to help drive operational efficiencies and prioritize capital expenditures. We will discuss the path Midco took from using GIS to record network information with very specific use cases, to an enterprise level platform with over 80% of employees directly using web delivered services.

## Content

### 1. Web-based GIS

Midco made a collection of apps, analyses, process changes, and data conversions to enable GIS to optimize operational efficiency. At Midco, GIS started out as a way to model the fiber optic network and expand business services. The potential was realized and plans to convert the coaxial network from Computer Aided Design (CAD) and Lode were developed. As the conversion was being planned, Midco continued to develop other use cases for GIS assisting with Sales, Marketing, and Construction Departments as well as rolling out web-based GIS. Operational transformation didn't occur at Midco from a single project, but rather several projects of varying scopes and complexities.

Web-based GIS started within Midco as only accessible via desktop computers and within the firewall. Eventually the system was upgraded to a portal environment that supported mobile devices and increased functionality through the use of widgets and app templates. This also gave end-users the ability to find meaningful content in a user-friendly manner.

Midco has seen significant usage of web-based GIS since implementing. The user base has grown to almost 80% of all employees within the company. Additionally, there are hundreds of views per day on GIS data and websites. This pattern of deployment has transformed GIS from a tool that a handful of employees use to a tool that is critical in daily operations at Midco.

Mobile access plays a large role in the uptick in GIS usage. Not only are web-browser based GIS apps supported, but also native iOS apps are used every day by Field Operations. One such example is Midco's Field Vision app which allows field staff to capture points on a map and record information about any issues going on in the plant that need to be addressed. Support for attaching pictures is included and any data captured is instantaneously transmitted back to the enterprise and can be assigned to the proper personnel to resolve the issue. A status on all work-orders is available and once the issue is resolved it's closed. Reports are available for managers to review and see what work is being done and by whom, giving further visibility into what tasks are being done. This simple app replaces the combination of paper and spreadsheets, which was prone to error.

Deploying GIS based web applications goes through a similar cycle as any application. First, a use case for the app must be identified and requirements defined. Then a beta version of the app is deployed to the



stakeholders and any feedback is applied back into the app. Finally, the app is deployed into production and then as enhancements are requested they can be applied to the app. Midco also implemented an app retirement process, which gives a workflow to delete apps that are no longer being used, or have been replaced by newer apps.

Having a way to clean up obsolete content is critical to maintaining a web GIS deployment. Data, code, and any other applicable information are all archived in case the app needs to be revived in the future. It's important to note that data may be used by multiple apps or content in a GIS environment, so only if that data is not being used by anything or anyone else should it be archived. Organization and metadata are important to ensuring that the correct steps are being taken during the retirement process.

## **2. Coax Data Conversion and Billing Integration**

Midco's journey toward operational transformation with GIS was greatly and positively impacted by converting data from a CAD environment to GIS. This project involved taking design files, CAD maps, and billing data and combining them into GIS. The resulting product was a system that could design and draft simultaneously and integrate with a billing system. This single project opened up many more opportunities within the company to integrate with other systems as well as drive workflows via the enterprise GIS.

A project of this magnitude needed to have an internal Return on Investment (ROI) review created. The primary driver that could be identified was savings through bringing coax design in-house, rather than continuing to contract it out. Doing this meant that resources would have to be staffed and trained as well as a design software selected or developed. Midco chose to partner with a software development company to create a new design product, native to the existing enterprise GIS. Creating this product on the already implemented platform allowed Midco to avoid potentially troublesome Extract Transform Load (ETL) processes and also get data instantly into the production database.

The GIS conversion at Midco was the digital transformation that brought the rest of the network (the coaxial portion) up to date with the fiber optic portion. Having a single platform allowed all entities that interact with the GIS system to operate much more efficiently. It also gave visibility at the address scale and opened up possibilities to gather data from various systems and associate it to a spatial point within the GIS.

Three distinct systems (billing, design, and drafting) were combined during the conversion. Data discrepancies were found among the systems and rules were established about what to do in the most common scenarios. Often times, a hierarchy of the originating systems was used in order to determine which source would be migrated. Other times, data transformation would happen and the newly created information was written into GIS. The primary takeaway is that a plan must always be established during a migration for handling conflicting, missing, or data that doesn't fit into the target system. Establishing this upfront by utilizing data exploration tools is critical to the success of a migration.


The billing integration was a way for Midco to avoid the duplicate data entry of addresses from walk-out into GIS and then later again into the billing system. It also kept attributes that GIS maintains (serviceability, network information, etc.) automatically up to date among systems. In the past, these were maintained via spreadsheet. Midco averages hundreds of changes per day, so automating this process was identified as a way to greatly reduce time and errors associated to transferring this data.

### 3. Plant Analysis and Visualization

Midco has experienced a lot of success with plant analysis and visualization. Fiber and coax mileage data is compiled weekly and presented in a non-spatial dashboard. This is one example of taking spatial data and tools and compiling it in such a way that is useful to end-users who don't want it on a map. Midco has found that converting data from spatial to non-spatial is a simpler task than converting it from non-spatial to spatial. This type of analysis can be accomplished through a variety of spatial geoprocessing tools, structured query language (SQL) queries, and scripting languages to create high quality outputs on a regular, automated schedule. For example, the coax plant report used to be manual and took over 200 working hours to complete. Since it was so labor intensive, it could only be done once a year vs the once a week cadence it's calculated at now.

Proactive Network Maintenance (PNM) is another area of opportunity that Midco took advantage of to use GIS to operate more efficiently. The integration that Midco deployed flows in two directions. In the first direction PNM software consumes web-based map services so that plant can visually be displayed within the PNM application. In the second, the GIS consumes data from PNM and displays it at the node level for use by field and engineering staff to find spatial patterns. Node health and other statistics are also derived from other systems and displayed within the same spatially enabled web application. The application also allows users to jump back to a specific date and view the statistics for that day.

Within the same application, data usage by address is compiled and displayed for a number of time windows (1 month, 3 months, 6 months, and 12 months) and is displayed as proportional points on the map (Figure 1). The ranks of addresses' data usage for the top ten users within each node are also displayed. Having this data allows for the creation of heat maps of cities, showing which regions have the highest data usage. The data is also transformed into a space-time cube, where in-depth analytics and machine learning can be utilized to find spatial trends over time. Ultimately, this enables better predictions of where in the network usage will increase faster in the future. The goal of all these tools is to help make better decisions on where and when to augment the network to take care of current troublesome areas and anticipate where future issues may arise.



**Figure 1 - Proportional Symbols of Data Usage by Address**

## Conclusion

In Midco's experience, operational transformation does not occur from a single project and can't be narrowed down to a single point in time. Operational transformation is a collection of projects and processes whose end-products continue to evolve. The transformation seen from GIS has brought Midco from a company where maps were referenced to a company where maps are still referenced, but the data contained within the GIS drives business processes downstream. It has gone from a nice to have tool to an enterprise level system that is part of larger processes and cannot simply be removed.

GIS adds a visual element to analyses and simply displaying data spatially helps detail out powerful stories. However, it goes far beyond that by being the engine through which phenomena can be examined through a spatial lens in ways not possible in standard reporting software. As more systems and workflows are integrated with GIS, the possibilities for the problems it can help to solve also increase. An effective strategy for growing a GIS within an organization is to first find a use case that is spatial in nature and then create the data, systems, and processes needed to solve that use case. Then, that new product can be utilized for something completely new, which leads to increased investment, use, and visibility of GIS. This snowball effect is one that was successful at Midco and can be replicated elsewhere.

## Abbreviations

CAD	Computer Aided Drafting
ETL	Extract Transform Load
GIS	Geographic Information Systems
MSO	Multiple Systems Operator
PNM	Proactive Network Maintenance
ROI	Return on Investment
SQL	Structured Query Language

# **Customer Safety Initiative (CSI)**

A Technical Paper prepared for SCTE•ISBE by

**Matt Carothers**

Security Architecture

Cox Communications

6305A Peachtree Dunwoody, Atlanta, GA 30328

**Damien Whaley,**

Cox Communications

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
The Problem.....	3
Secure Product Development Life Cycle (PDLC) .....	3
Identifying Malicious Control Channels .....	4
1. Stating the Problem.....	4
2. Identifying the Malicious Traffic .....	10
3. Scoring Example.....	11
4. Score Interpretation.....	11
Case Studies .....	12
5. Mirai.....	12
6. Case Study – Web Scrapers .....	14
Code Example.....	14
7. Output.....	15
References.....	16

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Product Development Lifecycle .....	3
Figure 2 – Example of a Single Clear Host.....	4
Figure 3 – Real Network Graph .....	5
Figure 4 – Illustration of simple Domain Name System.....	6
Figure 5 – Illustration of Bengin Traffic .....	7
Figure 6 – Adding a Control Group to Identify Benign Traffic .....	8
Figure 7 – Problems Caused by Infected Hosts in the Control Group .....	9
Figure 8 – Problems Caused by Sampling .....	10
Figure 9 – Scoring Example.....	11
Figure 10 – Analysis of DNS requests.....	12
Figure 11 – Zooming in on the Analysis of DNS requests.....	13

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Netflow Results .....	14

# The Problem

In today's world, more and more companies create internet-connected products. It is estimated that the global internet population includes upwards of 17 billion connected devices (Leuth, 2018). The so-called "internet of things" (IoT) means that even common household items such as light bulbs and refrigerators feature internet connectivity. Companies with no prior networking experience now rush to market with little thought for security. Their inexperience or outsourcing to the lowest bidder creates a fertile ground for cybercrime. Criminals write worms to infect devices such as home routers, cameras, and even teapots. They compromise millions of vulnerable devices, join them together in a network called a "botnet," and use them to launch cyber-attacks. Such attacks are growing rapidly. During the first half of 2018 alone, IoT malware grew three-fold (Spadafora, 2018).

As internet service providers, we play an important role in fighting these botnets. First, we perform penetration tests against devices we deploy to our customers in order to avoid becoming part of the problem. Second, we work with third parties who report malicious activity in order to identify the Command and Control (C2) servers for botnet infections on our network.

## Secure Product Development Life Cycle (PDLC)

Security is embedded into the product development lifecycle to proactively identify vulnerabilities and ensure compliance with security requirements.

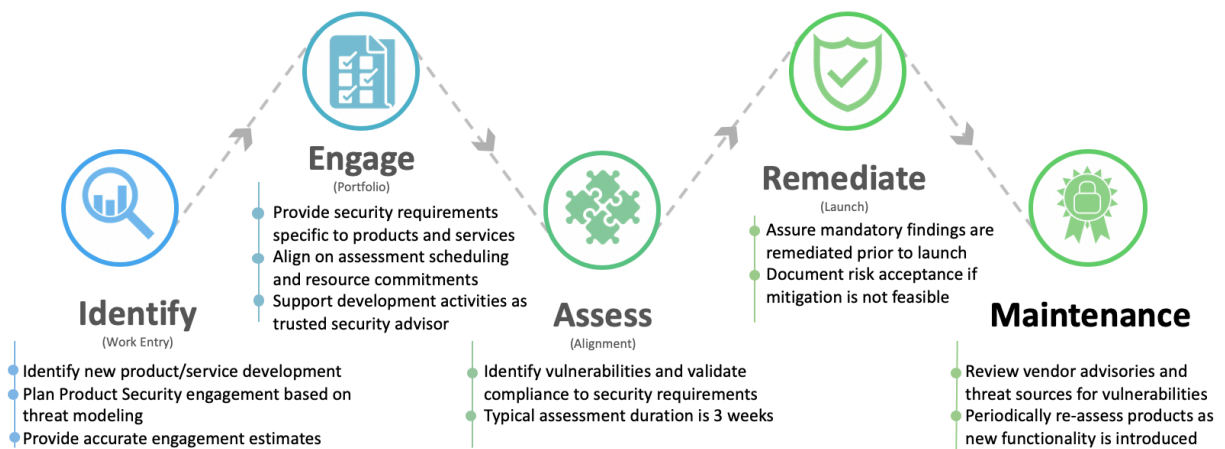



Figure 1 – Product Development Lifecycle

# Identifying Malicious Control Channels

## 1. Stating the Problem


We often receive lists of Internet Protocol (IP) addresses participating in some malicious activity, such as sending spam, port scanning, or launching Distributed Denial-of-Service (DDOS) attacks. Third parties request that we examine network traffic in order to identify the server controlling the bots. We would like to see something like this, a single, clear host in common with all the bots:



**Figure 2 – Example of a Single Clear Host**




However, real network graphs look more like this:



**Figure 3 – Real Network Graph**


While manual investigation is possible, we need an automated process.

As a naïve first approach, we simply sample some Domain Name System (DNS) requests from known infected IP addresses and find out what they have in common. We hope to find a clear picture like this:



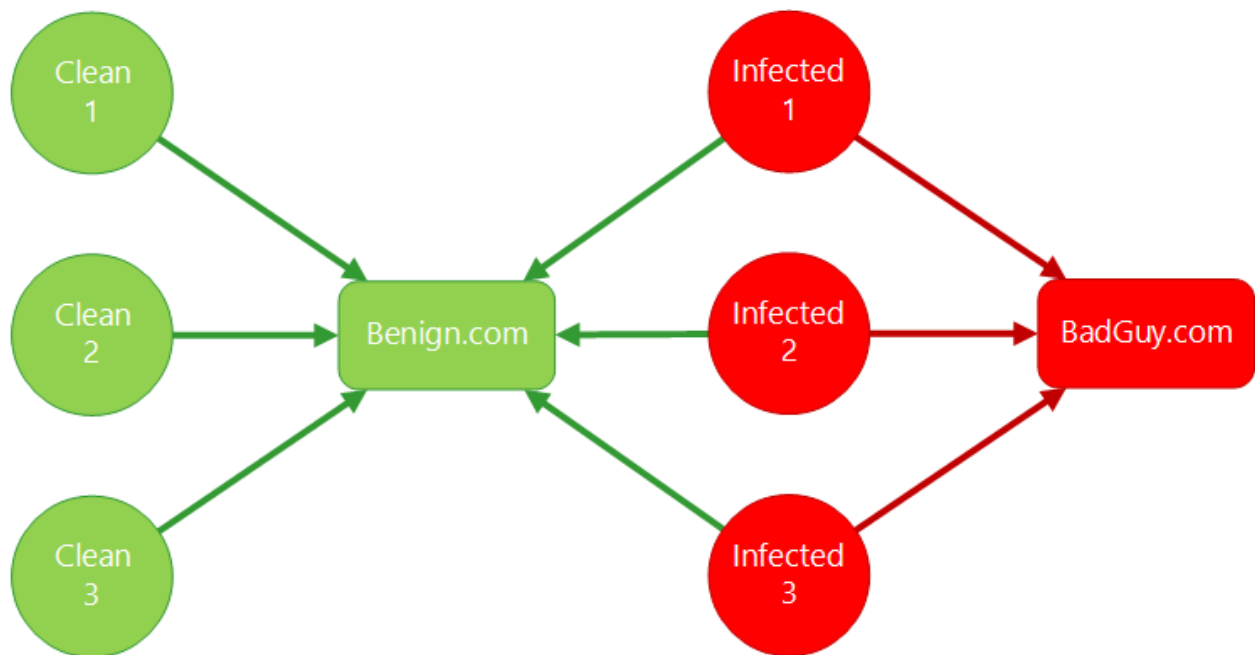
**Figure 4 – Illustration of simple Domain Name System**

Unfortunately, such an approach fails because most subscribers have devices contacting popular destinations such as Google, Facebook, Microsoft, Twitter, and various Content Delivery Networks (CDNs) in addition to anything malicious they might share. Simply flagging everything a group of infected devices have in common will therefore generate many false positives.




**Figure 5 – Illustration of Benign Traffic**

Our second approach adds a control group. Perhaps we can map the connections from a group of noninfected devices and eliminate those results from the set?




**Figure 6 – Adding a Control Group to Identify Benign Traffic**

This too fails because we can never say for certain that a given device is not infected. At best we can say we do not yet know it is infected. The presence of an infected device in the control group spoils this approach.



**Figure 7 – Problems Caused by Infected Hosts in the Control Group**

Sampling compounds the problem. We sample our netflow at 1:2000, meaning that we only send flow data for 1 out of every 2000 connections. Logging all DNS data all the time is not practical, so we sample DNS for a short period of time and hope that the malicious activity takes place during that time window. Thus, we cannot assume that every device will be found to connect to a given target during our sampling period.



**Figure 8 – Problems Caused by Sampling**

## 2. Identifying the Malicious Traffic

In order to find the malicious control channel, we need a score rather than a binary yes or no. Our score must increase for a given target as more known infected devices communicate with it. Our score must decrease as more presumed uninfected devices communicate with it. One might be tempted to apply machine learning to the problem and attempt to cluster the devices, but we developed a much easier solution: fractions. Our score is simply the percentage of known infected hosts contacting a target divided by the number of presumed uninfected hosts:

Percentage of infected hosts contacting a target

---

Percentage of clean hosts contacting a target

### 3. Scoring Example

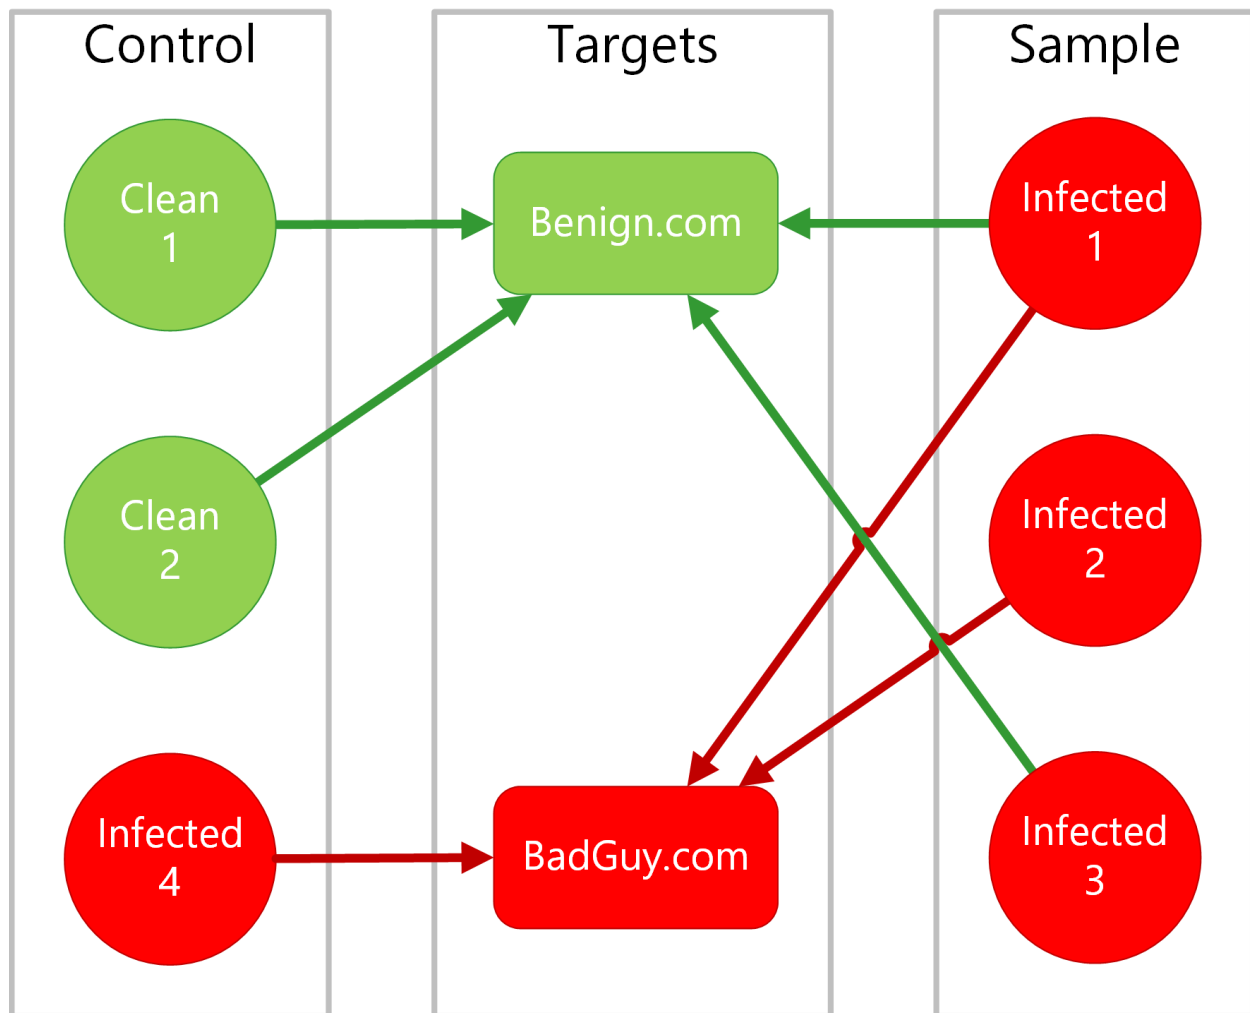


Figure 9 – Scoring Example

- 2 / 3 of infected hosts connecting to benign.com and 2 / 3 of clean hosts as well.  
**Score = (2/3) / (2/3) = 1**
- 2 / 3 of infected hosts connecting to badguy.com and 1 / 3 of “clean” hosts.  
**Score = (2/3) / (1/3) = 2.**

### 4. Score Interpretation


- A score of less than 1 indicates the target is under represented in the sample/infected set versus the control/clean set. I.e. the percentage of infected hosts connecting to a site is smaller than the percentage of hosts in the general population.
- A score around 1 indicates the target is found equally in both the sample/infected and control/clean sets.

- A score greater than 1 indicates the target is over represented in the sample/infected set vs. the control/clean set. I.e. the percentage of infected hosts connecting to a site is higher than the percentage of hosts in the general population.

## Case Studies

### 5. Mirai

In 2016 Cox received a list of IP addresses participating in a DDOS attack. Analysis of DNS requests showed this:



**Figure 10 – Analysis of DNS requests**

Zooming in, we see the top 4 hosts are Mirai controllers. The next 4, while not malicious themselves, are often associated with malicious activity.






Figure 11 – Zooming in on the Analysis of DNS requests

## 6. Case Study – Web Scrapers

In 2019 a trusted partner reported to us that Cox IPs were scraping their web site in order to conduct fraud. Analyzing netflow for the addresses gave us these results:

**Table 1 – Netflow Results**

TargetIP	Port	Score	Standard Deviations	ASN	Org
23.59.30.52	443	3043.648	23.396	9498	BHARTI Airtel Ltd.
54.223.56.208	443	3043.648	23.396	55960	Beijing Guanghuan Xinwang Digital
52.80.75.244	443	2790.011	21.446	55960	Beijing Guanghuan Xinwang Digital
104.31.95.106	443	2790.011	21.446	13335	Cloudflare Inc.

The first IP address in the list was most likely the control channel. More interestingly, analysis of Transport Layer Security (TLS) certifications on the second and third IP addresses showed them to belong to jide[.]com, home of "Remix OS Player - The Most Advanced Android Game Emulator for PC." We thus concluded that the malware is android-based. Also of interest is the fourth IP address, which is Pastebin. Pastebin may have also been a control channel, or it may have been used to exfiltrate data.

## Code Example

```
from missinglink import MissingLink

# Instantiate the linker with optional
# labels for the sample and control groups.
linker = MissingLink("infected", "clean")

# Designate some entities as part
# of our test group. All other entities
# are assumed to be part of the control
# group.
linker.label("10.0.0.1")
linker.label("10.0.0.2")
linker.label("10.0.0.3")

# Add some malicious relationships. 6.6.6.6 is our fictitious
# malicious site.

linker.link("10.0.0.1", "6.6.6.6")
```

```
linker.link("10.0.0.2", "6.6.6.6")
# Add some benign relationships.
linker.link("10.0.0.1", "8.8.8.8")
linker.link("10.0.0.2", "8.8.8.8")
linker.link("10.0.0.3", "8.8.8.8")
linker.link("10.0.0.4", "8.8.8.8")
linker.link("10.0.0.5", "8.8.8.8")
linker.link("10.0.0.6", "8.8.8.8")
linker.link("10.0.0.6", "9.9.9.9")
# Analyze the results
linker.analyze()
```

```
# Analyze the results
linker.analyze()
# Output the results
for result in linker.results:
    print(json.dumps(result))
```

## 7. Output

```
{
  "target": "6.6.6.6",
  "score": 2.0,
  "deviations_from_mean": 1.224744871391589,
  "infected_count": 2,
  "infected_percent": 0.6666666666666666,
  "clean_count": 0,
  "clean_percent": 0.0
```

}

Target	Score	Deviations from mean	Infected count	Infected percent	Clean count	Clean percent
6.6.6.6	2	1.2	2	67%	0	0%
8.8.8.8	1	0.0	3	100%	3	100%
9.9.9.9	0	-1.2	0	0%	1	33%

## References

Spadafora, A. (2018, September 19) *IoT malware grew significantly during the first half of 2018*.

Retrieved from

<https://www.techradar.com/news/iot-malware-grew-significantly-during-the-first-half-of-2018>

Leuth, K. (2018, August 8). *State of the IoT 2018*. Retrieved from <https://iot-analytics.com/state-of-the-iot-update-q1-q2-2018-number-of-iot-devices-now-7b/>

# **Mid-band Spectrum Opportunities And Challenges**

## **Balancing Coverage and Capacity for 5G Deployments in the US Market**

A Technical Paper prepared for SCTE•ISBE by

**Craig Schwechel**

Principal

inCode, a division of Ericsson

6300 Legacy Dr, Plano, TX 75024

(972) 583-0000

[cschwechel@incodeconsulting.com](mailto:cschwechel@incodeconsulting.com)

**Marth Wilson**, inCode Consulting, a division of Ericsson

**Lauren Buhl**, inCode Consulting, a division of Ericsson

**Tomislav Marcinko**, inCode Consulting, a division of Ericsson

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Synopsis .....	3
Introduction: .....	4
1. Defining the Mid-Band Spectrum Frequency Range .....	5
2. Mid-band Spectrum in Focus.....	6
2.1. CBRS 3550MHz-3700MHz Overview.....	7
2.1.1. CBRS Deployment and Business Case Considerations.....	7
2.2. C-Band 3700-4200MHz Overview.....	8
2.2.1. C-Band Deployment and Business Case Considerations.....	9
2.3. BRS/EBS 2496-2690 MHz Overview .....	9
2.3.1. BRS/EBS Deployment and Business Case Considerations.....	10
3. Mid-band Use Case and Competitive Considerations .....	10
4. Business Model Considerations .....	15
4.1. Cost Considerations .....	15
4.2. Mid-band Deployment Cost Drivers.....	15
4.2.1. Primary drivers for mid-band deployment cost variability.....	16
4.2.2. Secondary drivers for mid-band deployment cost variability .....	16
4.3. Mid-band Spectrum Revenue Benefit.....	17
Conclusion .....	18
Abbreviations.....	19
Bibliography & References .....	20

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1. New range of mid-band is 2.5-8GHz .....	5
Figure 2. Average spectrum holdings for major operators by low, mid, and high band breakout.....	6
Figure 3. Illustrative 3.5GHz CBRS deployment in a low to mid-density area to achieve 50Mbps service with 6.5km ISD.....	8
Figure 4. Enhancement of key capabilities from IMT-Advanced to IMT-2020 with a perspective on mid-band's ability to deliver the capability.....	11
Figure 5. Usage scenarios of IMT for 2020 and beyond .....	12
Figure 6. Latency and bandwidth requirements to deliver mobility use cases.....	13
Figure 7. Urban eMBB use case on 20MHz of B4 LTE and 60MHz of B46 LAA.....	14

# Synopsis

To date the FCC has made ~633MHz of licensed spectrum (below 6GHz) available to mobile network operators, which they have deployed to achieve broadscale coverage. However, to deliver the coverage and capacity required for emerging enhanced mobile broadband (eMBB), augmented reality (AR), virtual reality (VR), fixed wireless access (FWA), vehicle-to-everything (V2X), etc. use cases, much deeper pools of spectrum will be required, with mid-band playing a critical and broad role. Broadband Radio Service / Education Broadband Service (BRS/EBS), Citizens Broadband Radio Service (CBRS), and C-band spectrum bands have the potential to provide over 500MHz of additional capacity, on a much more economically viable deployment footprint than mmWave. This paper will explore the opportunities and challenges mid-band spectrum presents and will share perspective across three categories:

- **Use Case and Competitive Considerations** – the depth of spectrum available in the mid-band increases the economic viability for use cases not traditionally aligned with 3GPP technology, including FWA and video distribution, presenting an opportunity and threat to cable's traditional businesses.
- **Business Model Considerations** – Innovative and flexible models for mid-band spectrum allocation makes operators build their business cases around usage and deployment scenarios, including neutral host and private networks, who bring their own value props. Spectrum acquisition cost is increasing for mid-band spectrum, especially if wide, continuous channels are made available. 5G use cases drive the revenue opportunities and we are just beginning to see the innovation in what can be built on 5G networks.
- **Deployment Considerations** – While some of the mid-band spectrum (BRS/EBS) can be deployed using a traditional macro approach, greatest performance will be achieved through densification. Line of sight, building penetration, and cell edge performance will drive mid-band site placement considerations.

# Introduction:

Mid-band spectrum is lauded as the Goldilocks bands for 5G uses with the just-right combination of coverage and capacity. **This paper will provide an overview of mid-band spectrum in the US, validate the technical claims of mid-band given typical deployments, and explore 5G use cases most aligned with mid-band spectrum with corresponding deployment and business case considerations.** We incorporate our experience with delivering consulting engagements supporting our clients' decision points around mid-band spectrum. We also include our view on how mid-band spectrum can complement 5G deployments in the coming years.

Most 5G deployments will benefit from combining low, mid, and high-band spectrum to leverage the best characteristics of each set of bands. Low band provides ubiquitous coverage and can serve as an anchor band for 5G network cores. Mid-band combines access to wide channels with workable propagation characteristics to deliver gigabit-capable service over cell areas of several square kilometers. High band, or millimeter wave, spectrum uses ultra wide channels to deliver multi gigabit-capable service in hot-spot areas. Mid-band spectrum is capable of providing 5G-capable low latency, gigabit service without the 5-10x cell site densification required for mmWave spectrum. 5G deployments should aim for 100MHz of aggregated spectrum per operator for services. This can come from carrier aggregation of existing smaller bands or aggregating larger channels of mid-band spectrum.

The US is currently limited in its available mid-band spectrum assets and the calls for more mid-band spectrum are growing. On average, US peer wireless markets will have 4x more mid-band spectrum available by the end of 2020 for deployments<sup>1</sup>. There is tremendous political motivation to free up more spectrum for 5G, including mid-band, as the infrastructure and technology improvements are thought to spur innovation. The administration and the FCC want to do what they can to ensure the next generation of apps and start-ups enabled by 5G are done in the US, hence the FCC's 5G FAST plan<sup>2</sup>. The FCC chairman's 5G strategy includes three key components: Pushing more spectrum into the marketplace, updating infrastructure policy; and modernizing outdated regulations. Making additional mid-band spectrum available falls into the first component of this strategy.

We will focus on the 2.5GHz, 3.5GHz, and 3.7GHz licensed spectrum bands in the US for mobile and fixed wireless uses including commercial considerations for bringing use cases on these bands to the market.


---

<sup>1</sup> <https://www.ctia.org/news/mid-band-spectrum-global-update>, <https://www.ctia.org/news/more-mid-band-spectrum-is-key-to-u-s-5g-leadership>

<sup>2</sup> <https://www.fcc.gov/5G>



# 1. Defining the Mid-Band Spectrum Frequency Range



**Figure 1. New range of mid-band is 2.5-8GHz**

2.5GHz licenses making up most of the current mid-band holdings. Other countries with a similar wireless market will have on average 4x more mid-band spectrum compared to the US by 2020<sup>4</sup>. Playing catch up requires pushing through more spectrum and for the operators deploying that spectrum to align it with their existing portfolios and corporate strategies. Acquiring more mid-band spectrum will help shape how operators define their strategies in the 5G era.

Mid-band spectrum in the US has evolved in its definition. Before millimeter wave (mmWave) spectrum was available for mobile use, the Clearwire 2.5GHz was considered high band spectrum. AWS and PCS were considered high band by some and mid-band by others. Broadcast, SMR, Cellular, and all bands below 1GHz were considered low band spectrum.

The definitions have shifted with mmWave spectrum now deployed for mobile use. While millimeter wave technically isn't single-digit mm wavelengths until 30GHz, convention has labeled the 24GHz and higher spectrum as mmWave.


Low band spectrum is still 600MHz, 700MHz, and Cellular, but now includes AWS and PCS in the 1700/2100MHz and 1900MHz frequencies, respectively. This grouping is driven by relative characteristics similarities and colocation occurrences, meaning AWS and PCS can work on a 700MHz cell site grid.

Mid-band spectrum can also be grouped by performance characteristic similarities and intended usage. Spectrum from 2.5GHz – 8GHz<sup>3</sup>, Figure 1, can be defined as mid-band spectrum in that it provides suitable propagation characteristics with wide channel bandwidths to provide a balance between coverage and capacity. The “sub-6” label has gained in usage and can be consider synonymous with this mid-band label.

Mid-band spectrum holdings in this defined range in the United States has been limited, Figure 2. Average spectrum holdings for major operators by low, mid, and high band breakout show Sprint's

<sup>3</sup> FCC, inCode, <https://www.fcc.gov/document/fcc-opens-inquiry-new-opportunities-mid-band-spectrum-0>

<sup>4</sup> FCC 5G FAST plan



**Figure 2. Average spectrum holdings for major operators by low, mid, and high band breakout**

Source: inCode, FCC

## 2. Mid-band Spectrum in Focus

The industry has recognized the importance of mid-band spectrum in this defined frequency range as it looks ahead to 5G deployments. The FCC has committed to bringing more mid-band spectrum on line with a combination of auctions and flexible use policy approaches. Industry leaders are calling for more spectrum and to speed up the process. FCC Commissioner Michael O’Rielly blogged “**more attention needs to be paid to the mid-bands**”<sup>5</sup> and USCC’s Meyers lobbied, “...as we continue to meet the growing demand for data services and further identify and define potential 5G use cases, **we implore the FCC to bring as much mid-band spectrum to market as possible, as soon as possible** and within a framework that will allow regional and smaller wireless carriers to continue to meaningfully participate in this industry.”<sup>6</sup>

We will provide an overview of three licensed mid-band spectrum bands and their deployment considerations to validate their place in a 5G spectrum portfolio.

<sup>5</sup> <https://www.fcc.gov/news-events/blog/2017/07/10/mid-band-spectrum-win-making>

<sup>6</sup> Kenneth Meyers, President, CEO & Director, USCC, Telephone and Data Systems, Inc., United States Cellular Corporation, Q2 2019 Earnings Call, Aug 02, 2019

## **2.1. CBRS 3550MHz-3700MHz Overview**

Citizens Broadband Radio Service (CBRS) is TDD spectrum in the 3550-3700MHz<sup>7</sup> band that was first reserved for military use but could serve 5G customers under a flexible spectrum sharing plan. There is up to 70 MHz of CBRS to be licensed in each county in the Priority Access License (PAL) auction and an additional 80 MHz of unlicensed or lightly-licensed spectrum available for General Authorized Access (GAA).

The CBRS auction is now expected June 25, 2020 following the posting of the final auction rules in October 2018<sup>8</sup> and pending a vote at the September 26, 2019 FCC meeting. inCode estimates a PAL auction price average across the counties at \$0.13/MHz-pop. This value is at a discount to the global mid-band spectrum average of \$0.18/MHz-pop due to the restrictions and preemption characteristics on the bands. Frictional costs of adding the Spectrum Access System (SAS) / Environmental Sensing Capability (ESC) costs to the CBRS business case could push the 70MHz of PAL spectrum auction price down to \$0.10/MHz-pop. Google has announced its FWA pricing for SAS services at \$2.25/HH/month<sup>9</sup>. Connected device and mobility pricing may be different.

CBRS is governed by a flexible-use, three-tiered spectrum authorization framework to accommodate a variety of commercial uses on a shared basis with incumbent federal and non-federal users of the band. The three tiers are: Incumbent access, PAL, and GAA.

Incumbent access users include authorized federal users, grandfathered Fixed Satellite Service earth stations, and, for a limited time, grandfathered wireless broadband licensees in the 3650-3700 MHz portion of the band. These users will be protected from harmful interference from PAL and GAA users through exclusion zones and management SAS and ESC services. These capabilities are a prerequisite for commercial CBRS service. Commercial services on the GAA portion of the band are expected to begin in September 2019<sup>10</sup>.

The priority access tier consists of PALs that will be assigned using competitive bidding within the 3550-3650 MHz portion of the band. The auction process has been defined with input from multiple rounds of public comment in the last two years. Each PAL is defined as a 10 year renewable authorization to use a 10MHz channel within a county. Up to seven total PALs may be assigned in any given county with up to four PALs going to any single applicant, or 40MHz of spectrum per county.

The GAA tier is licensed to permit open and flexible access to the band for a wide group of potential users including non-tradition operators and enterprises. GAA users are permitted to use any portion of the 3.5 GHz band not assigned to a higher tier user. The GAA users may use unused, previous assigned priority access channels. The SAS and ESC will help manage use across the geographic license area for each PAL. The SASs should provide aggregated spectrum usage data to the public upon request.

### **2.1.1. CBRS Deployment and Business Case Considerations**

CBRS PALs have features that make them attractive for predictable business planning. They do have a renewable 10-year term with the ability to partition and disaggregate PALs for more flexible use. The county-level, 10-year term makes CBRS more predictable from a business planning perspective.

---

<sup>7</sup> <https://www.fcc.gov/wireless/bureau-divisions/mobility-division/35-ghz-band/35-ghz-band-citizens-broadband-radio-service>


<sup>8</sup> <https://www.fcc.gov/document/fcc-acts-increase-investment-and-deployment-35-ghz-band-0>

<sup>9</sup> <https://www.google.com/get/spectrumdatabase/sas/>

<sup>10</sup> [https://www.cbbsalliance.org/news/cbbs-alliance-to-launch-ongo-commercial-services-in-3-5-ghz-cbbs-band/](https://www.cbrsalliance.org/news/cbbs-alliance-to-launch-ongo-commercial-services-in-3-5-ghz-cbbs-band/)

However, they are still subject to preemption if a higher-tier user needs the spectrum. Mission-critical use cases should be mindful of preemption.

CBRS can be deployed as a rural and suburban FWA broadband solution. A 50Mbps capable FWA service should be possible at inter-site distances (ISD) of 5-7km, Figure 3, based on preliminary data<sup>11</sup>. PALs are appropriate for this type of FWA service to prevent interference and provide reliable service. This level of service is suitable for rural and suburban markets where the subs per sq-km is <5,000. The TDD spectrum can be tuned to favor the downlink (DL) and better match DL/Uplink (UL) usage patterns in a given area. This gives the flexibility to tune the sector for business customers who may have more symmetric DL/UL ratios. Enterprises can also create their own private networks using CBRS spectrum assuming a build on GAA channels or unused PALs.



**Figure 3. Illustrative 3.5GHz CBRS deployment in a low to mid-density area to achieve 50Mbps service with 6.5km ISD**

Source: Ericsson, inCode

## 2.2. C-Band 3700-4200MHz Overview

C-band spectrum has traditionally been reserved for satellite downlink and uplink. Earth stations are still in use, but unused spectrum can be put to use with a carefully crafted clearing and migration plan. The C-Band Alliance (CBA), a coalition of incumbent satellite companies which includes Intelsat, SES, Eutelsat, and Telesat, is the organization representing the interests of the incumbents and works closely with the FCC and industry to ensure all interests are represented.

CBA has been in active negotiations regarding how much spectrum could be made available. Up to 500MHz is available, but the CBA is signaling 150-200MHz could be available for auction. CBA proposed a plan in 2019 to migrate satellite customers to a narrow portion of the C-band spectrum to free

---

<sup>11</sup> Ericsson

up spectrum for a C-band auction. The NPRM for C-Band spectrum may have been approved at the FCC's July 2018 Open Meeting<sup>12</sup> under the intention of a flexible use spectrum policy, but progress has been slow to incumbent pushback and multiple public commentary periods<sup>13</sup>.

### **2.2.1. C-Band Deployment and Business Case Considerations**

The main challenge in C-band deployment is coming up with a spectrum clearing plan that works for both the incumbent satellite interests and the wireless industry interests. Spectrum clearing is complicated by needing to fund and launch satellites and potential migrations to fiber transit, in addition to finalizing an incumbent user migration plan.

Spectrum acquisition cost for the C-band spectrum could be 2x the PAL auction clearing price. The spectrum is valuable because of its high channel widths, up to 100MHz. Auction prices for C-band spectrum licenses could be in the \$0.25-0.30/MHz-pop range for 100MHz licenses. Larger blocks of spectrum with motivated bidders could drive this up to \$0.40/MHz-pop, but will likely remain around \$0.25/MHz-pop due to the staged nature of the spectrum clearing and availability.

The C-band 3.7-4.2 GHz spectrum is actually the C-band downlink. There is an associated 5.925-6.425 GHz band known as the C-Band uplink which could also be in play. This 500MHz uplink is within the 1.2GHz wide 5.925-7.125 GHz ("6 GHz") band being considered for unlicensed spectrum use as part of the October 2018 NPRM<sup>14</sup>. Addition of these adjacent bands could make the currently planned C-band spectrum more valuable.

The cell site grid would be on par with the CBRS grid, but still denser than the 2.5GHz grid. Likely deployment approach would be to collocate C-band radios on existing macro or small cell towers and densify as needed based on pop density and cell edge performance specs. C-band would still be a lower deployment cost than mmWave and could still provide 500-1000Mbps service. C-band FWA service would benefit from outdoor mounted antennas for customer premise equipment (CPE). The band will like perform similar to CBRS, though the allowed power has not been set yet, and going from outdoor to indoor antennas could drop the DL line rate from 500 Mbps to 100 Mbps with the additional path loss.

### **2.3. BRS/EBS 2496-2690 MHz Overview**

The 2.5GHz Broadband Radio Service / Education Broadband Service (BRS/EBS) spectrum is the single largest band of contiguous spectrum (194MHz) below 3GHz. This band was historically reserved for educational TV and Tribal Nations, but much of the spectrum has gone unused for more than twenty years, particularly in rural areas. Educational institutions largely use the internet for their broadcast needs. Sprint is the largest BRS/EBS spectrum lessee in the United States, holding nearly 80% of license leases. Sprint is the largest holder of license leases in the BRS band.

The 2.5 GHz band, which extends from 2496-2690 MHz, is comprised of 20 channels allocated for EBS, 13 channels allocated for commercial BRS, and associated guard band channels<sup>15</sup>. EBS licensees operate in 114MHz of the 2.5GHz band; the remaining 80MHz is assigned to the BRS, totaling 194MHz.

Currently, there are 1,300 EBS licensees holding 2,193 licenses. Many of these licensees don't use their spectrum, but the rules allow them to lease out their excess capacity to non-educational entities to use for

---

<sup>12</sup> <https://www.fcc.gov/news-events/events/2018/07/july-2018-open-commission-meeting>

<sup>13</sup> <https://ecfsapi.fcc.gov/file/0719066596388/DA-19-678A1.pdf>

<sup>14</sup> <https://transition.fcc.gov/oet/ea/presentations/files/oct18/3.1-Rulemakings-JSP.PDF>

<sup>15</sup> <https://docs.fcc.gov/public/attachments/DOC-358065A1.pdf>

non-educational purposes, e.g. Sprint. Sprint uses over 1000 BRS licenses and leases approximately 1500 of the 2,193 EBS licenses, or 68% of all EBS licenses at 2.5 GHz to service as the workhorse spectrum of its tri-band 3G/4G LTE network using 800 MHz, 1.9 GHz, and 2.5 GHz. These leases are authorized to have terms of up to 30 years and can contain rights of first refusal or purchase options helping Sprint's position when negotiating on distribution terms for the remaining leases and overlays.

### **2.3.1. BRS/EBS Deployment and Business Case Considerations**

The 2.5GHz band can currently support significant capacity and throughput where deployed, including two- channel and three-channel carrier aggregation (CA), 2CA and 3CA, for up to 40MHz and 60MHz of spectrum, respectively. These CA profiles can provide 200-500Mbps of DL capacity on Sprint's existing site grid size.

Sprint is pushing to adopt counties, or Basic Trading Areas (BTA), and Partial Economic Areas (PEA), as the appropriate geographic service area unit for new overlay licenses. Almost all of the EBS license areas today are misshapen and irregularly configured which makes it difficult for deployment planning. This will benefit any of the new licensees or lessees as the network grid planning will be more predictable. The site deployment cost will be comparable to today's macro site colocation costs.

Spectrum acquisition costs for 2.5GHz have risen in the past 2 years as it became apparent that 5G deployments were in need of more mid-band spectrum. Sprint's 2.5GHz holdings were valued at \$0.27/MHz-pop in 2016. Auction prices for the remaining licenses from the total 194MHz of BRS/EBS spectrum could be similar to Sprint's current 2.5GHz holdings valuation at \$0.50/MHz-pop<sup>16</sup> if valued today. Future auctions for the remaining licenses could see a migration to a band plan with two sizes of licenses: a 100MHz block and a 16.5MHz block. This is still tbd, but makes the auction more valuable given the high channel widths that could become available in the mid-band.


## **3. Mid-band Use Case and Competitive Considerations**

5G use cases have a chicken-and-the-egg dilemma. Will the use cases wait for the technology in its final form or will the technology continue to be designed to fit the envisioned use cases? We have an idea of what use cases could take off, but we are probably wrong. With 4G and LTE no one anticipated the age of apps that was created when the industry transitioned from 3G voice to 4G data. The iPhone release in 2007 ushered in the age of apps by rewarding innovation with capable mobile infrastructure. The 5G era will foster innovation, but the mash up of creativity will be driven by the capabilities of this new network, Figure 4. ITU defined the capabilities of 5G networks in the IMT-2020 standard. For now, we can imagine what a two-order of magnitude improvement in traffic capacity could look like and then match it up to use cases like immersive VR and 8k eMBB. The end use cases may be off, but the functionality is likely directionally right.

Near term use cases for mid-band spectrum on 5G will include improved mobile broadband, Mobile Virtual Network Operator (MVNO) offload, private networks, and FWA. CBRS is available now for use cases and more will be available following the PAL auction. Intra and inter-band carrier aggregation with other mid-band spectrum assets as they come on line will further complement these initial use cases.

---

<sup>16</sup> J.P. Morgan estimates.



M.2083-01

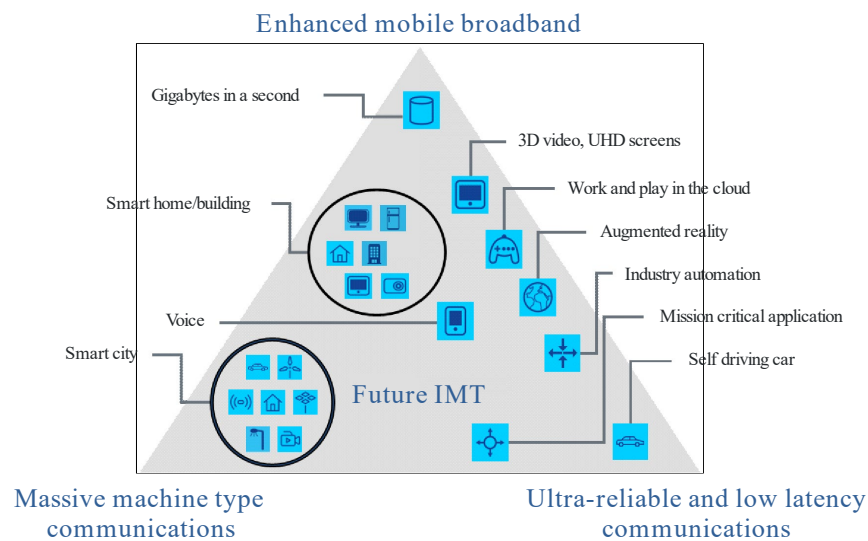
**Figure 4. Enhancement of key capabilities from IMT-Advanced to IMT-2020 with a perspective on mid-band's ability to deliver the capability.**

Source: Rec. ITU-R M.2083-0: IMT Vision - "Framework and overall objectives of the future development of IMT for 2020 and beyond"

The chart above depicts the expected increase in capabilities of the 5G network (IMT 2020) vs the 4.5 G network (IMT 2015). Peak data rates from 1-10 Gbps, high density IoT and single-digit latency are hallmarks of 5G. Improved traffic density and lower power consumption will also spur innovation in the realm of the possible with 5G networks. Mid-band spectrum is able to deliver all of these 5G capabilities at a high level, but maybe not as good as high and low band counterparts. High band spectrum with its 400MHz channels scores well for bandwidth, area capacity, and spectral efficiency, but lower for energy efficiency. Low band spectrum doesn't score well for capacity, but does well with the coverage metrics like Mobility. Mid-band spectrum balances both capacity and coverage metrics.

ITU had a use case view of the future of 5G that started with the overall objectives as simplified by the pillars of eMMB, Ultra Reliable Low Latency Communications (URLLC), and massive Machine Type Communications (mMTC), Figure 5. This way of defining the experience first helped shape how to design the technology.





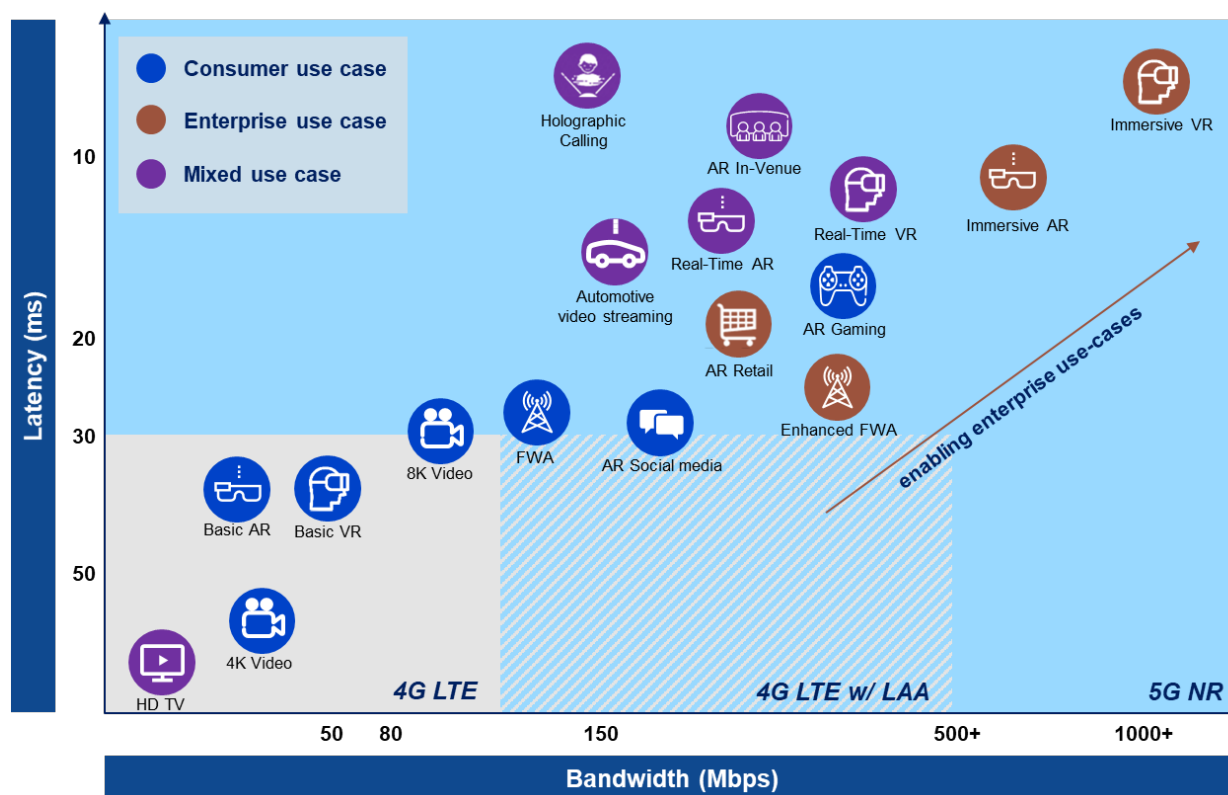
M.2083-02

**Figure 5. Usage scenarios of IMT for 2020 and beyond**

*Source: M.2083 : IMT Vision - "Framework and overall objectives of the future development of IMT for 2020 and beyond"*

5G use cases need capable spectrum, often in 100MHz and greater blocks. Mid-band spectrum isn't required for 5G, but it helps. Mid-band spectrum can support 5G capabilities up to 5-7km inter-site distances on 20, 40, and 100MHz wide channels. Cell edge spectral efficiencies of 3-9b/s/Hz can deliver gigabit speeds with up to single digit latency as determined by the path loss. This combination enables a multitude of eMBB, FWA, AR, and VR use cases, Figure 6.






**Figure 6. Latency and bandwidth requirements to deliver mobility use cases**

Source: inCode analysis

Current 4G networks are able to support eMBB data sessions and basic AR/VR capability, but they lack the low-latency capabilities to drive these and related use cases forward. Mid-band spectrum is better suited for use cases requiring high mobile data usage over a wide area. The combination of wide spectrum channels and capable signal propagation eases network design and planning constraints when building for these types of use cases.

**eMBB** The eMBB use case shines with mid-band spectrum on 5G. It builds on top of what is already capable for improved mobile data services. Real world eMBB results in an urban environment show 300-500Mbps is capable now on B4 LTE with B46 licensed-assisted access (LAA), giving optimism to 5G FWA use cases on mid-band spectrum. The LAA unlicensed band support benefits eMBB and high data usage use cases. LAA provides support for up to two licensed plus three unlicensed carriers enabling speeds above 1 Gbps<sup>17</sup>. T-Mobile US is using LAA capabilities today to transmit 4G LTE signals over unused 5GHz Wi-Fi channels to widen its available bandwidth when aggregated with its existing spectrum, Figure 7. T-Mobile is using 2CA up to 5CA carrier aggregation profiles of combine 20MHz channels of its 1700MHz/2100Mhz AWS Band 4 spectrum with three 20MHz channels of U-NII band 5GHz LAA spectrum. Mobile devices connect to the cell site and send UL data using the B4 spectrum, and then get DL data from a combination of B4 and LAA.

<sup>17</sup> Source: LAA, Ericsson - <https://www.ericsson.com/en/portfolio/networks/ericsson-radio-system/radio-system-solutions/licensed-assisted-access>



**Figure 7. Urban eMBB use case on 20MHz of B4 LTE and 60MHz of B46 LAA**

*Source: Ookla, T-Mobile LAA profile in NYC demonstrating 300Mbps mobile broadband with 4CA of 20MHz of B4 and 60MHz of B46.*

T-Mobile LAA on LTE today provides 300-500Mbps mobile broadband speeds on a 4CA to 5CA carrier aggregation profile with U-NII band B46 spectrum and its licensed bands. 5G NR on Rel.16 works to move LAA capabilities up to 1 Gbps with greater carrier aggregation and improved spectral efficiency.

**FWA**, is another use case where high bandwidth over a wide area is needed. We can look at rural, suburban, and urban FWA use cases with targeted DL bandwidth of 50, 100, and 300 Mbps, respectively, to evaluate the suitability of mid-band spectrum for this application.

**Rural FWA** deployments are typically low-density RAN buildouts on macro towers with line-of-sight (LoS) top-down deployments. Delivering a 50Mbps DL product with 50-75ms latency can be done on 4G LTE with low band TDD spectrum today with indoors CPE given a couple 10MHz channels. This approach is more challenged with higher adoption rates, CPE shift from the edge to the center of the home, or the served location is at the cell edge. Mid-band spectrum on 5G can help provide wider channels with improved spectral efficiency for better cell edge performance. CPE should be mounted outdoors to prevent excessive path loss. 40-60MHz of mid-band spectrum should be sufficient.

**Suburban FWA** is a medium-density RAN build but still on macro towers with LoS. Competitive FWA offers should be able to provide a 100Mbps product with 50ms latency. Mid-band spectrum with 3-5b/s/Hz cell edge spectral efficiency and 2-3 simultaneous attached users (SAU) per sector translates to a need for 100MHz of mid-band spectrum minimum. This is achievable in the US with just over 500MHz of mid-band spectrum available for use in the next few years. If the operator has ~20-40MHz of spectrum available from existing holdings then ~60-80MHz of new spectrum would be needed to enable this use case pending available carrier aggregation profiles. Mid-band spectrum on 5G is the best approach here as mmWave would be too expensive to build out in suburban areas.

**Urban FWA** use cases are typically high density RAN builds, often with small cells and non-line-of-sight situations. Offering a 300Mbps+ product with 5-10ms latency requires wide channels and a densified network site plan. mmWave spectrum with 3-7b/s/hz cell edge spectral

efficiency and 4-5 SAUs per sector translates to a required range of 300-500MHz total mmWave spectrum width less what can be aggregated from existing holdings. This is very likely with a mmWave 5G deployment, but the sites would need to have a <250m ISD to allow for handoffs without downgrading to 4G. This could be done with mid-band spectrum given the right channel width holdings or augmented with three B46 LAA channels. A mid-band build in an urban area would be at a lower cost than then much denser mmWave network build. Operators with owners economics on existing dense fiber networks in urban areas can offset some of the build cost for mmWave networks.

## **4. Business Model Considerations**

### **4.1. Cost Considerations**

The two main categories associated with mid-band spectrum deployment are getting the spectrum, and putting it to use. Each of these and their associated drivers go in the business model when planning for deployment

Spectrum acquisition costs, \$/MHz-pop or licensing costs, will be a key entry cost for mid-band spectrum deployment. However, the largest cost could be cell site densification to support mid-band spectrum pending the operator's existing grid. The existing tower grid for US operators average 12km between their macro tower sites outside of urban centers. This ISD drops to 0.6km for tower sites in and around urban centers on average. The 12km ISD works for 600MHz, 700MHz, Cellular, AWS, and PCS frequencies, but B41 2.5GHz would require densification to maintain higher data rates in rural deployments. The B41, CBRS B48 3550-3700, and higher mid-band spectrum should have ISDs in the 1-4km range depending on the cell edge performance needs. Some densification may be needed outside of urban centers, but existing sites in most urban centers should have adequate ISDs to provide both coverage and capacity on mid-band spectrum. Densification provides the added benefit of more spectrum reuse opportunities for efficient radio resource use for cell area covered pops.

Transmit power regulation on mid-band spectrum is a key factor in setting the grid size. CBRS has class B base stations for outdoor use which have a maximum Equivalent Isotropically Radiated Power (EIRP) of 47 dBm or about 50 watts. Unlicensed U-NII WiFi mid-band spectrum for outdoor point-to-point and point-to-multipoint is governed by FCC Part 15 rules<sup>18</sup> and maxes out at 30dBm or about 1 watt. Current licensed spectrum in the 2GHz range is allowed to transmit at a higher EIRP of up to 62 dBm. This translates to about 6-7km lower ISD for CBRS implying densification needs for macro towers.

Any densification needed for mid-band spectrum is far less than the densification needed for mmWave spectrum deployments. Spectrum starting in the 24GHz and 28GHz frequencies will need ISDs on the order of city blocks, or <200m, as seen with Verizon's 5G April 2019 launch in downtown Chicago.

### **4.2. Mid-band Deployment Cost Drivers**

There are many variables to consider when evaluating the operational and business costs of mid-band spectrum deployment. We'll look at a few key primary and secondary drivers for mid-band spectrum deployment cost variability.

---

<sup>18</sup> [eCFR Part 15](#)

#### **4.2.1. Primary drivers for mid-band deployment cost variability**

1. **Spectrum band** – Propagation and path loss characteristics vary by frequency. Migration from 600MHz to 2.5GHz to 3.5GHz to 24GHz bands drops the ISD by 59%, 22%, and 75%, sequentially. A cell site grid set up for 2.5GHz can more easily colo 3.5GHz radios than mmWave radios. Mid-band spectrum would need 2-3x densification over low band, while mmWave would need 10-12x densification over low band spectrum<sup>19</sup>. Each band has its corresponding transmit power requirements for base stations. Lower power means more cells sites are needed.
2. **Channel width** - Impacts the number of sites needed for densification. Depends on cell edge traffic throughput requirements. Ranges from 5MHz FDD up to 100Mhz TDD wide channels. FDD vs TDD impacts the asymmetry of a band where paired spectrum cannot respond to DL/UL asymmetry. Aim for 100MHz of mid-band spectrum to deliver many of the 5G use cases. CBRS will have up to 70MHz available for PAL with flexible use provisions for additional channels. The C-band 3.7-4.2GHz could free up 200MHz channels for operators pending release details. inCode's view is it will take a combination of low, mid, and high band spectrum to deliver the full capabilities of 5G, but many use cases will benefit from having at least 100 MHz of mid-band spectrum available for use. These 100MHz and wider channels make massive multiple-input and multiple-output (MIMO) possible.
3. **Carrier aggregation profiles** – CA profiles supported by 3GPP and the operator control what frequencies are available for signal modulation. Greater intra-band and inter-band CA aggregation creates wider channels and lowers the deployment costs.
4. **Spectral efficiency** - Cell edge vs cell best-case differs greatly due to path loss. Cell edge can vary from 1b/s/Hz to 7b/s/Hz now while 7-11b/s/Hz is achievable closer to the cell site. Holding other factors constant, moving from 2.5GHz spectrum to 24GHz spectrum increases the spectral efficiency 57% due to supporting higher orders of modulation on wider channel widths. Low band spectrum blocks typically have lower channel widths where side guardbands are a higher percentage of the overall channel width. Lower spectral efficiency means more cell sites are needed for a given level of service.

#### **4.2.2. Secondary drivers for mid-band deployment cost variability**

1. **Indoor vs outdoor antennas** - Mid-band spectrum will struggle with in-building penetration due to its power and propagation characteristics. Best to plan for outdoor antenna mounts for fixed wireless applications and use LAA profiles for mobility use cases. Mid-band spectrum fairs better than mmWave when it comes to going through windows and walls. Millimeter wave spectrum is blocked by the thin, reflective coating on low emissivity, or low-e, glass preventing indoor usage of 5G devices on these bands. Even 2.5GHz spectrum struggles to get into buildings. The 3.5GHz CBRS spectrum can get 50Mbps with outdoor mounted antennas in a rural deployment, but this number drops to less than 10Mbps when the CPE is placed indoors. Indoor CPE path loss is further hindered by CPE placement, often in the center of the structure to aid the WiFi signal emitted from the CPE. Mobile UEs have the same problem and require increase radio resource allocation to maintain a given DL/UL bit rate
  - a. A correctly installed outdoor CPE is directed to the best serving cell site, leading to a lower link budget path loss and increasing the value of mid-band and mmWave TDD spectrum. Macro sites above the canopy or small cells below the canopy can be LoS to outdoor antennas. This is the recommended approach.

---

<sup>19</sup> Ericsson

- b. Outdoor antennas deliver a large gain in signal quality as a result of the 10dB difference in antenna gain and the avoidance of 10–15dB in wall or window attenuation losses suffered by indoor devices.
  - c. Another contributor to signal attenuation in indoor devices is the deep indoor loss, as the device is likely to be placed in a hidden location or interior location to provide optimum Wi-Fi coverage. This could contribute another 5dB in path loss.
  - d. An indoor CPE is comparable to a smartphone in terms of spectrum efficiency. An outdoor antenna / indoor CPE combo is 2-3x more efficient. For the same data consumption rate, around 2-3x as many HHs can be served using outdoor rather than indoor units. And consequently, 2-3x as much spectrum is needed to serve indoor-only FWA households.
2. **LoS vs nLoS** – Line of sight is preferred, but not always possible. Outdoors non-line of sight (nLoS) benefits from beamforming and more capable phase array antennas. Multi-path MIMO improvements also benefit the higher ISD of mid-band spectrum
3. **Urban vs rural** – Area traffic capacity support improves two orders of magnitude with the migration to 5G. Mid-band spectrum propagates this benefit over a broader area. Urban areas are better positioned to benefit from 5G, but it requires densification. Urban clutter reflectance makes RF planning easier for mid to high bands because of beam forming and steering capabilities.
4. **DL loading** – Site densification varies with DL loading. Densification is required once DL loading approaches 100%. RF resource blocks get depleted at the site due to traffic increase and more users and traffic equals more required resource blocks.
5. **Busy hour traffic hours** – Sets the triggers for bandwidth management and service expectations. Determines worse-case capacity scenarios and minimum planning thresholds.
6. **Device and handset compatibility and availability** – Business cases on mid-band spectrum should consider UE availability for a given spectrum band. Baseband compatibility for 2.5GHz devices is more available than CBRS and C-band device due to Sprint's presence in that band for many years. There are 100s of 2.5GHz SKUs, but only 10s of CBRS SKU's in the market today<sup>20</sup>. The Google Pixel3 and Samsung S10 support CBRS.



### 4.3. Mid-band Spectrum Revenue Benefit

Revenue benefits from mid-band spectrum 5G services will likely be driven by eMBB use cases at first followed by cross-sell and up-sell to experience-based benefits enabled by 3GPP Rel.16. To date, the US market has not successfully monetized 5G. Mainly because it is still seen as more-capable eMBB, but still unproven. SK Telecom has been successful at monetizing 5G based on the experience benefit it provides to gaming packages. Mid-band spectrum for 5G can drive some of this experience benefit by completing the low / mid / high band mosaic for service delivery and ensure service continuity for coverage, capacity, and latency needs. CBRS is called the Innovation Band. New business models and revenue follow innovation.

<sup>20</sup> <https://www.gsmarena.com>

## Conclusion

In conclusion, mid-band spectrum for 5G is justified in having the just-right Goldilocks moniker. It is at the sweet spot of coverage and capacity for 5G deployments. The challenge will be getting enough 2.5-8GHz spectrum out into the US market in time for the 5G use cases to benefit. Many 5G use cases will need a combination of >100MHz spectrum channels, single-digit latency, and wide area coverage. Mid-band spectrum in the US is capable of meeting these criteria pending its availability schedule. The FCC is working with the industry to secure the deployment of additional mid-band spectrum to close this gap in the US spectrum portfolio. Companies should pursue mid-band spectrum acquisition to complement their existing spectrum portfolios.

# Abbreviations

3GPP	3rd Generation Partnership Project
5G	5 <sup>th</sup> generation cellular network technology
AR	Augmented Reality
AWS	Advanced Wireless Services
BRS/EBS	Broadband Radio Service / Education Broadband Service
BTA	Basic Trading Area
CA	Carrier Aggregation
CBA	C-Band Alliance
CBRS	Citizens Broadband Radio Service
CPE	Customer Premises Equipment
DL	Downlink
EDGE	Enhanced Data rates for GSM Evolution
EIRP	Equivalent Isotropically Radiated Power
eMBB	enhanced Mobile Broadband
ESC	Environmental Sensing Capability
FAST	Facilitate America's Superiority in 5G Technology
FCC	Federal Communications Commission
FDD	Frequency Division Duplex
FWA	Fixed Wireless Access
GAA	General Authorized Access
GPRS	General Packet Radio Service
GSM	Global System for Mobile Communications
HH	Households
HSPA	High Speed Packet Access
IMT	International Mobile Telecommunications
IoT	Internet of Things
ISD	Inter-Site Distance
ITU	International Telecommunication Union
LAA	License Assisted Access
LoS	Line of sight
LTE	Long-Term Evolution
MIMO	Multiple-Input and Multiple-Output
mMTC	massive Machine Type Communications
mmWave	Millimeter Wave
MSO	Multiple-System Operator
MVNO	Mobile Virtual Network Operator
nLoS	Non-line of sight
NPRM	Notice of Proposed Rulemaking
NR	New Radio
PAL	Priority Access License
PCS	Personal Communications Service
PEA	Partial Economic Area
Pop	Population
RAN	Radio Access Network

SAS	Spectrum Access System
SAU	Simultaneously Attached Users
SKU	Stockkeeping unit
SMR	Specialized Mobile Radio
TDD	Time Division Duplex
UE	User Equipment
UL	Uplink
UMTS	Universal Mobile Telecommunications Service
U-NII	Unlicensed National Information Infrastructure
URLLC	Ultra Reliable Low Latency Communications
V2X	Vehicle to everything
VR	Virtual Reality

## Bibliography & References

FCC (July 19, 2019). *Wireless Telecommunications Bureau, International Bureau, Office of Engineering and Technology, and Office of Economics and Analytics Seek Focused Additional ) Comment in 3.7-4.2 GHz Band Proceeding*. Retrieved from <https://ecfsapi.fcc.gov/file/0719066596388/DA-19-678A1.pdf>

NAB (August 14, 2019). Expanding Flexible Use of the 3.7 GHz to 4.2 GHz Band. Retrieved from <https://ecfsapi.fcc.gov/file/108140216309684/C-band%20focused%20reply%20comments.pdf>

Verizon (August 14, 2019). *Wireless Telecommunications Bureau, International Bureau, Office of Engineering and Technology, and Office of Economics and Analytics Seek Focused Additional ) Comment in 3.7-4.2 GHz Band Proceeding*. Retrieved from [https://ecfsapi.fcc.gov/file/1081494128931/Verizon%20Public%20Notice%20Reply%20Comments%20\(8.14.19\).pdf](https://ecfsapi.fcc.gov/file/1081494128931/Verizon%20Public%20Notice%20Reply%20Comments%20(8.14.19).pdf)

FCC (August 3, 2017). *FCC Opens Inquiry Into New Opportunities In Mid-Band Spectrum*. Retrieved from <https://www.fcc.gov/document/fcc-opens-inquiry-new-opportunities-mid-band-spectrum-0>

FCC (April 12, 2019). *The FCC's 5G FAST Plan*. Retrieved from <https://www.fcc.gov/5G>

FCC (April 21, 2015). Amendment of the Commission's Rules with Regard to Commercial Operations in the 3550- ) 3650 MHz Band. Retrieved from <https://docs.fcc.gov/public/attachments/FCC-15-47A1.pdf>

FCC (June 19, 2019). Transforming the 2.5 GHz Band. Retrieved from <https://docs.fcc.gov/public/attachments/DOC-358065A1.pdf>



# Low Latency DOCSIS

## Overview And Performance Characteristics

A Technical Paper prepared for SCTE•ISBE by

**Greg White**

Distinguished Technologist  
CableLabs  
g.white@cablelabs.com

**Karthik Sundaresan**

Distinguished Technologist  
CableLabs  
k.sundaresan@cablelabs.com

**Bob Briscoe**

Consultant  
CableLabs  
research@bobbriscoe.net

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
Low Latency DOCSIS.....	4
1. Latency in DOCSIS Networks.....	5
2. New Dual-Queue Approach.....	7
2.1. Low-Latency Aggregate Service Flows .....	8
2.2. Identifying NQB Packets—Default Classifiers .....	10
2.3. Coupled AQM.....	11
2.4. Queue Protection.....	11
3. Upstream Scheduling Improvements.....	12
3.1. Faster Request Grant Loop .....	13
3.2. Proactive Grant Service .....	13
4. Low Latency DOCSIS Performance .....	14
4.1. Traffic Models.....	14
4.2. Round Trip P99 Latency based on Traffic Mix.....	15
4.3. Round Trip Packet Delay and Packet Delay Variation Statistics .....	16
4.4. Proactive Grant Service Tradeoffs .....	17
4.5. Multiple Simultaneous Game Sessions .....	18
4.6. Uni-directional LLD .....	18
5. Deployment Considerations .....	19
5.1. Device Support.....	19
5.2. Packet Marking.....	20
5.3. Provisioning Mechanisms .....	20
5.3.1. Aggregate QoS Profiles.....	21
5.3.2. Migration Using Existing Configuration File and Service Class Name .....	21
5.3.3. Explicit Definition of ASF in the Configuration File .....	22
5.4. Latency Histogram Reporting.....	22
Conclusion .....	22
Abbreviations.....	23
Bibliography & References .....	24
Appendix A - Low Latency <i>and</i> High Bandwidth: L4S .....	25
Appendix B - Simulation Details.....	27

## List of Figures

Title	Page Number
Figure 1. Sources of Latency in DOCSIS 3.1 Networks .....	6
Figure 2. Example of Traditional and LLD Service Configurations .....	9
Figure 3. ASFs, SFs, and Classifiers.....	10
Figure 4. Coupled AQM across Service Flows .....	11
Figure 5. Queue Protection function.....	12
Figure 6. Request Grant Delay in DOCSIS Networks .....	12
Figure 7. MAP Interval and MAP Processing Time .....	13

Figure 8. Proactive Grant Service in LLD .....	14
Figure 9. Overall Round-Trip Latency Performance for NQB-Marked Traffic.....	17
Figure 10. Impact of PGS on NQB Traffic Latency .....	18
Figure 11. Impact of Multiple Simultaneous Game Sessions .....	18
Figure 12. Impact of Uni-directional LLD .....	19
Figure 13. ASFs, SFs, and the Low Latency DOCSIS Architecture .....	21
Figure 14. Example of Using a Latency Histogram Report to Estimate Latency Quantiles .....	22
Figure 15. NS3 Network Topology .....	27

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1. Evolution of Latency Performance in DOCSIS Networks (Round-Trip Time in milliseconds between the CM and CMTS) .....	5
Table 2. Background Traffic Mixes .....	15
Table 3. 99 <sup>th</sup> Percentile round-trip latency for nqb-marked traffic between the CM and CMTS .....	15
Table 4. Device Dependencies for LLD Features .....	19

# Introduction

Low Latency DOCSIS technology (LLD) is a specification developed by CableLabs in collaboration with DOCSIS vendors and cable operators that tackles the two main causes of latency in the network: queuing delay and media acquisition delay. LLD introduces an approach wherein data traffic from applications that aren't causing latency can take a different logical path through the DOCSIS network without getting hung up behind data from applications that are causing latency, as is the case in today's Internet architectures. This mechanism doesn't interfere with the way applications share the total bandwidth of the connection, and it doesn't reduce one application's latency at the expense of others.

In addition, LLD improves the DOCSIS upstream media acquisition delay with a faster request-grant loop and a new proactive scheduling mechanism. LLD makes the internet experience better for latency sensitive applications without any negative impact on other applications.

Users of current DOCSIS 3.1 equipment experience typical round trip latency performance of around 10 ms on the Access Network link. However, with normal usage patterns, the link can experience delay spikes of 100 ms or more. LLD systems can deliver a consistent 1 ms delay on the DOCSIS network for traffic that isn't causing latency, imperceptible for nearly all such applications. The experience will be more consistent with much smaller delay variation.

LLD functionality can be deployed by field-upgrading DOCSIS 3.1 cable modem and cable modem termination system devices with new software.

The technology includes tools that enable automatic provisioning of these new services, and it also introduces new tools to report statistics of latency performance to the operator.

Cable operators, DOCSIS equipment manufacturers, and application providers will all have to act in order to take advantage of LLD technology. This paper explains the technology and describes the role that each of these parties plays in making LLD technology a reality.

## Low Latency DOCSIS

Let's begin with bandwidth (or "speed"): the amount of data that can be delivered across a network connection over a period of time. Sometimes bandwidth is very important to the broadband experience, particularly when an application is trying to send or receive large amounts of data, such as watching videos on Netflix, downloading videos/music, syncing file-shares or email clients, uploading a video to YouTube or Instagram, or downloading a new application or system update. Other times, bandwidth (or bandwidth alone) isn't enough, and latency has a big effect on the user experience.

Latency is the time that it takes for a short message (a packet, in networking terminology) to make it across the network from the sender to the receiver and for a response to come back. Network latency is commonly measured as round-trip-time and is sometimes referred to as "ping time." Applications that are more interactive or real-time, like web browsing, online gaming, and video conferencing/chatting, perform the best when latency is kept low, and adding more bandwidth without addressing latency doesn't make things better.

When multiple applications share the broadband connection of one household (e.g., several users doing different activities at the same time), each of those applications can have an impact on the performance of

the others. They all share the total bandwidth of the connection (so more active applications mean less bandwidth for each one), and they can all cause the latency of the connection to increase.

It turns out that capacity-seeking applications today (i.e. those that want to send a lot of data all at once) do a reasonably good job of sharing the bandwidth in a fair manner, but they actually cause a pretty big latency problem when they do it because they send data too quickly and expect the network to queue it up. We call these applications "queue-building" applications, e.g., video streaming (Netflix). There are also plenty of other applications that don't send data too quickly, so they don't cause latency. We call these "non-queue-building" applications, e.g., video chatting (FaceTime).

LLD separates these two types of traffic into two logical queues, which greatly improves the latency experienced by the non-queue-building applications (many of which may be latency-sensitive) without having any downside for the queue-building applications. In addition, two queues allow LLD to support next-generation application protocols that can scale up to sending data at 10 Gbps and beyond while maintaining ultra-low queuing delay, which means that in the future, there may not be queue-building applications at all.

The Low Latency DOCSIS specifications were published earlier this year as updates to the DOCSIS 3.1 core specifications ([MULPIv3.1], [CCAP-OSSIV3.1], [CM-OSSIV3.1]), and DOCSIS equipment manufacturers are working on building support for the functionality. In addition, work is underway in the Internet Engineering Task Force to standardize this low-latency architecture so that application developers across the broader Internet ecosystem can take advantage of it.

## 1. Latency in DOCSIS Networks


Low Latency DOCSIS technology is the next step in a progression of latency improvements that have been made to the DOCSIS specifications by CableLabs in recent years. Table 1 provides a snapshot of the milestones in round-trip latency performance with DOCSIS technology from the first DOCSIS 3.0 equipment to the new LLD, which achieves ~1 ms of round-trip latency for non-queue-building traffic. The table references three metrics that describe the range of latencies added by the DOCSIS network link that would be experienced by a broadband user. The first, "When Idle," refers to a broadband connection that is not being actively used by the customer. The second, "Under Load," represents average latency while the user is actively using the service (e.g., streaming video). Finally, the third, "99<sup>th</sup> Percentile," gives an indication of the maximum latency that a customer would commonly experience in real usage scenarios. The table uses order-of-magnitude numbers because the actual performance will vary based on a number of factors including DOCSIS channel configuration and actual application usage pattern.

For latency-sensitive applications, the 99<sup>th</sup> percentile value can significantly impact user experience.

**Table 1. Evolution of Latency Performance in DOCSIS Networks (Round-Trip Time in milliseconds between the CM and CMTS)**

	When Idle	Under Load	99 <sup>th</sup> Percentile
DOCSIS 3.0 Early Equipment	~10 ms	~1000 ms	~1000 ms
DOCSIS 3.0 w/ Buffer Control	~10 ms	~100 ms	~100 ms
DOCSIS 3.1 Active Queue Management	~10 ms	~10 ms	~100 ms
Low Latency DOCSIS 3.1	~1 ms	~1 ms	~1 ms

The latency described in Table 1 is caused by a series of factors in the DOCSIS cable modem (CM) and cable modem termination system (CMTS). Figure 1 illustrates the range of latencies caused by those factors in well-managed DOCSIS 3.1 networks.



**Figure 1. Sources of Latency in DOCSIS 3.1 Networks**

The lowest two latency sources in Figure 1 have minor impacts on overall latency.

The “**Switching/Forwarding**” delay represents the amount of time it takes for the CM and CMTS to make the decision to forward a packet. This has a very minor impact on overall latency.

The “**Propagation**” delay (the amount of time it takes for a signal to travel on the HFC plant) is set by the speed of light and the distance from CM to CMTS. Not much can be done to affect latency from this source.

Of the sources in Figure 1, the top three significantly drive latency performance.

The range of the “**Serialization/Encoding**” delay comes from the upstream and downstream channel configuration options available to the operator. Some of these configurations provide significant robustness benefits at the expense of latency, whereas others may be less robust to noise but provide very low latency. The LLD specification does not modify the set of options available to the operator. Rather, operators should be encouraged to use the lowest latency channel configurations that they can, given the plant conditions.

The “**Media Acquisition**” delay is a result of the shared-medium scheduling currently provided by DOCSIS technology, in which the CMTS arbitrates access to the upstream channel via a request-grant mechanism. The figure illustrates the range of delays that would be present on an uncongested channel.

The “**Queuing**” delay is mainly caused by the current TCP protocol and its variants. Applications today that need to seek out as much bandwidth as possible use a transport protocol like TCP (or the TCP-replacement known as QUIC), which uses a “congestion control” algorithm (such as Reno, Cubic, or BBR) to adjust to the available bandwidth at the bottleneck link through the network. Typically, this will

be the last mile link—the DOCSIS link for cable customers—where the bandwidth available for each application often varies rapidly as the activity of all the devices in the household varies.

With today's congestion control algorithms, the sender ramps up the sending rate until it's sending data faster than the bottleneck link can support. Packets then start queuing in a buffer at the entrance to the link, i.e. the CM or CMTS. This queue of packets grows quickly until the device decides to discard some newly arriving packets, which triggers the sender to pause for a bit in order to allow the buffer to drain somewhat before resuming sending. This process is an inherent feature of the TCP family of Internet transport protocols, and it repeats over and over again until the file transfer completes. In doing so, it causes latency and packet loss for all of the traffic that shares the broadband link.

To be clear, the “congestion” that is being controlled here (and the latency/loss that results from it) is completely independent from CM to CM within a DOCSIS service group. In other words, this is “self congestion” that is caused by and also experienced by the set of applications in use by a single subscriber, and is not related to “shared channel congestion” resulting from the heavy utilization of an oversubscribed service group.

LLD technology tackles the two main causes of latency in the network: queuing delay and media acquisition delay.

- LLD addresses Queueing Delay by allowing non-queue-building applications to avoid waiting behind the delays caused by the current TCP or its variants. At a high level, the low-latency architecture consists of a dual-queue approach that allows both queues to share single pool of bandwidth.
- LLD cuts Media Acquisition Delay by using a faster request-grant loop and by adding support for a new proactive scheduler that can provide extremely low latency service.

In addition, LLD introduces detailed statistics on queueing delay via histogram calculations performed by the CM (for upstream) and CMTS (for downstream).

Furthermore, CableLabs is working with a broad cross-section of stakeholders in the IETF to standardize an end-to-end service architecture that can leverage LLD to enable even high bandwidth TCP flows to achieve ultra-low queuing delay. This technology will be important for future, interactive high-data-rate applications like cloud gaming, holographic light field experiences, as well as for enabling higher performance versions of today's applications like web and video conferencing.

The sections below describe these features in more detail.

## 2. New Dual-Queue Approach

Of all the LLD features, the dual-queue mechanism has by far the greatest impact on round-trip latency and latency variation. The concept of the dual-queue approach is that the majority of the applications that use the internet can be divided into two categories:

- **Queue-Building Applications:** These application traffic flows frequently send data faster than the path between sender and receiver can support. The most common instance of queue-building flows are flows that use the current TCP or QUIC protocols. As discussed above, these capacity-seeking protocols use a legacy congestion control algorithm that probes for available capacity on the path by sending data faster than the path can support and expecting the network to queue the excess data in internal buffers. The majority of traffic (by volume) today is queue-building. Some

examples of queue-building applications are video streaming (e.g., Netflix, YouTube) and application downloads.

- **Non-Queue-Building Applications:** These application traffic flows very rarely send data faster than the path can support. They come in two subcategories:
  - Today's self-limited, non-capacity-seeking apps, such as multiplayer online games and IP communication apps (such as Skype or FaceTime). These applications send data at a relatively low data rate and generally space their packets out in a manner that does not cause a queue to form in the network.
  - Future capacity-seeking TCP/QUIC or UDP applications that adopt the new L4S congestion control algorithm (see Appendix A - Low Latency *and* High Bandwidth: L4S) and so can immediately respond to fast congestion signals sent by the network. These applications are still in development, as networks must first support L4S before applications are able to take advantage, but some prime candidates are cloud gaming, web browsing, cloud VR, and interactive light field experiences.

Queue-building (QB) application flows are the source of queuing delay, and today's non-queue-building (NQB) apps typically suffer from the latency caused by the QB flows.

The purpose of the dual-queue mechanism is to segment QB traffic from NQB traffic in a manner that can be readily implemented in DOCSIS 3.1 equipment and that doesn't alter the overall bandwidth of the broadband service.

By segmenting these two types of applications into separate queues, each can get optimal performance. The QB traffic can build a queue as it needs to in order to achieve the necessary and expected throughput performance, and the NQB traffic can take advantage of the available lower latencies by avoiding the delay caused by the QB flows. It is important to note that this segmentation of traffic isn't for purposes of giving one class of traffic benefits at the expense of the other—it isn't a high-priority queue and a low-priority queue. Instead, each queue is optimized for the distinct features and requirements of the two classes of traffic, enabling increased functionality and adding value for the broadband user. This is smart network management at work.

## 2.1. Low-Latency Aggregate Service Flows

DOCSIS 3.1 equipment, like equipment built against earlier versions of the specification, supports a number of upstream and downstream Service Flows (SFs). These Service Flows are logical pipes that are defined by their configured Quality of Service (QoS) parameters (most commonly, the rate shaping parameters [MULPIv3.1] that specify the speed of user connections). Each Service Flow carries a subset of the traffic to/from a particular CM, as specified by a set of packet classifiers configured by the operator. Traditionally, each Service Flow provides nearly complete isolation of its traffic from the traffic transiting other Service Flows (those on the same CM as well as those on other CMs)—each Service Flow has its own buffer and queue and is scheduled independently by the CMTS.

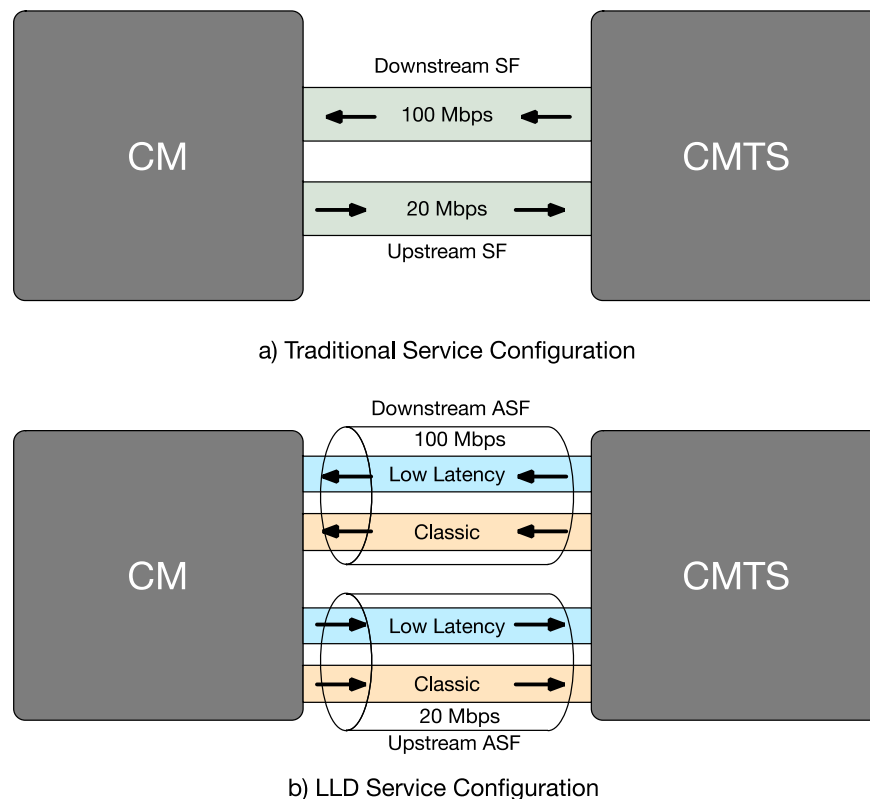
Typically, the operator defines a service offering to a customer via the configuration of a single upstream Service Flow and a single downstream Service Flow with rate shaping enabled, and all of the customer's traffic transits these two Service Flows.

The DOCSIS 3.1 specification already includes optional support in the CMTS for a mechanism to group any number of the Service Flows serving a particular CM. LLD leverages and extends this "Aggregate



Service Flow” (ASF) feature to establish (and group) a pair of Service Flows in each direction specifically to enable low-latency services. One of the Service Flows in the pair (the “Low Latency Service Flow”) will carry NQB traffic, and the other Service Flow (the “Classic Service Flow”) will carry QB traffic. The Aggregate Service Flow is configured for the service’s rate shaping setting, and the two constituent Service Flows inside the Aggregate have rate shaping disabled. The result is that the operator can configure the total aggregate rate of the service offering in each direction and does not have to configure (or even consider) how much of the user’s traffic is likely to be NQB vs QB.

Figure 2 illustrates an example configuration of broadband service as it might look in a current DOCSIS deployment, as well as how it would look with Low Latency DOCSIS service. In the traditional configuration, there is a single downstream Service Flow with a rate of 100 Mbps and a single upstream Service Flow with a rate of 20 Mbps. In the LLD configuration, there is a single downstream Aggregate Service Flow with a rate of 100 Mbps, containing two individual Service Flows, one for Low Latency traffic and one for Classic traffic. Similarly, there is single upstream Aggregate Service Flow with a rate of 20 Mbps, containing two individual Service Flows for Low Latency and Classic traffic.



**Figure 2. Example of Traditional and LLD Service Configurations**


The CMTS will enforce the Aggregate “Max Sustained Traffic Rate” (AMSR), and the end-user’s applications determine how much of the aggregate bandwidth they consume irrespective of which SF they use—just as they do today with a single DOCSIS SF.

As described later, Inter-Service-Flow scheduling is arranged to make the ASF function as a single pool of bandwidth.

## 2.2. Identifying NQB Packets—Default Classifiers

By default, the traffic within an Aggregate Service Flow is segmented into the two constituent Service Flows by a set of packet classifiers (Figure 3) that examine the Differentiated Services (DiffServ) Field and the Explicit Congestion Notification (ECN) Field, which are standard elements of the IPv4/IPv6 header [RFC3168]. Specifically, packets with an NQB DiffServ value<sup>1</sup> or an ECN field indicating either ECN Capable Transport 1 (ECT(1)) or Congestion Experienced (CE) will get mapped to the Low Latency Service Flow, and the rest of the traffic will get mapped to the Classic Service Flow.

The expectation is that non-queue-building traffic sources (applications) will either mark their packets with an NQB DiffServ value or support ECN.



**Figure 3. ASFs, SFs, and Classifiers**

Although the DiffServ Field is being used to indicate NQB behavior, that does not imply adoption of the Differentiated Services architecture as it is typically understood. In the traditional DiffServ architecture, applications indicate a desire for a particular treatment of their packets—often implemented as a priority level—which in essence conveys a value judgement as to the importance of that traffic relative to the traffic of other applications. Such an architecture can work just fine in a managed environment where all applications conform to a common view of their relative priority levels and so can be trusted to mark their packets appropriately. It fails, however, when applications need to send packets across trust boundaries between networks, where there would be no common view on their relative importance and no assurance that applications are marking appropriately. As a result, the DiffServ architecture is often used within managed networks (corporate networks, campus networks, etc.) but is not used on the Internet.

LLD's usage of the DiffServ Field to indicate NQB sidesteps this fundamental problem by eliminating the subjective value judgement on the relative importance of applications. Instead, this usage of the DiffServ Field describes objectively verifiable behavior on the part of the application—that it will not build a queue. Therefore, networks can verify that the marking has been applied properly before a packet is allowed into the Low Latency Service Flow queue (see Section 2.4).

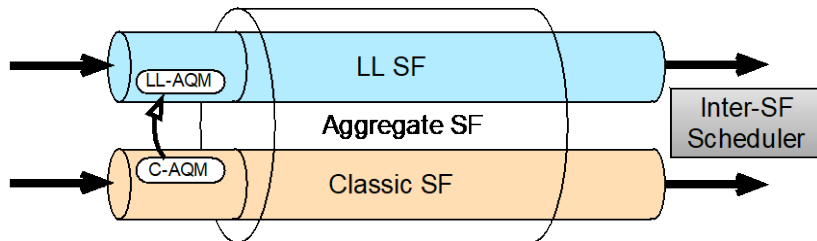
The ECN classifiers enable LLD's support of the IETF's Low-Latency Low-Loss Scalable throughput (L4S) service, which is an evolution of the original ECN facility, to support applications needing both high bandwidth and low latency (see Appendix A - Low Latency and High Bandwidth: L4S).

<sup>1</sup> As of the writing of this report, it is proposed that the DiffServ value 0x2A be standardized in IETF/IANA to indicate NQB [nqb-dscp]. Certain existing DiffServ values may also be classified as NQB by default, such as Expedited Forwarding (EF).

### 2.3. Coupled AQM

To manage queuing delay, both the Low Latency Service Flow queue and the Classic Service Flow queue support Active Queue Management (AQM) (Figure 4).

In the case of the Classic Service Flow, the queue implements the same state-of-the-art Active Queue Management techniques used in today's DOCSIS 3.1 networks. For upstream Classic Service Flows, the DOCSIS 3.1 specification mandates that the CM implement the DOCSIS-PIE (Proportional-Integral-Enhanced AQM Algorithm) [RFC8034], which introduces packet drops at an appropriate rate to drive the queue delay to the default target value of 10 ms. For downstream Classic Service Flows, the AQM in the CMTS is still vendor specific.




**Figure 4. Coupled AQM across Service Flows**

In the case of the Low Latency Service Flow, the queue supports L4S congestion controllers by implementing an Immediate Active Queue Management algorithm that utilizes ECN marking instead of packet drops. By default, the algorithm does not mark the packet if the queuing delay is less than 0.475 ms and always marks the packet if the delay is greater than 1 ms. Between those configurable values, the algorithm marks at a rate that ramps up from 0% to 100% over the range. In addition, per [aqm-dualq-coupled], the Immediate AQM in the Low Latency Queue is coupled to the Classic Queue AQM so that congestion in the Classic Queue will induce ECN marking in the Low Latency Queue in such a manner as to balance the per-flow throughput across all of the flows in both queues. L4S congestion control and the role of the dual-queue-coupled-aqm in providing this flow balance is described further in Appendix A - Low Latency *and* High Bandwidth: L4S.

To enable the Low Latency Queue to rapidly dequeue an arrived burst of traffic, the Inter-Service-Flow scheduler gives a higher weight to the Low Latency Queue than it does to the Classic Queue. The coupling to the Low Latency AQM counterbalances the weighted scheduler by making low-latency applications leave equal capacity for Classic applications. This ensures that the weighted scheduler does not give priority over bandwidth, as a traditional weighted scheduler would.

### 2.4. Queue Protection

Because of the small buffer size of the Low Latency Queue, classic TCP flows or other queue-building flows would see poor performance (due to high packet loss) if they were to end up in the Low Latency Queue. In addition, they would destroy the latency performance for the non-queue-building flows, negating the primary benefits of LLD.

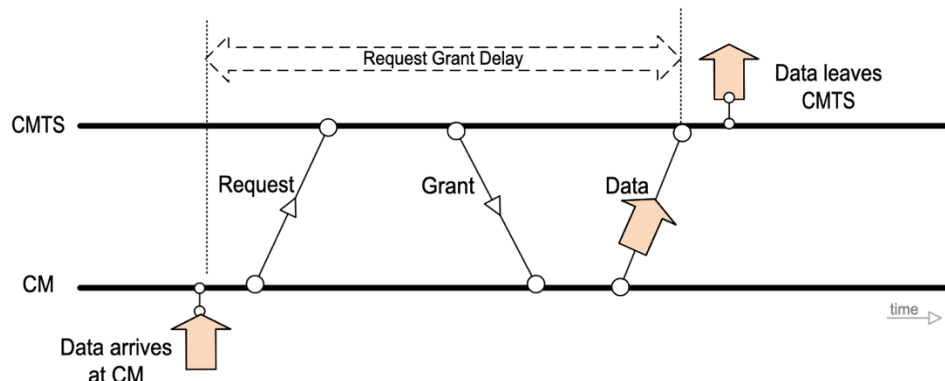


**Figure 5. Queue Protection function**

To prevent this situation, the packets that are classified to the Low Latency queue pass through a “Queue Protection” function (Figure 5), which scores each flow’s contribution to the growth of the queue. If the queue delay exceeds a threshold, the Queue Protection function identifies the flow or flows that have contributed most to the growth of the queue delay, and it redirects future packets from those flows to the Classic Service Flow. This mechanism is performed objectively and statistically, without examining the identifiers or contents of the data being transmitted. It is described more fully in [docsis-q-protection].

### 3. Upstream Scheduling Improvements

The DOCSIS upstream Media Access Control (MAC) Layer uses a request-grant mechanism. When data packets arrive at the CM to be transmitted in the upstream direction, a request message is sent from the CM to the CMTS. The CMTS then schedules the individual transmission bursts for all the CMs and communicates this via a bandwidth allocation map (MAP) message. Each MAP message describes the upstream transmission opportunities (grants) for a time interval and is sent shortly before the interval to which it applies.



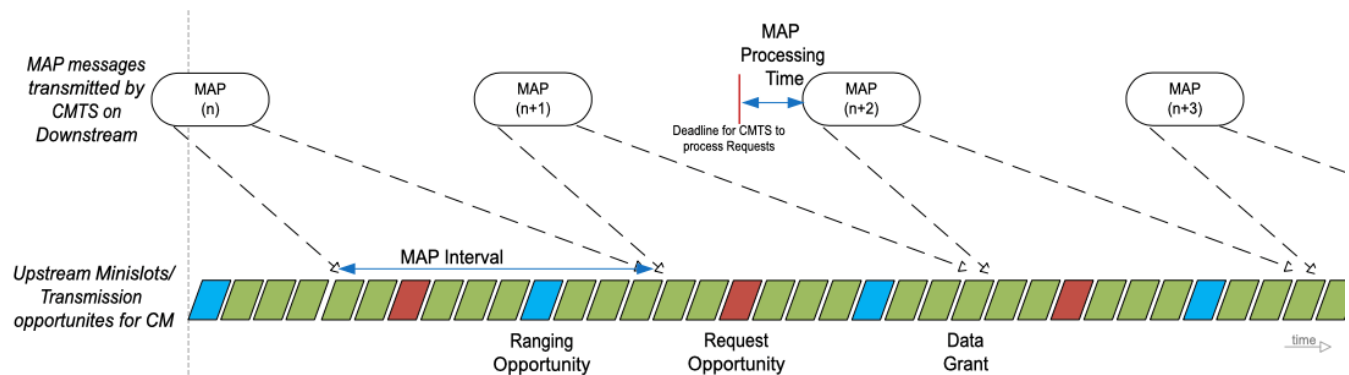
**Figure 6. Request Grant Delay in DOCSIS Networks**

When a CM has data to send, it waits for a "contention request" transmission opportunity. When that opportunity arrives, it sends a short request message indicating the amount of data it has to send. It then waits for a subsequent MAP message granting it a transmission opportunity in which to send its data. This time interval between the arrival of the packet at the CM and the time at which the data arrives at the CMTS on the upstream channel is known as the Request-Grant Delay (Figure 6). On uncongested channels, and in the absence of queuing delay, this delay is generally 2–8 ms.

### 3.1. Faster Request Grant Loop

LLD lowers the request-grant delay by requiring support for a shorter MAP Interval and a shorter MAP Processing Time (Figure 7).

The MAP interval is the amount of time that each MAP message describes. The MAP interval is also the time interval between consecutive MAP messages. Reducing the MAP interval means that the CMTS processes incoming requests more frequently, thus shortening the amount of time that a request might wait at the CMTS before being processed. A shorter MAP interval also means that grants are not scheduled as far into the future within each MAP message.



**Figure 7. MAP Interval and MAP Processing Time**

The MAP Processing Time is the amount of time the CMTS uses to perform its scheduling calculations. With a shorter MAP Processing Time, there is less delay between a request being received at the CMTS and the resulting grant being scheduled.

The LLD specification requires support for a nominal MAP interval of 1 ms or less for OFDMA upstream channels, in place of the 2-4 ms used previously. In certain configurations, a 1 ms MAP interval may introduce tradeoffs such as upstream and/or downstream inefficiency that will need to be weighed against the latency improvement.

### 3.2. Proactive Grant Service


DOCSIS scheduling services are designed to customize the behavior of the request-grant process for particular traffic types. LLD introduces a new scheduling service called Proactive Grant Service (PGS), which can eliminate the request-grant loop entirely (Figure 8).

In PGS, a CMTS proactively schedules a stream of grants to a Service Flow at a rate that is intended to match or exceed the instantaneous demand. In doing so, the vast majority of packets carried by the Service Flow can be transmitted without being delayed by the Request-Grant process. During periods when the CMTS estimates no demand for bandwidth for a particular PGS Service Flow, it can conserve bandwidth by providing periodic unicast request opportunities rather than a stream of grants.

The service parameters that are specific to PGS are Guaranteed Grant Interval (GGI), Guaranteed Grant Rate (GGR), and Guaranteed Request Interval (GRI). In addition, the traditional rate-shaping parameters,

such as Maximum Sustained Traffic Rate and Peak Rate, serve as an upper bound on the grants that can be provided to a PGS Service Flow.

PGS can eliminate the delay caused by the Request-Grant loop, but it comes at the price of efficiency. Inevitably, the CMTS will not be able to exactly predict the instantaneous demand for the Service Flow, so it may overestimate the capacity needed. When the shared channel is fully utilized, this could reduce the capacity available to other Service Flows.



**Figure 8. Proactive Grant Service in LLD**

The PGS scheduling type may appear at first to be similar to an existing DOCSIS upstream scheduling type “UGS/AD.” The main differences with PGS are that it sets a minimum floor on the level of granting (minimum grant spacing and minimum granted bandwidth) rather than setting a fixed grant pattern (fixed grant size and precise grant spacing), it supports the “Continuous Concatenation and Fragmentation” method of filling grants (where a contiguous sequence of bytes are dequeued to fill the grant, regardless of packet boundaries) rather than only carrying a single packet in each grant, and the CM is expected to continue to send Requests to the CMTS to inform it of packets that might be waiting in the queue.

## 4. Low Latency DOCSIS Performance

CableLabs has developed a simulator using the NS3 platform (<https://www.nsnam.org>) in order to evaluate the performance of different aspects of LLD. The simulator models a DOCSIS 3.1 link (OFDM/A channel types) between the CM and the CMTS and can be configured to enable or disable various components of the technology.

### 4.1. Traffic Models

Because the latency performance of the service depends on the mix of applications in use by the customer, we have developed a set of 10 traffic mix scenarios that represent what we believe to be common busy-hour behaviors for a cable customer. All traffic mixes include two bidirectional UDP sessions that are modeled after online games, but they could also represent VoIP or video conferencing/chatting applications. One of the sessions has its packets marked as NQB and the other does not, allowing us to see the benefit that the low-latency queue provides.

In addition, each traffic mix has a set of other applications that create background load, as summarized in Table 2 (see Appendix B - Simulation Details for details on the traffic types). All of this background load traffic utilizes the classic queue.

Some of these traffic mixes represent behaviors that may be very common for broadband users during busy hour, whereas others represent more extreme behaviors that users may occasionally engage in. When generating an overall view of the performance across all of the traffic mixes, we model the fact that they may not all be equally likely to occur by giving the more common mixes (1, 2, and 8) ten times the weight that we give to each of the other less common mixes.

**Table 2. Background Traffic Mixes**

Traffic Mix 1	1 web user
Traffic Mix 2	1 web user, 1 video streaming user
Traffic Mix 3	1 web user, 1 FTP upstream
Traffic Mix 4	1 web user, 1 FTP downstream
Traffic Mix 5	1 web user, 1 FTP upstream and 1 FTP downstream
Traffic Mix 6	1 web user, 5 FTP upstream and 5 FTP downstream
Traffic Mix 7	1 web user, 5 FTP up, 5 FTP down, and 2 video streaming users
Traffic Mix 8	5 web users
Traffic Mix 9	16 TCP down (speedtest)
Traffic Mix 10	8 TCP up (speedtest)

## 4.2. Round Trip P99 Latency based on Traffic Mix

Table 3 summarizes the 99<sup>th</sup> percentile per-packet latency for the NQB-marked game traffic across all ten traffic mixes, as well as the weighted overall performance, for four different systems:

1. a legacy DOCSIS 3.1 system with AQM disabled, 2 ms MAP interval;
2. a legacy DOCSIS 3.1 system with AQM enabled, 2 ms MAP interval;
3. a Low Latency DOCSIS 3.1 system without PGS, 1 ms MAP interval; and
4. a Low Latency DOCSIS 3.1 system with PGS configured for 5 Mbps GGR, 1 ms MAP interval.

We include LLD with and without PGS because some network operators may wish to deploy LLD without the overhead that comes with PGS scheduling.

**Table 3. 99<sup>th</sup> Percentile round-trip latency for nqb-marked traffic between the CM and CMTS**

	Legacy DOCSIS 3.1 with no AQM	Legacy DOCSIS 3.1 with AQM	Low Latency DOCSIS with no PGS	Low Latency DOCSIS with PGS
Traffic Mix 1	7.7 ms	7.7 ms	4.7 ms	0.9 ms
Traffic Mix 2	7.7 ms	7.7 ms	4.8 ms	0.9 ms
Traffic Mix 3	159.5 ms	36.6 ms	4.7 ms	0.9 ms

	Legacy DOCSIS 3.1 with no AQM	Legacy DOCSIS 3.1 with AQM	Low Latency DOCSIS with no PGS	Low Latency DOCSIS with PGS
Traffic Mix 4	7.8 ms	7.9 ms	4.7 ms	0.9 ms
Traffic Mix 5	159.6 ms	57.4 ms	4.7 ms	0.9 ms
Traffic Mix 6	253.7 ms	96.7 ms	4.7 ms	0.9 ms
Traffic Mix 7	253.9 ms	74.7 ms	4.7 ms	0.9 ms
Traffic Mix 8	7.7 ms	7.7 ms	4.7 ms	0.9 ms
Traffic Mix 9	259.3 ms	52.1 ms	4.8 ms	0.9 ms
Traffic Mix 10	254.0 ms	34.1 ms	4.8 ms	0.9 ms
<b>Weighted Overall P99</b>	<b>250.5 ms</b>	<b>32.4 ms</b>	<b>4.7 ms</b>	<b>0.9 ms</b>

As can be seen in this table, there are several traffic mixes (notably 1, 2, 4, and 8) for which the relatively light traffic load doesn't create the conditions for TCP to cause significant queuing delay at P99, so even the "Legacy DOCSIS 3.1 with no AQM" system results in fairly low latency. However, in the heavier traffic mixes, the benefit of AQM can be seen and the benefit of the dual-queue mechanism in LLD becomes very apparent. By separating the NQB-marked traffic from the queue-building traffic, the NQB-marked traffic is isolated from the delay created by the TCP flows entirely, and very reliable low latency is achieved. The system modeled by the right-most column, which additionally implements PGS, can eliminate the request-grant delay for the NQB traffic and thereby drive the round-trip latency below 1 ms at 99<sup>th</sup> percentile.


### 4.3. Round Trip Packet Delay and Packet Delay Variation Statistics

Figure 9 illustrates the weighted overall latency performance across all ten traffic mixes. The plot is a log-log complementary cumulative distribution function, with the y-axis labeled with the equivalent quantile values. Focusing, for instance, on the horizontal line indicating the 99<sup>th</sup> percentile (P99), it can be seen that LLD with PGS holds delay below 0.9 ms for 99% of packets. In contrast, a DOCSIS 3.1 network without AQM can only hold delay below 250 ms for 99% of packets. So, P99 delay is more than 250 times better with LLD.

Additionally, this plot illustrates the reduction in Packet Delay Variation (sometimes referred to as "jitter") provided by LLD. Packet Delay Variation (PDV) is defined in [RFC5481] as the variation in packet delay relative to the minimum packet delay experienced (i.e. relative to the P0 value in the plot). Thus we can see P99.9 round-trip PDVs of: 400  $\mu$ s for LLD w/PGS, 3 ms for LLD w/o PGS, 75 ms for D3.1 w/AQM, and 400 ms for D3.1 w/o AQM. This is an improvement by a factor of 1000.

We therefore see that LLD will bring a consistent, low-latency, responsive quality to cable broadband performance and user experiences for NBQ traffic.






**Figure 9. Overall Round-Trip Latency Performance for NQB-Marked Traffic**

#### 4.4. Proactive Grant Service Tradeoffs


Figure 10 illustrates the impact that the PGS “Guaranteed Grant Rate” (GGR) parameter has on the 99<sup>th</sup> percentile latency for the NQB traffic. The PGS implementation used for this experiment is an extremely basic one, it simply provides “unsolicited” grants at intervals of no more than 1 ms and with a grant size that is based on the GGR parameter. It also responds to direct requests by the CM for additional grants. The GGR value thus sets the amount of upstream capacity that is utilized by the low latency service flow, even when the actual user traffic rate may be very low. For example in these simulation cases, the “game” flow is sending about 30 kbps, yet with a GGR set to (e.g.) 5 Mbps, the amount of DOCSIS upstream channel capacity consumed by the customer’s service flow is 5 Mbps. In the simplified PGS implementation used here, the CMTS scheduler is making no attempt to track utilization and adjust its grant rate accordingly. It is expected that CMTS vendors will implement more sophisticated algorithms that can achieve better latency performance with less bandwidth overhead, so in a sense, this figure provides an idea as to the worst-case performance that could be expected with PGS.



**Figure 10. Impact of PGS on NQB Traffic Latency**

#### 4.5. Multiple Simultaneous Game Sessions

Figure 11 shows the latency performance result when there is more than one game session active simultaneously. The plot shows the 99<sup>th</sup> percentile round trip latency over the DOCSIS link, for anywhere from 1 to 100 simultaneous sessions. In this case, PGS was disabled.




**Figure 11. Impact of Multiple Simultaneous Game Sessions**

#### 4.6. Uni-directional LLD

The LLD dual-queue functionality can be enabled in one direction and not the other if the operator so chooses. Figure 12 illustrates the impact of having dual-queue in both directions (“usds”), upstream only

(“usonly”), downstream only (“dsonly”) or in neither direction (“aqm”). In this experiment, when dual-queue is disabled, the traditional DOCSIS-PIE AQM is utilized. As can be seen, dual-queue provides an independent benefit in each direction, particularly when one looks at the 90<sup>th</sup> percentile of performance. But, in order to achieve reliable low-latency service (e.g. 95<sup>th</sup> percentile and above), the dual-queue functionality is needed in both directions.



**Figure 12. Impact of Uni-directional LLD**

## 5. Deployment Considerations

### 5.1. Device Support

Deploying LLD in the MSO network can be accomplished via software-only upgrades to the existing DOCSIS 3.1 CMs and CMTSs. Table 4 shows which LLD features need implementation on the CM side, the CMTS side, or both. The Dual Queue feature in the upstream requires an upgrade to the CM as well as to the CMTS. The other features (Dual Queue in Downstream, Upstream Scheduling improvements) only require upgrades on the CMTS, so they can be deployed even to CMs that don’t support LLD (including DOCSIS 3.0 modems). As shown in the previous section, offering a Downstream-only Dual Queue for customers with older modems has some merit.

**Table 4. Device Dependencies for LLD Features**

LLD Feature	Downstream Latency Improvements		Upstream Latency Improvements	
	CMTS upgrade?	CM upgrade?	CMTS upgrade?	CM upgrade?
Dual Queue (ASF, Coupled AQM, QP)	Required	Not required	Required	Required
Upstream Scheduling (Faster Req-Grant Loop, PGS)	Not applicable		Required	Not required

## 5.2. Packet Marking

The design of LLD takes the approach that applications are in the best position to determine which flows or which packets are non-queue-building. Thus, applications such as online games will be able to tag their packets with the NQB DiffServ value to indicate that they behave in a non-queue-building way, so that LLD will be able to classify them into the Low Latency Service Flow.

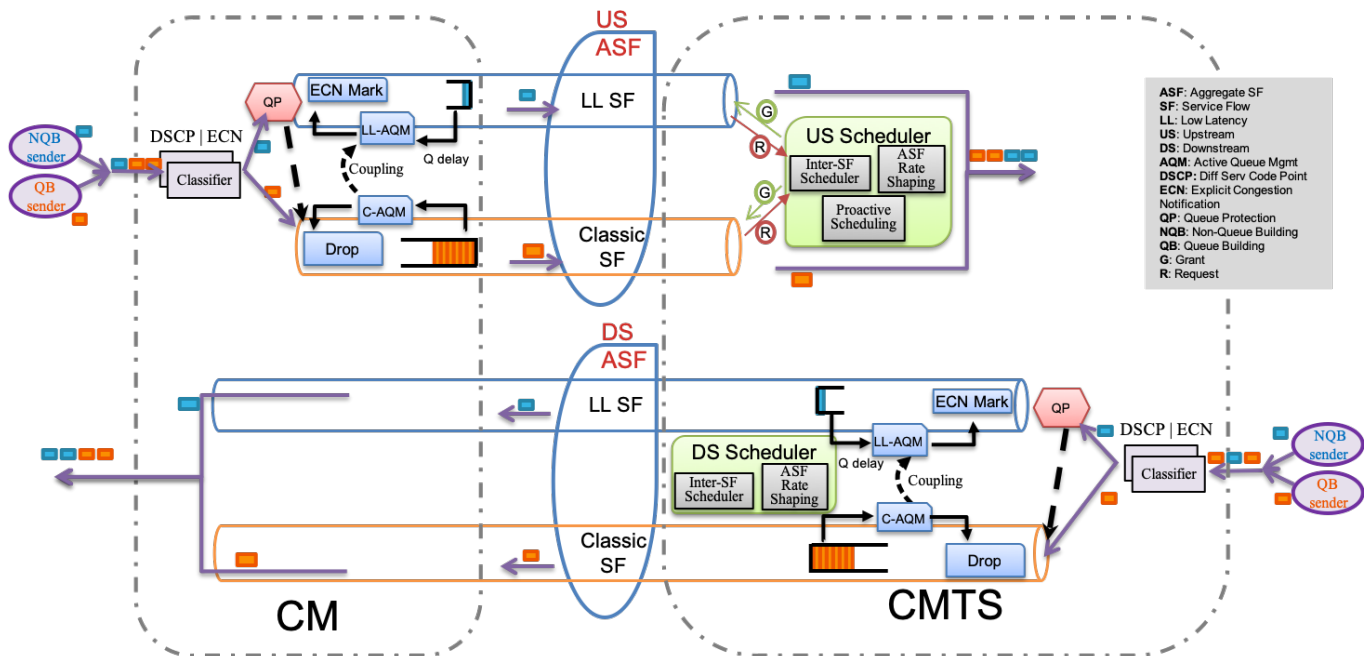
For these packet markings to be useful for the LLD classifiers, they will need to survive the journey from the application source to the CM or CMTS. In some cases, operators today clear the DiffServ Field in packets entering their network from an interconnecting network, which would prevent the markings making their way to the CMTS. This practice is presumably driven by the view that DiffServ Field usage is defined by each operator for use within its network, in which case preserving another network's markings has no value. As was described in Section 2.2, it is proposed that a single globally standard value be chosen to indicate NQB so that operators that intend to support LLD can ensure that this specific value traverses their inbound interconnects and their network and then arrives at the CMTS intact.

Although application marking is preferable, some network operators might want to provide immediate benefits to applications that behave in a non-queue-building way, in advance of application developers introducing support for NQB tagging. It might be possible to repurpose the queue protection function to identify NQB behavior even if the packets are not tagged as NQB, e.g., by assuming that all non-TCP traffic is likely to be NQB and relying on queue protection to redirect the QB flows. The effectiveness of this approach is currently an area of research.

Further, it is possible that intermediary software or devices (either installed by the user or provided by the operator) could identify flows that are expected to be NQB and mark the packets on behalf of the application. This approach is actively being pursued.

## 5.3. Provisioning Mechanisms

The LLD specifications include provisioning mechanisms to allow an MSO to deploy low-latency features with minimal operational impact. Figure 13 shows all the pieces needed to build a low-latency service in the upstream and downstream direction. Although it is possible to define a Low Latency ASF, its constituent Classic and Low Latency SFs, and the associated classifiers explicitly in the CM's configuration file, a new feature known as the Aggregate QoS Profile can make this configuration automatic in many cases. Default classifiers will be created and default parameters for AQM and queue protection will be used, or any of these can be overridden by the operator as needed.



**Figure 13. ASFs, SFs, and the Low Latency DOCIS Architecture**

### 5.3.1. Aggregate QoS Profiles

Similar to Service Class Names that are expanded by the CMTS into a set of QoS parameters for a Service Flow during the registration process, an operator can create an Aggregate QoS Profile (AQP) on the CMTS to describe the parameters of an Aggregate Service Flow, its constituent Service Flows, and the classifiers used to identify NQB traffic.

Just like with Service Class Names, the operator can also provide explicit values in the configuration file for any ASF or SF parameters that they wish to “override.”

### 5.3.2. Migration Using Existing Configuration File and Service Class Name

One very straightforward way to migrate to LLD configurations may not involve any changes to the CM configuration file. This method involves the automatic expansion of a Service Flow definition to a Low Latency ASF via the use of a Service Class Name and matching AQP definition.

When the CMTS sees a Service Class Name in a Service Flow definition from the CM’s config file, if the CM indicates support for LLD, then the CMTS will first use the Service Class Name as an AQP Name and look for a matching entry in the AQP Table. If it finds a matching entry, it will automatically expand the Service Flow into an ASF and two Service Flows.

This mechanism allows the operator to deploy LLD by simply updating the CMTS to support the feature and configuring AQP Table entries that match the Service Class Names in use in CM config files. Then, as CMs are updated over time to include support for LLD, they will automatically start being configured with a Low Latency ASF.


### 5.3.3. Explicit Definition of ASF in the Configuration File

An operator can also encode a Low Latency ASF in a CM configuration file directly using an Aggregate Service Flow TLV (70 or 71). The ASF TLV could have an AQP Name that is used by the CMTS to look up a definition of the ASF in its AQP Table. It could also have ASF parameters that would explicitly define the ASF or would override the AQP parameters. A configuration could also have explicit individual Service Flow TLVs (24 or 25) that are linked to the ASF via the Aggregate Service Flow Reference TLV.

## 5.4. Latency Histogram Reporting

As part of the AQM and Queue Protection operations, CMs and CMTSs generate estimates of the queuing latency for the upstream and downstream Service Flows, respectively. The latency histogram reporting function exposes these estimates to the operator to provide information that can be utilized to characterize network performance, optimize configurations, or troubleshoot problems in the field.

This latency histogram reporting can be enabled via a configuration file setting or can be initiated by setting a MIB object on the device. The operator configures the bins of the histogram, and the CM or the CMTS logs the number of packets with recorded latencies into each of the bins. The CM implements histograms for upstream Service Flows, and the CMTS implements histograms for downstream Service Flows. (This function can be enabled even for Service Flows for which AQM is disabled.) The latency estimates from the AQM are represented in the form of a histogram as well as a maximum latency value.



**Figure 14. Example of Using a Latency Histogram Report to Estimate Latency Quantiles**

## Conclusion

LLD enables a huge leap in latency performance and will improve the Internet experience overall. With LLD technology, online gaming will become more responsive and video chats will cease to be “choppy.” This technology will enable a range of new applications that require real-time interface between the cyber and physical worlds, such as vehicular communications and remote health care services.

To realize the benefits of LLD, a number of parties need to take action. DOCSIS equipment manufacturers will need to develop and integrate the LLD features into software updates for CMTSs and CMs. Cable operators need to plan the roll-out of software updates and configurations to DOCSIS

equipment and set up the network to support those services (e.g., carrying DiffServ/ECN markings through the network). Application and operating system vendors will need to adopt packet marking for NQB traffic and/or adopt the L4S congestion controller. Each element of the Internet ecosystem will make these decisions independently; the faster that all take the necessary steps, the more quickly the user experience will improve.

The cable industry has provisioned its network with substantial bandwidth and is poised to take another leap forward with its 10G networks. But more bandwidth is only part of the broadband performance story. Latency is becoming crucial to the evolution of broadband. That is why LLD technology is a cornerstone of cable's 10G future.

## Abbreviations

AMSR	aggregate maximum sustained traffic rate
AQM	active queue management
AQP	aggregate QoS profile
ASF	aggregate service flow
BBR	bottleneck bandwidth and RTT
CCDF	complementary cumulative distribution function
CE	congestion experienced
CM	cable modem
CMTS	cable modem termination system
DASH	dynamic adaptive streaming over HTTP
DCTCP	data center TCP
DiffServ	differentiated services
DOCSIS	Data-Over-Cable Service Interface Specification
DOCSIS-PIE	DOCSIS Proportional Integral Enhanced AQM
ECN	explicit congestion notification
ECT	ECN capable transport
EF	expedited forwarding
FTP	file transfer protocol
Gbps	gigabits per second
GGI	guaranteed grant interval
GGR	guaranteed grant rate
GRI	guaranteed request interval
HFC	hybrid fiber-coax
HTTP	hypertext transfer protocol
IANA	Internet Assigned Numbers Authority
IETF	Internet Engineering Task Force
IPv4	internet protocol version 4
IPv6	internet protocol version 6

kbps	kilobits per second
L4S	low-latency low-loss scalable throughput
LLD	Low Latency DOCSIS
MAC	media access control
Mbps	megabits per second
MIB	management information base
ms	milliseconds
MSO	multiple system operator
NQB	non-queue-building
NS3	Network Simulator 3
OFDM	orthogonal frequency division multiplexing
OFDM/A	orthogonal frequency division multiplexing or multiple access
OFDMA	orthogonal frequency division multiple access
P99	99th percentile
PGS	proactive grant service
QB	queue-building
QoS	quality of service
RTT	round-trip time
SF	service flow
TCP	transmission control protocol
TLV	type length value
UGS/AD	unsolicited grant service with activity detection
VR	virtual reality

## Bibliography & References

[MULPIv3.1] MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I17-190121, January 21, 2019, Cable Television Laboratories, Inc.

[CCAP-OSSIV3.1] DOCSIS 3.1 CCAP Operations Support System Interface Specification, CM-SP-CCAP-OSSIV3.1-I14-190121, January 21, 2019, Cable Television Laboratories, Inc.

[CM-OSSIV3.1] DOCSIS 3.1 Cable Modem Operations Support System Interface Specification, CM-SP-CM-OSSIV3.1-I14-190121, January 21, 2019, Cable Television Laboratories, Inc.

[RFC3168] Ramakrishnan, K., “The Addition of Explicit Congestion Notification (ECN) to IP”, [RFC3168](#), DOI 10.17487/RFC3168, September 2001.

[RFC5481] Morton, A. and B. Claise, “Packet Delay Variation Applicability Statement”, [RFC5481](#), DOI 10.17487/RFC5481, March 2009.



[RFC8034] White, G. and R. Pan, “Active Queue Management Based on Proportional Integral Controller Enhanced (PIE) for Data-Over-Cable Service Interface Specifications (DOCSIS) Cable Modems”, [RFC8034](#), DOI 10.17487/RFC8034, February 2017.

[RFC8311] Black, D., "Relaxing Restrictions on Explicit Congestion Notification (ECN) Experimentation", [RFC8311](#), DOI 10.17487/RFC8311, January 2018.

[aqm-dualq-coupled] Schepper, K., B. Briscoe, and G. White, "DualQ Coupled AQMs for Low Latency, Low Loss and Scalable Throughput (L4S)", [draft-ietf-tsvwg-aqm-dualq-coupled](#) (work in progress), July 2019.

[docsis-q-protection] Briscoe, B. and G. White, “Queue Protection to Preserve Low Latency”, [draft-briscoe-docsis-q-protection](#) (work in progress), July 2019.

[ecn-l4s-id] Schepper, K. and B. Briscoe, "Identifying Modified Explicit Congestion Notification (ECN) Semantics for Ultra-Low Queuing Delay (L4S)", [draft-ietf-tsvwg-ecn-l4s-id](#) (work in progress), July 2019.

[nqb-dscp] White, G. and T. Fossati, “Identifying and Handling Non Queue Building Flows in a Bottleneck Link”, [draft-white-tsvwg-nqb](#) (work in progress), June 2019.

[web-user-model] 3GPP standardized web user model, 3GPP2-TSGC5, "HTTP, FTP and TCP models for 1xEV-DV simulations", 2001.

## Appendix A - Low Latency *and* High Bandwidth: L4S

How can LLD support applications that want maximum speed, and low latency too? This is achievable through a technology called L4S: Low Latency Low Loss Scalable throughput.


L4S improves many of today's applications (e.g., video chat, everything on the web), but it will also enable future applications that will need both high bandwidth and low delay, such as cloud gaming, HD video conferencing, cloud-rendered interactive video, cloud-rendered virtual reality, augmented reality, remote presence with remote control, interactive light field experiences, and others yet to be invented.

L4S involves incremental changes to the congestion controller on the sender and to the AQM at the bottleneck. The key is to indicate congestion by marking packets using Explicit Congestion Notification (ECN) rather than discarding packets. L4S uses the 2-bit ECN field in the IP header (v4 or v6) and defines each marked packet to represent a lower strength of congestion signal [RFC8311] than the original ECN standard [RFC3168]. All the benefits of L4S follow from that.

- **Low Latency:** The sender's L4S congestion controller makes small but frequent rate adjustments dependent on the proportion of ECN marked packets, and the L4S AQM starts applying ECN-marks to packets at a very shallow buffer threshold. This means an L4S queue can ripple at the very bottom of the buffer with sub-millisecond queuing delay but still fully utilize the link. Small, frequent adjustments could not even be considered if packet discards were used instead of ECN—they would induce a prohibitively high loss level. Further, AQMs could not consider a very shallow threshold if small adjustments were not used, as severe link under-utilization would result.

- **Low Loss:** By definition, using ECN eliminates packet discard. In turn, that eliminates retransmission delays, which particularly impact the responsiveness of short web-like exchanges of data. Using ECN eliminates both the round-trip delay repairing a loss and the delay while detecting a loss. In addition, an L4S AQM can immediately signal queue growth using ECN, catching queue growth early. In contrast, classic AQMs hold back from discarding a packet for 100–200 ms because if a burst subsides of its own accord, a loss in itself could cause more harm than the good it would do as a signal to slow down. Furthermore, eliminating packet discard eliminates the collateral damage caused to flows that were not significantly contributing to congestion.
- **Scalable Throughput:** Existing congestion control algorithms don't scale, so applications need to open many simultaneous connections to fully utilize today's broadband connections. An L4S congestion controller can rapidly ramp up its sending rate to match any link capacity. This is because L4S uses a "scalable congestion controller" that maintains the same frequency of control signals (2 ECN marks per round trip on average) regardless of flow rate. With classic congestion controllers, the faster they try to go, the longer they run blind without any control signals.

The technology behind L4S isn't new; it is based on a scalable congestion control called Data Center TCP (DCTCP) that is currently used in data centers to get very high throughputs with ultra-low delay and loss. What is new is the development of a way that scalable traffic can coexist with the existing TCP and QUIC traffic on the Internet—the key that unlocks a transition to L4S. Until now, DCTCP has been confined to data centers because it would starve any classic flows sharing a link.



**Figure 12. Coexistence of L4S and Classic Traffic Using Coupled AQMs**

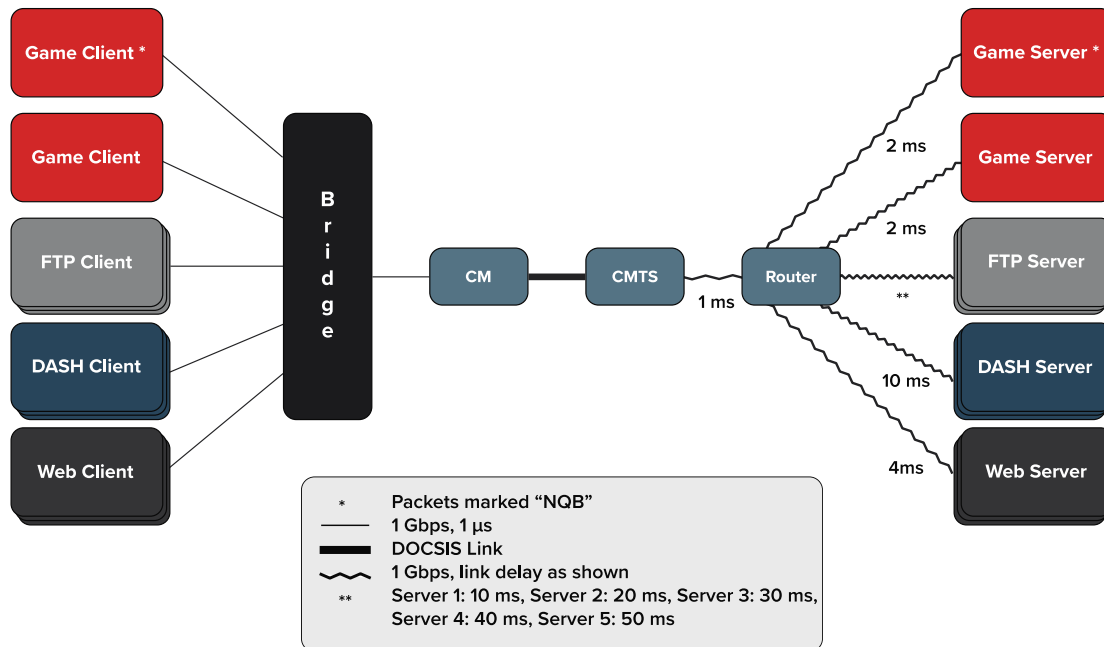
Separation into two queues serves two purposes: (1) it isolates L4S flows from the queuing of classic TCP and QUIC and (2) it sends each type of traffic appropriately scaled congestion signals. This results in any number of application flows (of either type) all getting roughly equal bandwidth each, as if there were just one aggregate pool of bandwidth, with no division between the Service Flows.

The approach couples the levels of ECN and drop signaling, as shown in Figure 15. The packet rate of today's classic congestion controls conforms to the well-known square-root rule ( $r_c$  on the left of the figure), whereas the packet rate of L4S congestion control conforms to a linear rule ( $r_l$ ). So, the Low Latency AQM applies an ECN marking level to L4S traffic that is coupled to the square root of the dropping level being applied to Classic traffic. This ensures that the packet rates of the two types of flow turn out roughly the same.

Supporting L4S in LLD equipment is relatively straightforward. All that is needed is to classify L4S flows into the Low Latency SF and support the logic in the Low Latency SF to perform immediate ECN marking of packets (see Section 2.2).

## Appendix B - Simulation Details

For the results reported in this paper, we set up the following network with 5 types of client devices behind the CM and a set of servers behind the CMTS. The link delays shown are 1-way values. The DOCSIS link is configured in the most latency-efficient manner (short interleavers, small OFDMA frame sizes) and models a plant distance of 8 km. The service is configured with a Maximum Sustained Traffic Rate (rate limit) of 50 Mbps in the upstream direction and 200 Mbps in the downstream direction.



**Figure 15. NS3 Network Topology**

The upstream game traffic model involves normally distributed packet interarrival times ( $\mu=33$  ms,  $\sigma=3$  ms) and normally distributed packet sizes ( $\mu=110$  bytes,  $\sigma=20$  bytes) constrained to discard draws of packet size  $<32$  bytes or  $>188$  bytes. The downstream game traffic model involves normally distributed packet interarrival times ( $\mu=33$  ms,  $\sigma=5$  ms) and normally distributed packet sizes ( $\mu=432$  bytes,  $\sigma=20$  bytes) constrained to discard draws of packet size  $<32$  bytes or  $>832$  bytes.

The background load traffic is configured as follows. The web user is based on the 3GPP standardized web user model [web-user-model]. The video streaming model is an abstracted model of a Dynamic Adaptive Streaming over HTTP (DASH) streaming video user where the video stream is 6 Mbps and is implemented as a 3.75 MB file download every 5 seconds. Each FTP session involves the sender selecting a file size using a log-normal random variable ( $\mu=14.8$ ,  $\sigma=2.0$ , leading to a median file size of 2.7 MB), opening a TCP connection, sending the file, closing the TCP connection, then pausing for 100 ms before repeating the process. Although we refer to this model as an FTP model, the intention is that it models TCP usage across all applications other than web browsing, speed test, and video streaming. The speed-test model utilizes the FTP client and server nodes, but launches 16 simultaneous long-running (unbounded size) TCP connections for the downstream test, or 8 simultaneous long-running TCP connections for the upstream test.

# **Experiment Results for Supporting LTE-FDD, LTE-TDD, and 5G Timing Synchronization Over DOCSIS CAA and DAA**

A Technical Paper prepared for SCTE•ISBE by

**Yair Neugeboren**

Director, Systems Engineering  
CommScope  
32 Hamelacha Street, Netanya, Israel  
+972 542 205 051  
yair.neugeboren@commscope.com

**Greg Cyr**

Engineering Fellow  
CommScope  
2400 Ogden Ave., Suite 180, Lisle, IL 60532  
+1 630 281 3031  
greg.cyr@commscope.com

**Chris Zettinger**

Principal Systems Engineer  
CommScope  
2400 Ogden Ave., Suite 180, Lisle, IL 60532  
+1 630 281 3272  
chris.zettinger@commscope.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	4
Content .....	4
1. MBH over DOCSIS Use Cases and Timing Requirements.....	4
1.1. LTE-FDD and APTS Support with an IEEE 1588-Unaware DOCSIS Network .....	4
1.1.1. Synchronization Requirements.....	4
1.1.2. Proposed Solution over DOCSIS.....	7
1.2. LTE-TDD and 5G Support with an IEEE 1588-Aware DOCSIS Network.....	7
1.2.1. Synchronization Requirements.....	8
1.2.2. Proposed Solution over DOCSIS.....	10
2. Equipment and Lab Setup.....	11
2.1. LTE-FDD and APTS Support with an IEEE 1588-Unaware DOCSIS Network .....	12
2.1.1. I-CCAP .....	12
2.1.2. Remote PHY.....	13
2.2. LTE-TDD and 5G Support with an IEEE 1588-Aware DOCSIS Network.....	14
2.2.1. I-CCAP .....	14
2.2.2. Remote PHY.....	15
3. Procedure and Results.....	17
3.1. LTE-FDD and APTS Support with an IEEE 1588-Unaware DOCSIS Network .....	17
3.1.1. I-CCAP .....	17
3.1.2. Remote PHY.....	19
3.2. LTE-TDD and 5G with an IEEE 1588-Aware DOCSIS Network .....	20
3.2.1. Remote PHY.....	20
Conclusion .....	26
Abbreviations.....	27
Bibliography and References.....	28

## List of Figures

Title	Page Number
Figure 1 – ITU-T G.8261.1 Reference Model for LTE-FDD.....	5
Figure 2 – ITU-T G.8261.1 Network Limits for LTE-FDD .....	6
Figure 3 – ITU-T G.8271.2 Reference Model for APTS .....	6
Figure 4 – Proposed Solution for LTE-FDD and APTS Support over DOCSIS.....	7
Figure 5 – End-to-End View of Full Timing Support for Phase Sync .....	8
Figure 6 – End-to-End View of Partial Timing Support for Phase Sync .....	8
Figure 7 – ITU-T G.8271.1 Dynamic Time Error Network Limit (MTIE) .....	10
Figure 8 – DTP Reference Architecture.....	10
Figure 9 – PTP Frequency Delivery (G.8265.1) with I-CCAP .....	12
Figure 10 – PTP Frequency Delivery (G.8265.1) with R-PHY .....	13
Figure 11 – PTP Time and Frequency Delivery (G.8275.2) with I-CCAP.....	15
Figure 12 – PTP Time and Frequency Delivery (G.8275.2) with R-PHY.....	16
Figure 13 – PDV for DOCSIS 3.0 with I-CCAP .....	17
Figure 14 – G.8261.1 MTIE Results for DOCSIS 3.0 I-CCAP .....	18
Figure 15 – PDV for DOCSIS 3.1 with I-CCAP .....	18
Figure 16 – G.8261.1 MTIE Results for DOCSIS 3.1 I-CCAP .....	19
Figure 17 – PDV for D3.1 with R-PHY.....	19
Figure 18 – PTP Time and Frequency Delivery (G.8275.2) with R-PHY.....	21
Figure 19 – Phase Transfer Stability Between RPD and CM 1PPS .....	21
Figure 20 – Phase Transfer Stability Between RPD and CM 1PPS (After CM Reset) .....	22
Figure 21 – Phase Transfer Stability Between R-PHY and CM 1PPS (After both CM and RPD Reset) ...	22
Figure 22 – TE of the Recovered Phase at the Slave Probe with 3' of Coax .....	23
Figure 23 – TE of the Recovered Phase at the Slave Probe with DTP Compensation .....	24
Figure 24 – TE of the Recovered Phase at the Slave Probe with 400' of Coax .....	24
Figure 25 – G.8271.1 MTIE Measurements Calculated by the Slave Probe with 400' of Coax .....	25
Figure 26 – PTP Probe (Probe to CM) .....	25
Figure 27 – PTP Probe (CM to Probe) .....	26

## List of Tables

Title	Page Number
Table 1 – 4G/LTE and 5G Synchronization Requirements.....	8
Table 2 – ITU-T G.8271.1 Example of Time Error Allocation .....	9
Table 3 – Noise Generation Estimation for a Pair of Media Converters.....	11
Table 4 – Summary of Phase Transfer Stability Measurements.....	23

# Introduction

The ability to support the transport of accurate timing information over DOCSIS networks has become an important and urgent need for both cable and telco operators as mobile backhaul becomes a key service offering of the future.

The DOCSIS network (without any modifications) is not an appropriate access infrastructure for delivering accurate timing information for both frequency and phase. The latency asymmetry between the DOCSIS upstream (US) and downstream (DS) and the large packet delay variation (PDV) in upstream grant cycles cannot guarantee the timing requirements of LTE; 16 ppb of frequency accuracy for LTE-FDD and 1.5 microsecond of phase accuracy for LTE-TDD and 5G.

There is active research on-going within the cable industry to define the requirements needed for the DOCSIS network to support these stringent mobile backhaul requirements. Much of this work is based on the DOCSIS Timing Protocol (DTP), which was previously defined within DOCSIS 3.1 and the CableLabs MBH (Mobile Backhaul) Synchronization Techniques specification.

This paper will describe modeling and experimental results for accurate phase and frequency delivery over DOCSIS, covering multiple different use cases as LTE-FDD, LTE-TDD, and 5G. Some approaches require minimal changes to the current DOCSIS implementations, while others require incorporating the DOCSIS 3.1 DOCSIS Timing Protocol (DTP). This paper will show how the experiments results prove that DOCSIS Distributed Access Architecture (DAA) and Centralized Access Architecture (CAA) could support the MBH requirements for the various use cases.

In this paper we will focus on two use cases, one for frequency delivery and one for phase delivery. We will describe the mobile sync accuracy requirements in general from 3GPP and ITU-T and the proposed solutions for those use cases over DOCSIS.

Following that, we will describe in detail the MBH over DOCSIS setups we used to explore the performance and show the results for each use case.

## Content

### 1. MBH over DOCSIS Use Cases and Timing Requirements

#### 1.1. LTE-FDD and APTS Support with an IEEE 1588-Unaware DOCSIS Network

##### 1.1.1. *Synchronization Requirements*

Frequency synchronization is needed in two main use cases. One is for LTE-FDD support and the other is aimed to provide frequency assistance for Assisted Partial Timing Support (APTS) clocks.

LTE-FDD mobile cells use two frequencies for transmitting and receiving data simultaneously, so it requires the frequency to be very accurate so there will not be any interference between the transmissions.

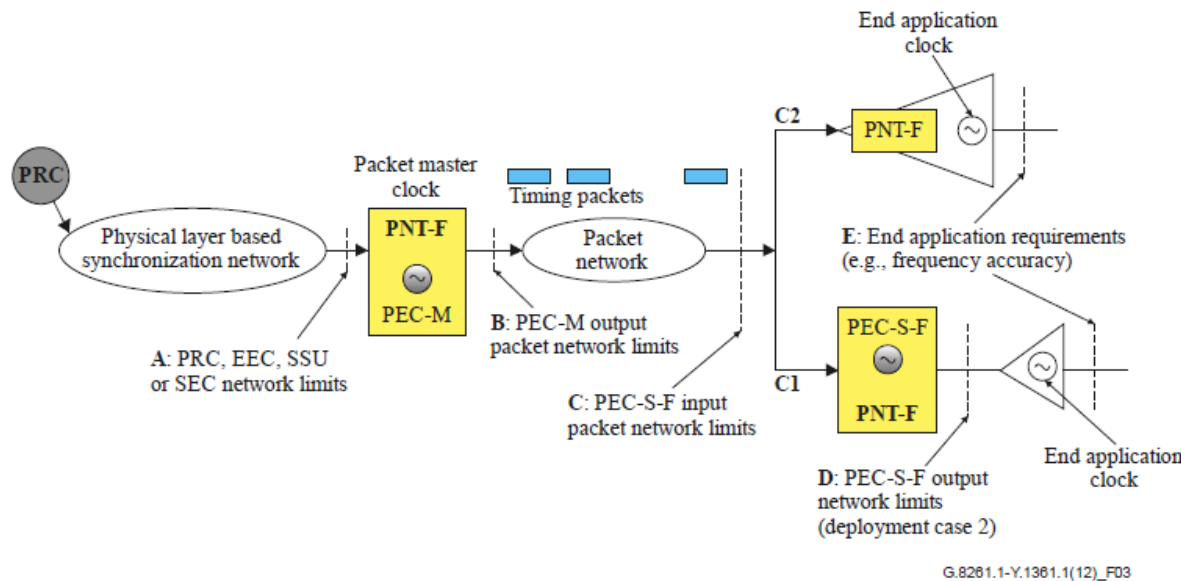
Frequency synchronization over the network can be achieved mainly in one of two methods:

- A. Physical layer frequency synchronization, for example Synchronous Ethernet (SyncE). This method requires each network element in the path to support this physical layer synchronization.
- B. Using Precision Timing Protocol (PTP) to achieve frequency synchronization with or without phase synchronization. ITU-T G.8261.1 defines the packet delay variation network limits applicable to packet-based methods for frequency synchronization. ITU-T G.8265.1 is the precision time protocol telecom profile used for frequency synchronization.

**This paper will focus on the PTP method for achieving frequency synchronization for LTE-FDD.**

The LTE-FDD frequency accuracy defined by 3GPP at the air interface is  $\pm 50$  ppb. The frequency accuracy requirement between the frequency source and the end application is  $\pm 16$  ppb.

As mentioned above, ITU-T G.8261.1 defines the packet delay variation network limits applicable to packet-based methods for frequency synchronization to achieve the 3GPP required frequency accuracy. The reference model for the network limits is shown in Figure 1.



**Figure 1 – ITU-T G.8261.1 Reference Model for LTE-FDD**


Reference point C is the point where the LTE-FDD cell is connected (and where the DOCSIS network ends). The packet delay variation network limit given in the ITU-T G.8261.1 is defined as follows:

With window interval  $W = 200$  seconds and fixed cluster range  $\delta = 150 \mu\text{s}$  starting at the floor delay, the network transfer characteristic should satisfy a Floor Packet Percentage (FPP)  $(n, W, \delta) \geq 1\%$ . This means that for any window interval of 200 seconds, at least 1% of transmitted timing packets will be received within a fixed cluster, starting at the observed floor delay, and having a range of  $150 \mu\text{s}$ .

The Maximum Time Interval Error (MTIE) mask is shown in Figure 2:



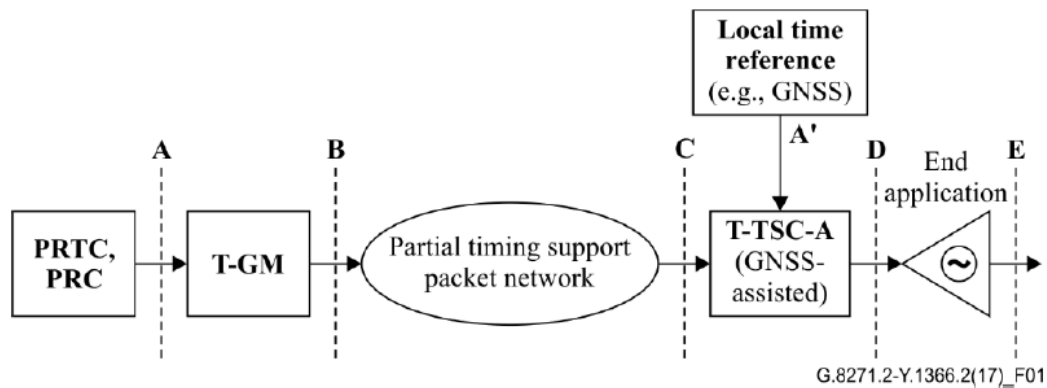
Observation interval $\tau$ (s)	MTIE requirement ( $\mu$ s)
$0.05 < \tau \leq 0.2$	$46 \tau$
$0.2 < \tau \leq 32$	9
$32 < \tau \leq 64$	$0.28 \tau$
$64 < \tau \leq 1\,125$	18
$\tau > 1\,125$	$0.016 \tau$



**Figure 2 – ITU-T G.8261.1 Network Limits for LTE-FDD**

In the APTS configuration, PTP is used as a backup timing source to a local time reference (e.g., Primary Reference Time Clock (PRTC) based on the Global Navigation Satellite System (GNSS)) for durations up to 72 hours. It is not intended to use PTP as the primary timing source. In this case, the PTP source is used to assist with holding an accurate frequency at the APTS clock when its primary phase source fails, avoiding it from drifting.

ITU-T G.8271.2 defines the network limits the network limits to achieve a reliable APTS secondary source. The reference model for the network limits is shown in Figure 3.



**Figure 3 – ITU-T G.8271.2 Reference Model for APTS**

Reference point C is the point where the APTS cell is connected (and where the DOCSIS network ends). The packet delay variation network limit given in ITU-T G.8271.2 is defined as below.

The network limit value and the metric processing parameters for assisted partial timing support use the metric pktSelected2wayTE (described in ITU-T G.8260) as follows:

- Peak-to-peak pktSelected2wayTE <1,100 ns
- Selection window = 200 seconds
- Selection percentage = 0.25%

This limit makes sure that when the GNSS source is lost, the absolute Time Error (TE) will remain within expected limits of **100 ns** while PTP is the selected fallback source.


To summarize, the access network must provide a PDV of less than 150  $\mu$ s and a fractional frequency offset of less than 16 ppb to provide an accurate frequency reference. For APTS support, the MTIE must also be less than 100 ns.

### 1.1.2. Proposed Solution over DOCSIS

There are mainly two options of supporting frequency delivery over DOCSIS:

- A. DOCSIS as a frequency aware network. In this option the DOCSIS Cable Modem Termination System (CMTS) or Remote PHY (R-PHY) locks its internal DOCSIS clock to an external source via SyncE or PTP. The DOCSIS Cable Modem (CM) is natively locked to the CMTS/R-PHY clock through the DOCSIS symbol clock. Lastly, the end slave clock is locked to the CM clock via SyncE or PTP. This option takes advantage of the fact that the DOCSIS system already provides accurate frequency sync between its components. This option is similar to the frequency delivered via physical clocks (SyncE) as described above.
- B. DOCSIS as frequency unaware network. In this option, the DOCSIS segment does not contain any of the frequency synchronization elements. Frequency delivery passes “over the top” as regular data packets using PTP messages. The main challenge in delivering timing information over the top of DOCSIS is that there is large latency asymmetry between the DOCSIS upstream and downstream paths. In addition, there is a significant PDV over DOCSIS mainly on the upstream path due to the DOCSIS scheduling. However, for frequency information delivery, asymmetry is not a factor. In addition, taking the fact that G.8265.1 allows the use of 1-way PTP delivery for frequency sync, eliminates the upstream PDV problem and routes the PTP traffic only over the limited PDV downstream path. The PDV can be reduced further by dedicating a DOCSIS service flow for the PTP traffic (UDP port 319).

In this paper, we will focus on the second option. The proposed solution for frequency delivery over the top of DOCSIS is shown in Figure 4.



**Figure 4 – Proposed Solution for LTE-FDD and APTS Support over DOCSIS**

## 1.2. LTE-TDD and 5G Support with an IEEE 1588-Aware DOCSIS Network


There are two synchronization approaches that allow phase synchronization for mobile backhaul.

- Full Timing Support for Phase Synchronization

- Partial Timing Support for Phase Synchronization

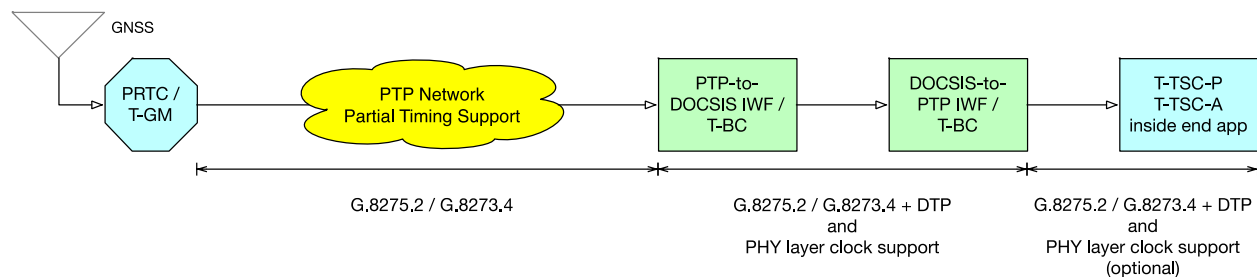
Full Timing Support networks are composed exclusively of network elements that support IEEE 1588 protocol operation. Partial Timing Support networks can include network elements that are not IEEE 1588 aware.

Figure 5 shows an end-to-end view of the synchronization flow for mobile backhaul with full timing support providing phase synchronization service to an end application mobile base station. PTP and SyncE are fully terminated at the PTP-to-DOCSIS Inter Working Function (IWF). DTP is used to transfer the synchronization information between DOCSIS equipment. The DOCSIS-to-PTP IWF regenerates synchronization information in a suitable format for use by the end application mobile base station.



**Figure 5 – End-to-End View of Full Timing Support for Phase Sync**

Figure 6 shows an end-to-end view of the synchronization flow for mobile backhaul with partial timing support providing phase synchronization service to an end application mobile base station. The DOCSIS interworking functions perform identical functions in a partial timing support network as in a full timing support network.



**Figure 6 – End-to-End View of Partial Timing Support for Phase Sync**

### 1.2.1. Synchronization Requirements

Phase synchronization is required for LTE-TDD, LTE-Advanced, and 5G.

LTE-TDD mobile cells use the same frequency for transmitting and receiving data, so it requires the time (phase) to be very accurate so there will no interference between different cell's transmissions.

LTE-TDD and 5G frequency and phase accuracy defined by 3GPP as shown in Table 1.

**Table 1 – 4G/LTE and 5G Synchronization Requirements**

	Frequency	Phase	Notes
4G LTE TDD	±50 ppb (air)	10 μs (wide: cell radius >3	Phase: 3GPP TS 36.133 §7.4.2

	Frequency	Phase	Notes
	$\pm 16$ ppb (network)	km $3 \mu\text{s}$ (local: cell radius <3 km)	Frequency: 3GPP TS 36.922 §6.4.1.2
<b>5G TDD</b>	$\pm 50$ ppb (air) $\pm 16$ ppb (network)	$\leq 3 \mu\text{s}$	3GPP TS 38.104 Table 6.5.1.2.1

ITU-T G.8271.1 provides examples of the time error budget allocations to meet the LTE-TDD and 5G requirements. Those are shown in Table 2.

\*Note that the time error specified in 3GPP is  $3 \mu\text{s}$  while the ITU-T budget is  $1.5 \mu\text{s}$  as the latter splits the overall phase error requirement between each of the Grand Master (GM) to cell paths.


**Table 2 – ITU-T G.8271.1 Example of Time Error Allocation**

Budget component	Failure scenario (a) ( <i>T-GM rearrangement</i> )		Failure scenario (b) ( <i>Short GNSS interruption</i> )	
<b>PRTC (<math>ce_{ref}</math>)</b>	100 ns		100 ns	
<b>Holdover and rearrangements in the network (<math>TE_{HO}</math>)</b>	NA		400 ns	
<b>Random and error due to synchronous Ethernet rearrangements (<math>dTE</math>)</b>	200 ns		200 ns	
<b>Node constant including intrasite (<math>ce_{ptp\_clock}</math>) (Notes 1 and 2)</b>	Type A 550 ns	Type B 420 ns	Type A 550 ns	Type B 420 ns
<b>Link asymmetries (<math>ce_{link\_asym}</math>) (Note 3)</b>	250 ns	380 ns	100 ns	230 ns
<b>Network limit at reference point C (<math>TE_C</math>)</b>	1 100 ns		1 350 ns (Note 4)	
<b>Rearrangements and short holdover in the end application (<math>TE_{REA}</math>)</b>	250 ns		NA	
<b>End application (<math>TE_{EA}</math>)</b>	150 ns		150 ns	
<b>Total limit at reference point E (<math>TE_E</math>)</b>	1 500 ns		1 500 ns	

From Table 2 above, the network is budgeted  $1.1 \mu\text{s}$  of the total  $1.5 \mu\text{s}$  for constant phase error. This includes all the network components between the GM and the end slave. If DOCSIS is part of that network, it will be part of this  $1.1 \mu\text{s}$  budget.

Beside the constant time error budget allocation, G.8271.1 specifies a dynamic low frequency TE network limit in terms of MTIE and shown in Figure 7.


MTIE limit (ns)	Observation interval, $\tau$ (s)
$100 + 75\tau$	$1.3 < \tau \leq 2.4$
$277 + 1.1\tau$	$2.4 < \tau \leq 275$
580	$275 < \tau \leq 10\,000$



**Figure 7 – ITU-T G.8271.1 Dynamic Time Error Network Limit (MTIE)**

### 1.2.2. Proposed Solution over DOCSIS

The DOCSIS Time Protocol (DTP) was introduced in the DOCSIS 3.1 specifications and allows for the passing of the IEEE 1588 protocol operation over the DOCSIS network with no jitter from network buffering. The DTP reference architecture is shown in Figure 8.



**Figure 8 – DTP Reference Architecture**

The CMTS and CM in Figure 8 corresponds to the PTP-to-DOCSIS IWF and the DOCSIS-to-PTP IWF respectively in Figure 5 and Figure 6. In R-PHY networks, the R-PHY replaces the CMTS in the timing path to the Customer Device.

The DOCSIS Time Protocol provides mechanisms to distribute frequency and phase information from the CMTS/R-PHY to a DTP-capable cable modem. The CMTS/R-PHY generates the downstream baud rate

from its system clock to allow DTP-enabled cable modems to lock to that frequency and obtain frequency synchronization. For phase information, the CMTS/R-PHY synchronizes the DOCSIS Extended Timestamp to the received PTP timestamp. The cable modem recovers this timestamp and aligns its PTP timestamp to it. The DTP protocol computes the downstream delay while accounting for the upstream asymmetry so that the resulting PTP timestamp from the cable modem is closely aligned to the received PTP timestamp at the CMTS/R-PHY. To perform this computation and obtain this close alignment, delay values for components in the CMTS/R-PHY, cable modem, and the DOCSIS network are required.

The DOCSIS CMTS/R-PHY will have a PTP stack, and the CM will have its own PTP stack, so the CMTS/R-PHY will appear as one PTP hop, and the DOCSIS CM will appear as a second PTP hop. As a result, the DOCSIS network will increment the PTP hop count by two, once for the CMTS/R-PHY and once for the CM.

In order to insert IEEE 1588-capable DOCSIS equipment within the network, the total DOCSIS synchronization portion of the chain typically should meet the following performance requirements when measured in isolation between the PTP input on the first IWF (PTP-to-DOCSIS) and the PTP output on the second IWF (DOCSIS-to-PTP). The performance is covered in G.8273.2, Appendix V, “Performance Estimation for Cascaded Media Converters acting as T-BCs” that summarizes noise generation estimation for a pair of media converters based on the Class A and Class B T-BC noise generation specifications and the source for the data in Table 3.

**Table 3 – Noise Generation Estimation for a Pair of Media Converters**

	Based on Class A T-BC			Based on Class B T-BC		
	Single T-BC	Pair	Pair DOCSIS Class A IWF	Single T-BC	Pair	Pair DOCSIS Class B IWF
<b>cTE (ns)</b>	±50	±100	±500	±20	±40	±250
<b>dTE<sub>L</sub> MTIE (ns)</b>	40	60	60	40	60	60
<b>dTE<sub>L</sub> TDEV (ns)</b>	4	6	6	4	6	6
<b>dTE<sub>H</sub> (peak-to-peak, ns)</b>	70	70	70	70	70	70
<b>max TE  (ns)</b>	100	160	560	70	100	310

## 2. Equipment and Lab Setup

The following equipment was used for the different experiments:

- Grandmaster clock with a GPS reference supporting G.8265.1 and G.8275.2 profiles
- IEEE 1588 unaware switch
- CMTS for both Integrated Converged Cable Access Platform (I-CCAP) and CCAP Core functionality with PTP and DTP support
- R-PHY
- DOCSIS 3.0 CMs
- DOCSIS 3.1 CMs with DTP and PTP support
- IEEE 1588 slave probe with a GPS reference
- 4-channel Oscilloscope to measure the 1PPS differences between the GM, CCAP/R-PHY, CM and the slave probe

## 2.1. LTE-FDD and APTS Support with an IEEE 1588-Unaware DOCSIS Network

ITU-T has defined the G.8265.1 telecom profile to provide frequency-only distribution using PTP. A telecom profile defines the parameters needed to guarantee protocol interoperability between implementations. It also specifies the optional features and default values that must be supported. It does not guarantee a specific level of performance.

In this networking scenario, the network is a single PTP domain, from the PTP master to the clock probe. The intermediate networking equipment is not aware that PTP traffic is being carried between the master clock and the slave clock. It does not terminate or monitor the messages or monitor the timestamps in the messages, thus it is defined as non-participating with the PTP protocol.

Some of the relevant configuration parameters of the G.8265.1 profile are listed below:

- Unicast message transmission
- One-way or two-way messaging
- One-step or two-step timestamping
- Message rates, from 1 per 16 seconds to 128 per second

The networking equipment classifies the PTP traffic as expedited forwarding traffic, which is the highest priority traffic in the network. In these scenarios, a low priority background traffic was not added to compete for queuing and scheduling resources.

As the number of non-participating nodes increases, the accumulation of packet delay and packet delay variation will reduce the ability of the slave clock to recover the master clock's frequency beyond an acceptable level.

A clock probe is used as the slave clock to measure the clock frequency and phase accuracy. Clock probes measure time error between a source timing signal and a measurement reference timing signal in terms of standard TIE/TE and MTIE measurements.

To reduce the PDV over the DOCSIS downstream path, the PTP packets were put into a dedicated DOCSIS service flow with highest priority. In addition, both DOCSIS 3.1 (192 MHz OFDM) and DOCSIS 3.0 (24 Annex B bonded SC-QAM) were tested.

### 2.1.1. I-CCAP

The first test configuration is a Centralized Access Architecture scenario. The configuration is shown in Figure 9.

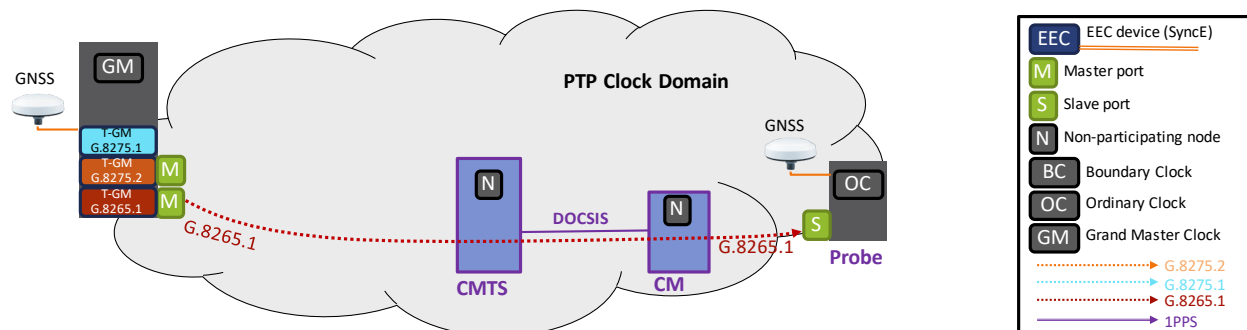


Figure 9 – PTP Frequency Delivery (G.8265.1) with I-CCAP

The network is a single PTP clock domain. The master clock sends PTP messages to the network side interface of the CMTS, which forwards them over DOCSIS to the cable modem. A clock probe is attached to the CMCI of the cable modem and terminates the PTP messages. The master clock and the clock probe are referenced to GPS, providing a common traceable source of frequency and time/phase information. This allows the clock recovery behavior of the slave clock to be measured against the network characteristics.

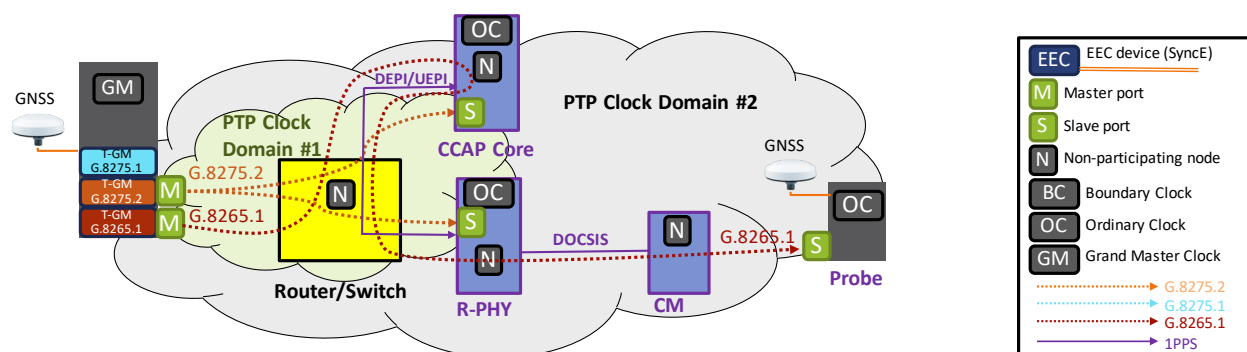
The PTP timestamped packets are generated at the T-GM and processed by the clock probe. The intermediate networking equipment is unaware that PTP messages are being carried, so they are sent over-the-top, and the CMTS and cable modem are unaware and non-participating in the protocol. The T-GM was connected to the I-CCAP via a 1GbE electrical interface. The cable modem was also connected to the clock probe via a 1GbE electrical interface.

The configured G.8265.1 parameters for this scenario are listed below:

- Unicast message transmission
- One-way messaging
- One-step timestamping
- Message rates: 128 per second (Sync)

### 2.1.2. Remote PHY

The second test configuration is a distributed access architecture scenario with a CCAP Core and an R-PHY connecting to a cable modem over DOCSIS. The configuration is shown in Figure 10.



**Figure 10 – PTP Frequency Delivery (G.8265.1) with R-PHY**

The network consists of two clock domains, one domain provides phase and frequency synchronization for the CCAP Core and the R-PHY for DAA operation and a second for over-the-top frequency synchronization of the clock probe. G.8275.2 was used as the profile for the DAA timing domain while G.8265.1 was used for the LTE-FDD (PTP over the top) timing domain.

The G.8275.2 master clock sends PTP messages through the non-participating switch to the network side interface of the CMTS and to the network side interface of the R-PHY. The CMTS and R-PHY are both slave clocks in the first PTP clock domain.

The G.8265.1 master clock sends PTP messages through the non-participating switch to the CMTS for encapsulation in DEPI pseudowires, back through the switch to the network side interface of the R-PHY, which de-encapsulates them from DEPI and forwards them over DOCSIS to the cable modem. A clock probe is attached to the CMCI of the cable modem and terminates the PTP messages. The G.8265.1



master clock and the clock probe are referenced to GPS, providing a common traceable source of frequency and time/phase information. This allows the clock recovery behavior of the slave clock to be measured against the network characteristics.

The PTP timestamped packets are generated at the T-GM and processed by the clock probe. The intermediate networking equipment is unaware that PTP messages are being carried, so they are sent over-the-top, and the switch, CMTS, R-PHY, and cable modem are unaware and non-participating in the PTP protocol. The T-GM was connected to the switch via a 1GbE electrical interface. The switch connects to the R-PHY through a 10GbE optical interface. The cable modem was also connected to the clock probe via a 1GbE electrical interface.

The configured G.8265.1 parameters for this scenario are listed below:

- Unicast message transmission
- One-way messaging
- One-step timestamping
- Message rates: 128 per second (Sync)

## **2.2. LTE-TDD and 5G Support with an IEEE 1588-Aware DOCSIS Network**

LTE-TDD uses a single frequency for transmitting and receiving data at different times, so it requires the highly accurate distribution and synchronization of frequency information and time or phase information between network elements across the network. ITU-T has defined the G.8275.1 and G.8275.2 telecom profiles to provide time and frequency distribution using PTP. In G.8275.1, timing support is provided by all elements in the network, while in G.8275.2, not all nodes need to provide timing support. For both profiles, physical layer functions like Synchronous Ethernet, may be used to stabilize the frequency operation of the clocks in the nodes and be provided on the CMCI output of the cable modem. Synchronous Ethernet was not used in either of the associated testing scenarios of this section.

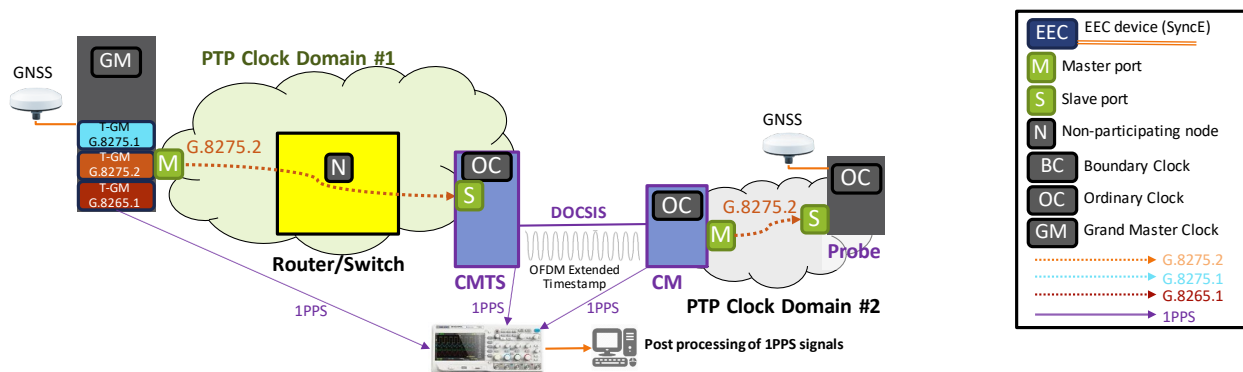
The DOCSIS Time Protocol provides mechanisms to distribute frequency and phase information from the CMTS to a DTP-capable cable modem. The CMTS can derive the downstream baud rate from its system clock to allow cable modems to lock to that frequency and obtain frequency synchronization. For phase information, the CMTS synchronizes the DOCSIS Extended Timestamp to the received PTP timestamp. The cable modem recovers the DOCSIS timestamp and aligns its PTP timestamp to it. The DTP protocol computes the downstream delay while accounting for the downstream and upstream asymmetries so that the resulting PTP timestamp from the cable modem is closely aligned to the received PTP timestamp at the CMTS.

A PTP Clock Probe uses the timestamp carried in PTP messages from the CM to measure the TE/TIE between the clock under test and the reference timing signal for PTP clock analysis. The PTP clock probe provides monitor access to the various PTP timestamps.

Due to lack of SyncE support on all DOCSIS components at the time, SyncE was not used for the tests; only PTP was used.

### **2.2.1. I-CCAP**

This test configuration is a Centralized Access Architecture scenario. This configuration is shown in Figure 11.



**Figure 11 – PTP Time and Frequency Delivery (G.8275.2) with I-CCAP**

The network consists of two PTP clock domains, one domain that provides phase and frequency synchronization of the CMTS to the T-GM and a second provides phase and frequency synchronization of the clock probe to the cable modem. G.8275.2 was used for both clock domains.

The T-GM clock sends PTP messages directly to the network side interface of the CMTS. The CMTS is a PTP slave clock in the first domain. It terminates the PTP messages and synchronizes the DOCSIS Extended Timestamp to the received PTP timestamp. The cable modem recovers this timestamp and aligns its PTP timestamp to it. It is the master clock for the second PTP domain and is connected to the slave clock probe using PTP through the CMCI.

The T-GM was connected directly to the CMTS through a 1GbE optical interface. The cable modem was connected directly to the slave clock probe via a 1GbE electrical interface.

The CMTS and the CM exchange DTP messages to calculate the DOCSIS asymmetry. The delay values in the DTP messages was determined manually.

The configured G.8275.2 parameters for the GM to CMTS domain are listed below:

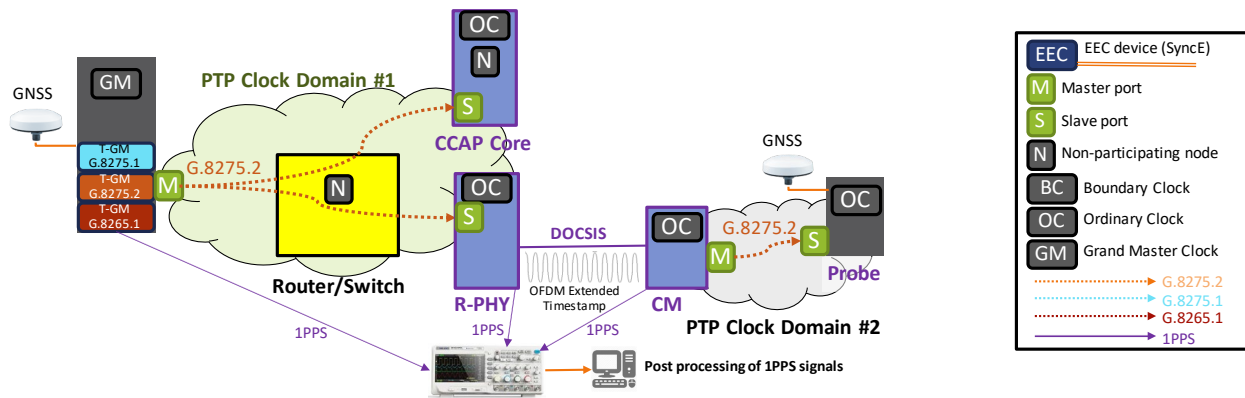
- Unicast message transmission
- Two-way messaging
- One-step timestamping
- Message rates: 64 per second (Sync and Delay)

The configured G.8275.2 parameters for the CM to slave probe domain are listed below:

- Unicast message transmission
- Two-way messaging
- Two-step timestamping
- Message rates: 128 per second (Sync and Delay)

### **2.2.2. Remote PHY**

The second test configuration is a distributed access architecture scenario with a CCAP Core and an R-PHY connecting to a cable modem over DOCSIS. This configuration is shown in Figure 12.



**Figure 12 – PTP Time and Frequency Delivery (G.8275.2) with R-PHY**

The network consists of two PTP clock domains, one domain that provides phase and frequency synchronization of the CCAP Core and the R-PHY to the T-GM for DAA operation and a second provides phase and frequency synchronization of the clock probe to the cable modem. G.8275.2 is used for both clock domains.

The G.8275.2 T-GM clock sends PTP messages through the non-participating switch to the network side interface of the CCAP Core and to the network side interface of the R-PHY. The CCAP Core and R-PHY are slave clocks in the first PTP clock domain. The R-PHY terminates the PTP messages and synchronizes the DOCSIS Extended Timestamp to the received PTP timestamp. The cable modem recovers this timestamp and aligns its PTP timestamp to it. It is the master clock for the second PTP domain and is connected to the slave clock probe using PTP through the CMCI.

Two alignment procedures were performed as part of this test. The first compensated for the delay asymmetry present in the first PTP clock domain between the R-PHY and the T-GM, primarily due to the different ethernet speeds used on the R-PHY (10GbE) and GM (1GbE). The second alignment procedure determined the appropriate DOCSIS Time Protocol values used during this test.

The T-GM was connected to the switch via a 1GbE electrical interface. The switch connects to the CMTS through a 10GbE optical interface. The cable modem connects to the slave clock probe via a 1GbE electrical interface.

The CMTS and the CM exchanged DTP messages to calculate the DOCSIS asymmetry. The delay values in the DTP messages was determined manually.

The configured G.8275.2 parameters for the GM to R-PHY domain are listed below:

- Unicast message transmission
- Two-way messaging
- One-step timestamping
- Message rates: 64 per second (Sync and Delay)

The configured G.8275.2 parameters for the CM to slave probe domain are listed below:

- Unicast message transmission
- Two-way messaging
- Two-step timestamping
- Message rates: 128 per second (Sync and Delay)

### 3. Procedure and Results

#### 3.1. LTE-FDD and APTS Support with an IEEE 1588-Unaware DOCSIS Network

As mentioned in section 2.1, PTP packets were sent over the top of DOCSIS using a 1-way G.8265.1 messaging between the GM and the probe.

Tests were performed on both I-CCAP and R-PHY setups and measurements of MTIE and PDV were taken and compared to ITU-T G.8261.1 and G.8271.2 requirements for LTE-FDD and APTS use cases respectively.

##### 3.1.1. I-CCAP

In the I-CCAP case, two experiments were performed. One using a DOCSIS 3.0 CM with 14 Annex B SC-QAM bonded downstream channels and one using a DOCSIS 3.1 CM with 192 MHz OFDM downstream channel. Results are shown for both.

##### 3.1.1.1. DOCSIS 3.0 I-CCAP

Figure 13 and Figure 14 show the PDV and MTIE results measured at the output of the CM by a probe over SC-QAM channels with DOCSIS 3.0.





Figure 13 – PDV for DOCSIS 3.0 with I-CCAP




**Figure 14 – G.8261.1 MTIE Results for DOCSIS 3.0 I-CCAP**

As can be seen from the above figures, the measured PDV at the output of the CM is roughly 20  $\mu$ s to 50  $\mu$ s, and the results meet the MTIE mask of G.8261.1 for LTE-FDD.

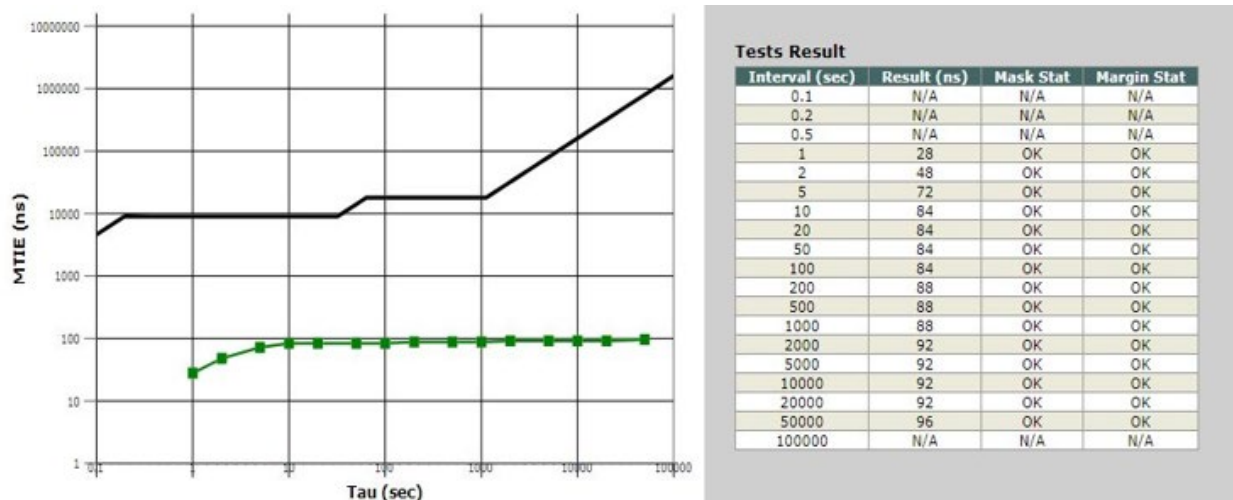
However, the MTIE is ~12  $\mu$ s, which does not meet the 100 ns requirement for APTS.

### 3.1.1.2. DOCSIS 3.1 I-CCAP

Figure 15 and Figure 16 show the PDV and MTIE results measured at the output of the CM by a probe over OFDM channels with DOCSIS 3.1.



**Figure 15 – PDV for DOCSIS 3.1 with I-CCAP**



**Figure 16 – G.8261.1 MTIE Results for DOCSIS 3.1 I-CCAP**


As can be seen from the above figures, the measured PDV at the output of the CM is roughly 10  $\mu$ s to 15  $\mu$ s, and the results meet the MTIE mask of G.8261.1 for LTE-FDD.

In addition, the MTIE is ~96 ns, which does meet the 100 ns requirement for APTS.

### 3.1.2. Remote PHY

In the R-PHY case, only the DOCSIS 3.1 experiment was performed to see whether the DAA degrades the performance or not.

Figure 17 shows the PDV results measured at the output of the CM by a probe.



**Figure 17 – PDV for D3.1 with R-PHY**

As can be seen from the above figures, the measured PDV at the output of the CM was very high with peaks going into the millisecond range. This is at least one order of magnitude higher than the threshold required by G.8261.1.

From the results it seems that the DEPI encapsulation/de-encapsulation in the Core and R-PHY adds a significant amount of PDV to the PTP packets.

When trying to connect a slave clock to this PTP output, the slave clock could not achieve frequency lock.

Therefore, neither LTE-FDD nor APTS can be supported with PTP over the top of a DOCSIS Remote PHY system.

### **3.2. LTE-TDD and 5G with an IEEE 1588-Aware DOCSIS Network**

There were no significant differences in results between the I-CCAP and R-PHY use cases. Therefore, the results shown are from the R-PHY setup only.

#### **3.2.1. Remote PHY**

##### **3.2.1.1. Test Calibration**

Two alignment procedures were performed as part of this test. The first compensated for the delay asymmetry present in the first PTP clock domain (GM to R-PHY). The second alignment procedure determined the appropriate DTP values to be used during this test.

The delay asymmetry value is configured to compensate for the difference in latency between the ingress and egress path between the R-PHY and its IEEE 1588 GM. When the ingress path is slower than the egress path, a positive asymmetry value is configured. A negative value is configured when ingress path is faster than the egress path. The delay asymmetry value was adjusted until the delay between the 1PPS output from the T-GM to the 1PPS output from the R-PHY was less than  $\pm 50$  ns.

The DTP algorithm is based on the true ranging offset (TRO) calculated by the CM based on the ranging information. The TRO is effectively a round trip delay measurement.

In the DTP algorithm, both the CMTS and CM are reporting their different US and DS timing path delays. To eliminate any HFC asymmetries, a (nearly) zero-length HFC plant was used.

In the R-PHY setup, the CCAP core is the one communicating DTP with the CM. However, the DS and US delay values were effectively tuned to those of the R-PHY.

Using an oscilloscope, the 1PPS output from the R-PHY and CM were aligned as close as possible limited by the 1PPS accuracy of the CM and the oscilloscope resolution.

The TRO value that was calculated by the CM for the zero-length plant was also read. The DTP math algorithm was performed to determine the required time adjustment value to be used by the CM. Once the appropriate delay values are calculated, the differences between a zero-length plant and a real plant that has the same R-PHY and CM are reflected in the TRO value.

The DTP alignment procedure was performed with a coax cable less than 3 feet long to represent the zero-length plant and those values were also used to initialize the R-PHY and CM pair with a 400-foot cable. The same DTP values were also used to initialize a second DTP-capable CM.

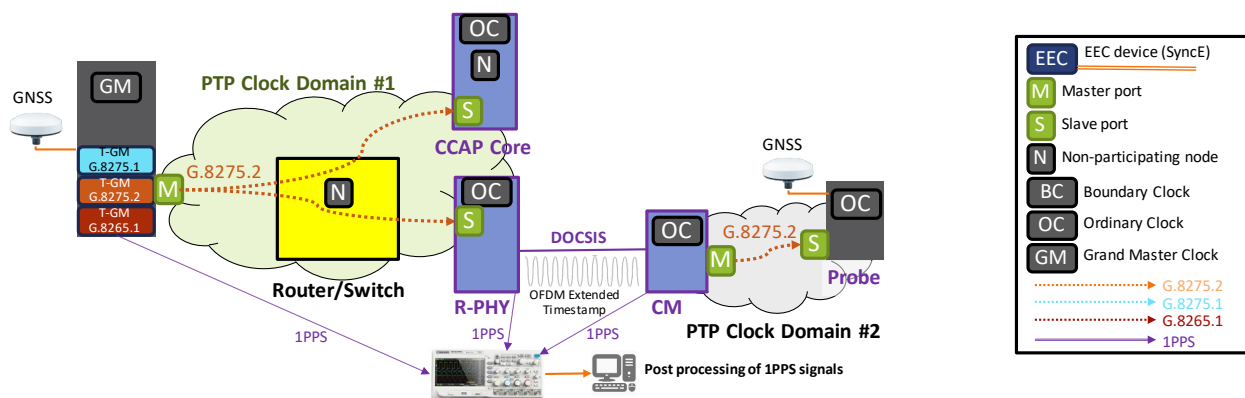


### 3.2.1.2. Key Measurements

The following measurements were taken throughout the tests:

- 1PPS differences between the R-PHY and CM in long runs to check the phase transfer stability. This was measured using an oscilloscope
- 1PPS differences between the R-PHY and CM after R-PHY/CM reset. This was measured using an oscilloscope
- Overall TE/TIE and MTIE measured by the PTP probe compared to GPS. This gives an estimate of the overall DOCSIS network phase delivery performance


Figure 18 shows the network configuration for PTP Time and Frequency Delivery using R-PHY. This is the same as Figure 12.



### Figure 18 – PTP Time and Frequency Delivery (G.8275.2) with R-PHY

### 3.2.1.3. Measurement Results


Measurements were made on the difference between the 1PPS signals from the R-PHY and the CM to determine the phase transfer stability. The zero-length plant coax cable was used. Figure 19 shows a set of 20 measurements collected after the DTP messages were exchanged and the CM adjusted its time accordingly. Those measurements spanned 23 ns peak-to-peak, ranging from  $-17$  ns to  $+6$  ns with respect to the average of the group. This was very consistent.



### Figure 19 – Phase Transfer Stability Between RPD and CM 1PPS




Figure 20 shows a set of 20 measurements that were taken after only the CM was reset and initialized using the same DTP parameters as in the initial set of measurements. The TRO calculated by the CM increased by 20 ns for this test. The 1PPS phase error spanned 20 ns peak-to-peak, ranging from –12 ns to +8 ns with respect to the average of the group. The average value for this group of measurements was 12 ns more than the average calculated in the first set of 20 measurements. This was also very consistent.



**Figure 20 – Phase Transfer Stability Between RPD and CM 1PPS (After CM Reset)**

Figure 21 shows the set of 20 measurements that were taken after both the RPD and CM were reset and initialized using the same DTP parameters as in the initial set of measurements. The TRO calculated by the CM increased by 39 ns in this test. The phase error spanned 22 ns peak-to-peak ranging from –10 ns to +12 ns with respect to the average value of the group. The average value for this group of measurements was 47 ns more than the average value calculated in the first set of 20 measurements. This is still very consistent, but it shows that there is some variability in the values after RPD reset. In any case, the difference is less than 50 ns.



**Figure 21 – Phase Transfer Stability Between R-PHY and CM 1PPS (After both CM and RPD Reset)**

Table 4 summarizes the phase transfer stability measurements. The average value for the measurements after the CM resets and after both CM and RPD resets is relative to the average for the initial set of

measurements. Similarly, the TRO measurements are similarly relative to the TRO from the initial set of measurements. The absolute values are not shown.

As can be seen the phase transfer from the RPD to the CM is stable with variances of less than 50 ns in long run and after CM and RPD resets. The variance in TRO indicated small variances in the DS and/or US paths across reboots. This variance may be reduced in the future after further study.


**Table 4 – Summary of Phase Transfer Stability Measurements**

	Average Phase Error (ns)	Phase Error Range (ns)	TRO (ns)
<b>Initial set of measurements</b>	0	–17 to +6	–
<b>Measurements after CM reset (relative to initial set of measurements for average and TRO)</b>	+12	–12 to +8	+20
<b>Measurements after both R-PHY and CM reset (relative to initial set of measurements for average and TRO)</b>	+47	–10 to +12	+39

Figure 22 shows the TE of the recovered phase at the slave probe referenced to GPS time. The measurement was taken with a 3-foot coax between the RPD and the CM to approximate the zero-length plant for calibration.

In the graph, it can be seen the recovered phase has a TE of roughly 220 ns with a variation of 100 ns peak-to-peak.


These results are well below the 500 ns TE budget from Section 1.2.2 for a class B DOCSIS system defined in the CableLabs Synchronization Techniques specification.



**Figure 22 – TE of the Recovered Phase at the Slave Probe with 3' of Coax**

In order to reduce the 220 ns cTE artificially, the DTP parameters were adjusted to take into account this 200 ns TE.

Figure 23 shows the TE of the recovered phase at the slave probe compared to GPS time after the DTP parameters were calibrated. The figure shows a compensated TE of roughly 30 ns with a variation of 60 ns peak-to-peak.




**Figure 23 – TE of the Recovered Phase at the Slave Probe with DTP Compensation**

In order to check the consistency of the DTP and TRO measurement accuracy, we changed the path between the RPD and the CM to a 400' length of coax. The DTP parameters were unchanged. The TRO increased by approximately 900 ns over the 3-foot plant values.

Figure 24 shows the TE of the recovered phase at the slave probe compared to GPS with a 400-foot coax between the RPD and the CM.

The figure shows a TE of roughly 10 ns with a variation of 50 ns peak-to-peak.


These results are quite consistent with the result of the 3-foot plant and shows that the DTP and TRO calculations are accurate.



**Figure 24 – TE of the Recovered Phase at the Slave Probe with 400' of Coax**

Figure 25 shows the MTIE measurement taken by the slave probe.

The MTIE performance of at the output of the CM is below 100 ns and meets the MTIE requirements for phase delivery defined in G.8271.1.




**Tests Result**

Interval (sec)	Result (ns)	Mask Stat	Margin Stat
0.1	N/A	N/A	N/A
0.2	N/A	N/A	N/A
0.5	N/A	N/A	N/A
1	24	N/A	N/A
2	40	OK	OK
5	44	OK	OK
10	48	OK	OK
20	48	OK	OK
50	56	OK	OK
100	60	OK	OK
200	80	OK	OK
500	88	OK	OK
1000	88	OK	OK
2000	88	OK	OK
5000	N/A	N/A	N/A
10000	N/A	N/A	N/A
20000	N/A	N/A	N/A
50000	N/A	N/A	N/A
100000	N/A	N/A	N/A


**Figure 25 – G.8271.1 MTIE Measurements Calculated by the Slave Probe with 400' of Coax**

In order to estimate the PDV and asymmetry of the PTP output from the CM, we used the PTP probe to compare the offset between the T1 and T4 timestamps to GPS time. This measures the PDV of the PTP packets between the CM and PTP slave probe for the forward and reverse direction. It can also show the delay asymmetry of the path.

Figure 26 and Figure 27 show the PTP probe for Probe-CM (upstream direction) and CM-Probe (downstream direction) paths respectively.



**Figure 26 – PTP Probe (Probe to CM)**



**Figure 27 – PTP Probe (CM to Probe)**

As can be seen when comparing the two PTP paths, the PDV is quite similar at 25 ns. The difference in delay between the upstream and downstream paths is roughly 400 ns. The asymmetry value is roughly 400 ns with a median value of ~200 ns. This is as expected since the recovered clock has a n offset of 200 ns. Each path has a 200 ns delay from the recovered clock on opposite directions.

## Conclusion

In this paper, we constructed and tested DOCSIS CAA and DAA networks to provide sufficiently accurate synchronization for various LTE-FDD, LTE-TDD, and 5G use cases.

We showed that for LTE-FDD and APTS, a DOCSIS CAA system can carry over the top PTP information and still maintain accurate synchronization. This approach can be deployed today with minimal changes.

The IEEE 1588-unaware over-the-top DOCSIS DAA network did not meet the required performance for LTE-FDD and APTS.

For LTE-TDD and 5G, both DOCSIS CAA and DAA systems can provide accurate sync delivery using DTP and meet the requirements specified for MBH. This solution requires the full support of PTP and DTP on the CMTS/R-PHY and CM as defined in the CableLabs Synchronization Techniques specification.

Network and module improvement such as using SyncE as part of the synchronization chain as well as embedding a switch within the CM with their impact on the overall frequency and phase accuracy for these and other applications are areas for additional studies.

# Abbreviations

1GbE	1 Gb/s Ethernet
10GbE	10 Gb/s Ethernet
1PPS	1 Pulse Per Second
3GPP	3rd Generation Partnership Project
APTS	Assisted Partial Timing Support
CAA	Centralized Access Architecture
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMCI	Cable Modem to Customer Premises Equipment Interface
CMTS	Cable Modem Termination System
cTE	Constant Time Error
DAA	Distributed Access Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data-Over-Cable Service Interface Specifications
DS	Downstream
dTE	Dynamic Time Error
DTP	DOCSIS Time Protocol
FPP	Floor Packet Percentage
GM	Grand Master
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HFC	Hybrid Fiber Coax
I-CCAP	Integrated CCAP
IEEE	Institute of Electrical and Electronic Engineers
IEEE 1588	IEEE Std 1588-2008 (PTP)
ITU-T	Telecommunication Standardization Sector of the International Telecommunication Union
IWF	Inter Working Function
LTE	Long-Term Evolution
LTE-A	Long-Term Evolution Advanced
LTE-FDD	Long-Term Evolution Frequency Division Duplex
LTE-TDD	Long-Term Evolution Time Division Duplex
MBH	Mobile Backhaul
MTIE	Maximum Time Interval Error
OFDM	Orthogonal Frequency Division Multiplexing
OTT	Over The Top
PDV	Packet Delay Variation
PHY	Physical
ppb	parts per billion
PTP	Precision Time Protocol
QAM	Quadrature Amplitude Modulation

R-PHY	Remote PHY
RPD	Remote PHY Device
SC-QAM	Single Carrier QAM
SyncE	Synchronous Ethernet
T-BC	Telecom Boundary Clock
T-GM	Telecom Grand Master
TDEV	Timing Deviation
TE	Time Error
TIE	Time Interval Error
TRO	True Ranging Offset
UEPI	Upstream External PHY Interface
US	Upstream

## Bibliography and References

3GPP TS 36.104, “Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) radio transmission and reception.”

3GPP TS 36.133, “Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management.” IEEE Std 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, IEEE Instrumentation and Measurement Society, 24 July 2008.

*IEEE Std 1588-2008, “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems”.*

*ITU-T G.8261.1/Y.1361.1 Packet delay variation network limits applicable to packet-based methods (Frequency synchronization) (05/14).*

*ITU-T G.8265.1/Y.1365.1, Precision time protocol telecom profile for frequency synchronization (07/14).*

*ITU-T G.8271.1/Y.1366.1, Network limits for time synchronization in packet networks (03/18).*

*ITU-T G.8271.2/Y.1366.2 Network limits for time synchronization in packet networks with partial timing support from the network (11/18).*

*ITU-T G.8275.1/Y.1369.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network (08/17).*

*ITU-T G.8275.2/Y.1368.2, Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network (08/17).*

*Data-Over-Cable Service Interface Specifications Modular Headend Architecture v2 DOCSIS Remote Out-of-Band Specification, CM-SP-R-PHY-III-180926, September 26, 2018, Cable Television Laboratories, Inc.*

*DOCSIS 3.1, Media Access Control (MAC) and Upper Layer Protocols Interface Specification, CM-SP-MULPhv3.1-III-180926, September 26, 2018, Cable Television Laboratories, Inc.*

*Data-Over-Cable Service Interface Specifications Modular Headend Architecture v2 Remote Upstream External PHY Interface Specification CM-SP-R-UEPI-I09-180926*, September 26, 2018, Cable Television Laboratories, Inc.

*Data-Over-Cable Service Interface Specifications Modular Headend Architecture v2 Remote Downstream External PHY Interface Updates CM-SP-R-DEPI-I11-180926*, September 26, 2018, Cable Television Laboratories, Inc.

*Data-Over-Cable Service Interface Specifications Modular Headend Architecture v2 Remote DOCSIS Timing Interface CM-SP-RDTI-I10-180509*, May 9, 2018, Cable Television Laboratories, Inc.

*Synchronization Techniques for DOCSIS® Specification CM-SP-SYNC-D02-190419*, April 19, 2019, Cable Television Laboratories, Inc.

John T. Chapman, Jennifer Andreoli-Fang, “Mobile Backhaul Synchronization Architecture,” proceedings of SCTE Fall Technical Forum, 2017.



# **Capacity Planning, Traffic Engineering, and HFC Plant Evolution for the Next 25 Years**

## **How and When HFC Network Technologies Will Need to Adapt to Support Future Bandwidth Growth**

A Technical Paper prepared for SCTE•ISBE by

**Tom Cloonan**

CTO- Network Solutions

CommScope

2400 Ogden Ave.- Suite 180, Lisle, IL 60532

630-281-3050

tom.cloonan@commscope.com

**Ayham Al-Banna**, Engineering Fellow, CommScope,

ayham.al-banna@commscope.com

**Frank O’Keeffe**, Distinguished System Engineer, CommScope,

frank.o’keeffe@commscope.com

**John Ulm**, Engineering Fellow, CommScope,

johh.ulm@commscope.com

**Ruth Cloonan**, CEO, BlueOpus,

blueopus11@yahoo.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	5
Content .....	8
1. Overview of Bandwidth Trends and Attributes Driving the Future of HFC.....	8
2. Overview of Some Key Technology Candidates for Future Upstream/Downstream Bandwidth Augmentation.....	17
2.1. Traditional-Full-Duplex DOCSIS (Traditional-FDX) for Upstream Augmentation .....	18
2.2. 204 MHz Frequency Division Duplex (FDD) High-Splits for Upstream Augmentation.....	20
2.3. 204-684 MHz Frequency Division Duplex (FDD) Ultra-Splits for Upstream Augmentation .....	21
2.4. Static Soft-Full-Duplex DOCSIS (Static Soft-FDX) for Upstream and Downstream Bandwidth Augmentation.....	23
2.5. Dynamic Soft-Frequency-Division Duplex (Dynamic Soft-FDX) for Upstream and Downstream Bandwidth Augmentation .....	25
2.6. Extended Spectrum DOCSIS (ESD) for Upstream and Downstream Bandwidth Augmentation.....	30
2.7. Distributed Node Architectures and Active Taps for Futuristic Upstream and Downstream Bandwidth Augmentation .....	33
3. Analysis of Migration Paths for Different Architectures.....	34
Conclusions.....	47
Abbreviations.....	55
Bibliography & References .....	57

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Competitor PON Solutions & Capabilities .....	6
Figure 2 – Different Technology Paths MSOs May Utilize in the Future .....	7
Figure 3 – Downstream Average Bandwidth Trends.....	8
Figure 4 – Upstream Average Bandwidth Trends .....	9
Figure 5 – Downstream & Upstream Maximum Bandwidth Trends (Nielson’s Law).....	10
Figure 6 – Downstream & Upstream Maximum Bandwidth Trends (“Slowed” Nielson’s Law for Asymmetrical Services).....	11
Figure 7 – Downstream & Upstream Maximum Bandwidth Trends (“Slowed” Nielson’s Law for Symmetrical Services).....	12
Figure 8 – Changes in Four Different Areas to Support Future Bandwidth Growth.....	18
Figure 9 – Spectrum for 85 MHz FDD Mid-Split & 396 MHz Traditional-FDX (w/ First-Order Guestimates on Bandwidth Capacities) .....	19
Figure 10 – Interference Group Elongation Problem That May Occur With FDX Amplifiers (Not A Problem For FDD High-Split or FDD Ultra-Split Approaches).....	20
Figure 11 – Spectrum for 85 MHz FDD Mid-Split & 204 MHz FDD High-Split (w/ First-Order Guestimates on Bandwidth Capacities) .....	21

Figure 12 – Spectrum for 85 MHz FDD Mid-Split & 396 MHz FDD Ultra-Split (w/ First-Order Guestimates on Bandwidth Capacities) .....	23
Figure 13 – Static Soft-FDX .....	24
Figure 14 – Examples of Yearly Spectrum Changes using Static Soft-FDX .....	25
Figure 15 – Dynamic Soft-FDX .....	26
Figure 16 – Examples of Second-by-Second Spectrum Changes using Dynamic Soft-FDX.....	26
Figure 17 – Interference Groups in Traditional-FDX and RF Leg-based Dynamic Soft-FDX (for Node+0 and Node+1 Systems).....	27
Figure 18 – FDX Operation vs Sliding FDD Operation at Various Levels of FDD Systems, Traditional-FDX Systems, & Dynamic Soft-FDX Systems.....	28
Figure 19 – Taxonomy of Upstream Augmentation Solutions with a Continuum of FDX Solutions (including Tradition-FDX and Dynamic Soft-FDX and Static Soft-FDX) .....	29
Figure 20 – Launch Power Spectral Density and Total Composite Power for a 3 GHz Extended Spectrum DOCSIS System .....	32
Figure 21 – OFDM Bit-loading & Corresponding Throughput for a 3 GHz Extended Spectrum DOCSIS Node+0 System .....	33
Figure 22 – Predicted 1.8 GHz Bandwidth Capacities as a Function of Amp-to-Amp Hard-line Coaxial Length and Amplifier Output Power Levels and Cascade Lengths .....	33
Figure 23 – Migration Path for Architecture 1a (Traditional-FDX- Asymmetric SLA).....	38
Figure 24 – Migration Path for Architecture 2a (Static Soft-FDX Baseline w/ 1.2 GHz Affinity- Asymmetric SLA) .....	39
Figure 25 – Migration Path for Architecture 3a (Static Soft-FDX w/ Node+3 Affinity - Asymmetric SLA) .....	39
Figure 26 – Migration Path for Architecture 4a (Static Soft-FDX w/ 15% DS Tmax CAGR- Asymmetric SLA) .....	40
Figure 27 – Migration Path for Architecture 5a (Static Soft-FDX w/ Reduced US Tmax- Asymmetric SLA) .....	40
Figure 28 – Migration Path for Architecture 6a (Static Soft-FDX w/ Selective Subscriber Migration- Asymmetric SLA) .....	41
Figure 29 – Migration Path for Architecture 7a (Static Soft-FDX w/ Guard-band Elimination- Asymmetric SLA) .....	41
Figure 30 – Migration Path for Architecture 8a (Dynamic Soft-FDX Baseline- Asymmetric SLA) .....	42
Figure 31 – Migration Path for Architecture 9a (Static Soft-FDX w/ Active Taps- Asymmetric SLA).....	42
Figure 32 – Migration Path for Architecture 1b (Traditional-FDX- Symmetric SLA).....	43
Figure 33 – Migration Path for Architecture 2b (Static Soft-FDX Baseline w/ 1.2 GHz Affinity- Symmetric SLA).....	43
Figure 34 – Migration Path for Architecture 3b (Static Soft-FDX w/ Node+3 Affinity - Symmetric SLA) ....	44
Figure 35 – Migration Path for Architecture 4b (Static Soft-FDX w/ 15% DS Tmax CAGR- Symmetric SLA) .....	44
Figure 36 – Migration Path for Architecture 5b (Static Soft-FDX w/ Reduced US Tmax- Symmetric SLA) .....	45
Figure 37 – Migration Path for Architecture 6b (Static Soft-FDX w/ Selective Subscriber Migration- Symmetric SLA).....	45
Figure 38 – Migration Path for Architecture 7b (Static Soft-FDX w/ Guard-band Elimination- Symmetric SLA) .....	46
Figure 39 – Migration Path for Architecture 8b (Dynamic Soft-FDX Baseline- Symmetric SLA).....	46
Figure 40 – Migration Path for Architecture 9b (Static Soft-FDX w/ Active Taps- Symmetric SLA).....	47

Figure 41 – “Big Changes” for Architecture 1a (Traditional-FDX- Asymmetric SLA) .....	47
Figure 42 – “Big Changes” for Architecture 2a (Static Soft-FDX Baseline w/ 1.2 GHz Affinity- Asymmetric SLA) .....	48
Figure 43 – “Big Changes” for Architecture 3a (Static Soft-FDX w/ Node+3 Affinity - Asymmetric SLA) .....	48
Figure 44 – “Big Changes” for Architecture 4a (Static Soft-FDX w/ 15% DS Tmax CAGR- Asymmetric SLA) .....	48
Figure 45 – “Big Changes” for Architecture 5a (Static Soft-FDX w/ Reduced US Tmax- Asymmetric SLA) .....	48
Figure 46 – “Big Changes” for Architecture 6a (Static Soft-FDX w/ Selective Subscriber Migration- Asymmetric SLA) .....	49
Figure 47 – “Big Changes” for Architecture 7a (Static Soft-FDX w/ Guard-band Elimination- Asymmetric SLA) .....	49
Figure 48 – “Big Changes” for Architecture 8a (Dynamic Soft-FDX Baseline- Asymmetric SLA) .....	49
Figure 49 – “Big Changes” for Architecture 9a (Static Soft-FDX w/ Active Taps- Asymmetric SLA) .....	49
Figure 50 – “Big Changes” for Architecture 1b (Traditional FDX- Symmetric SLA) .....	50
Figure 51 – “Big Changes” for Architecture 2b (Static Soft-FDX Baseline w/ 1.2 GHz Affinity- Symmetric SLA) .....	50
Figure 52 – “Big Changes” for Architecture 3b (Static Soft-FDX w/ Node+3 Affinity - Symmetric SLA) ....	50
Figure 53 – “Big Changes” for Architecture 4b (Static Soft-FDX w/ 15% DS Tmax CAGR- Symmetric SLA) .....	50
Figure 54 – “Big Changes” for Architecture 5b (Static Soft-FDX w/ Reduced US Tmax- Symmetric SLA) .....	51
Figure 55 – “Big Changes” for Architecture 6b (Static Soft-FDX w/ Selective Subscriber Migration- Symmetric SLA) .....	51
Figure 56 – “Big Changes” for Architecture 7b (Static Soft-FDX w/ Guard-band Elimination- Symmetric SLA) .....	51
Figure 57 – “Big Changes” for Architecture 8b (Dynamic Soft-FDX Baseline- Symmetric SLA) .....	51
Figure 58 – “Big Changes” for Architecture 9b (Static Soft-FDX w/ Active Taps- Symmetric SLA) .....	52

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Relationships between Fiber Depth, # of HHP, & # of Subs .....	16
Table 2 – US & DS BW Capacities for Various FDD Ultra-Split Frequency Bands Illustrating the Zero-Sum Game .....	23
Table 3 – Examples of Extended Spectrum DOCSIS Variants & their Predicted Performance Levels (Zero-Order Models) .....	31

# Introduction

Since their inception, Hybrid Fiber-Coax (HFC) networks have long been an evolving and changing infrastructure, continually adding new incremental changes in technology and delivering ever-increasing Bandwidth Capacities to accommodate the needs of their various services (Video, High Speed Data (HSD), and Voice). MSOs have long recognized that the HFC plant contains vast quantities of un-tapped Bandwidth Capacity, which can usually be enabled in a gradual, cost-effective, “just-in-time” fashion using minor evolutionary transitions applied intelligently to selected piece-parts of the network. This relatively low-cost, evolutionary approach to network transition has been quite successful and is usually preferred over more expensive revolutionary changes (ex: switching to FTTH) that attempt to change (or replace) a large amount of the HFC plant equipment and head-end equipment and in-home equipment all at once.

The last few years have produced a vast array of new and exciting technology ideas for the future, and they all have value. However, the list of options has been deemed to be too long and quite confusing by many MSOs. The following are examples of some candidate technologies being considered in the confusingly large list of options:

- Management/MAC Placement Variations
  - Centralized Access Architectures (CAAs)
    - Integrated CCAPs
    - M-CMTSS
    - Physical CCAP Cores + Remote PHY Shelves
    - Virtual CCAP Cores + Remote PHY Shelves
    - Chassis-based OLTs
  - Distributed Access Architectures (DAAs)
    - Physical CCAP Cores + Remote PHY Devices (RPDs)
    - Virtual CCAP Cores + Remote PHY Devices (RPDs)
    - Remote MACPHY Devices (RMDs)
    - Remote MAC Cores (RMCs)
    - Remote OLTs (R-OLTs)
  - Virtualization
    - Virtualization of Control/Management Planes (vMgr)
    - Software Defined Networking (SDN)
    - Virtualization of Data Planes (vCore)
- Upstream (US) and Downstream (DS) Bandwidth (BW) Augmentation Variations
  - Full Duplex DOCSIS (FDX)
  - FDX Amps
  - Upstream Extended Spectrum DOCSIS (ESD) using 204-684 MHz Ultra-Splits and beyond
  - Downstream Extended Spectrum DOCSIS (ESD) using 1.8 GHz spectra and beyond
  - Static Soft Frequency Division Duplex (Static Soft-FDD)
  - Dynamic Soft Frequency Division Duplex (Dynamic Soft-FDD)
  - Active Taps
- Fiber Depth Variations
  - Node+0
  - Node+Non-Zero
  - Distributed Node Architectures (DNA)
  - Fiber-To-The-Last-Active (FTTLA)

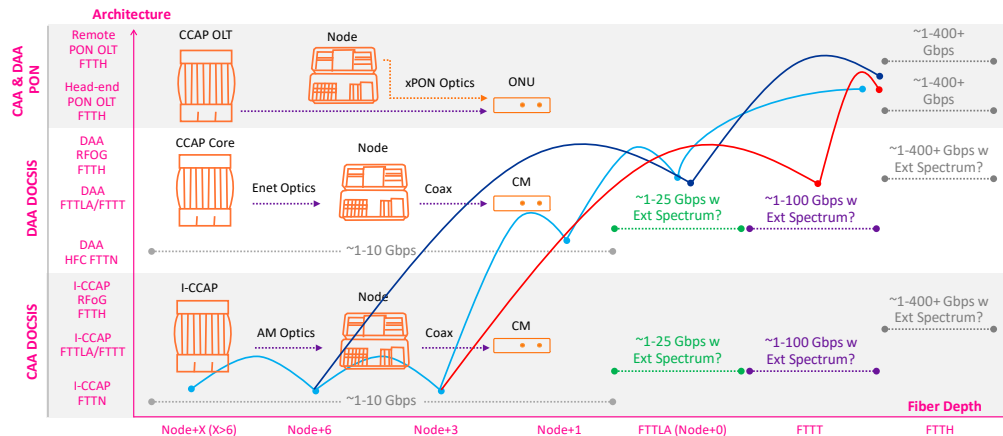
- Fiber-To-The-Tap (FTTT)
- Fiber-To-The-Home (FTTH).

How do we (as an industry) make sense of all of these options? And how do we stay competitive in an ever-changing market-place? In coming years, Multiple System Operators (MSOs) will also be challenged by a set of HSD competitors offering new wireless-based 5G solutions and fiber-based (Passive Optical Network or PON) solutions. Some details on the competitor PON solutions are shown below in Figure 1. The rows shown in red are likely to be the near-term competitors with which HFC plants will need to contend. The rows shown in black are likely to be the longer-term competitors with which HFC plant augmentations will need to contend.

Type of OLT	Down-stream Capacity	Up-stream Capacity
<b>GPON &amp; Turbo-Mode EPON</b>	<b>2.5 Gbps</b>	<b>1.25 Gbps</b>
<b>XG-PON</b>	<b>8.6 Gbps</b>	<b>2.4 Gbps</b>
<b>Asym XGS-PON (low-cost)</b>	<b>8.6 Gbps</b>	<b>2.4 Gbps</b>
<b>Sym XGS-PON (high-cost)</b>	8.6 Gbps	8.6 Gbps
<b>NG-PON2</b>	8.6 Gbps	8.6 Gbps
<b>10G EPON TDMA</b> Downstream Dual-rate WDM Upstream Dual-rate TDMA	9.7 Gbps 10.7 Gbps	8.6 Gbps
<b>10G EPON WDMA</b> Downstream Dual-rate WDM Upstream Dual-rate WDMA	9.7 Gbps 10.7 Gbps	9.6 Gbps
<b>10G EPON WDMA "Mixed Mode"</b> Downstream Dual-rate WDM Upstream Dual-rate WDMA Mixed Mode	9.7 Gbps 10.7 Gbps	9.6 Gbps

**Figure 1 – Competitor PON Solutions & Capabilities**

The higher Bandwidth Capacities permitted by these new competitor offerings must be carefully considered by MSOs as they plan their HFC network evolutions and other technology evolutions for the next decade and beyond. MSOs must be able to easily evolve their network to compete, and according to some analysts and the Cable Industry's own 10G proponents, the Bandwidth Capacities of the future may require 10 Gbps (or higher) by 2030 (or sooner). Some also predict a need to support Symmetrical Services in the very near future (implying equal or close-to-equal Maximum Throughputs (Tmax's) on Upstream and Downstream Service Level Agreements (SLAs)- possible with 10 Gbps Upstream and Downstream SLAs needed in the 2030s). Low-latency transport will also become imperative to support gaming and Virtual Reality and autonomous vehicle navigation and various advanced mobile services. Supporting these features (and others) will likely require several phased evolutionary changes to the HFC network over the next ten to twenty years, with each MSO choosing a potentially different path (as shown in Figure 2).



**Figure 2 – Different Technology Paths MSOs May Utilize in the Future**

Thus, it is clear that MSOs have arrived at an epoch in time where important decisions must be made to help the MSOs select their optimal paths. It is likely that different MSOs will have differing constraints, so there are likely to be various paths selected by different MSOs. However, it is probably beneficial for the industry to begin paring down the large, confusing list above. This would permit MSOs and vendors alike to focus on optimizing and cost-reducing the small subset of remaining solutions that will truly make the most positive impact on the cable industry going forward.

As MSOs make their HFC plant evolution decisions, they must consider their current constraints; but they must also consider the longer-term impact of their decisions on their future HFC network evolution path. Typically, Cable Modems (CMs) and Cable Modem Termination Systems (CMTSs) can be changed out and upgraded quite regularly as new technologies materialize (once every 5-7 years), but the general rule-of-thumb in the industry is that Outside Plant equipment (Nodes, Amplifiers, Taps, etc.) should remain in the field for 10-15 years or longer. As a result, Outside Plant changes deployed in 2029 may need to live in the field until 2044. For that reason, this paper will attempt to look out 25 years into the future to the year 2044- this will undoubtedly result in guestimates that are likely to be incorrect, but the hope is that this paper stimulates important industry-wide discussions on the evolution into that unclear future.


Thus, using reasonable estimates on a) future bandwidth growth, b) future technology challenges from competitor service providers, and c) possible future technologies permitted within the HFC plant, this paper will attempt to provide MSOs with a guide as they make the important decisions that will help them to migrate their HFC plants over the next 25 years.

# Content

## 1. Overview of Bandwidth Trends and Attributes Driving the Future of HFC

In a very real sense, the Cable Industry is predominantly an industry focused on managing bandwidth delivery to and from subscribers. This is true for all services- High-Speed Data, Voice, and Video. As a result, predicting the Bandwidth Capacity trends for these future services is a difficult, but important task. Some key trends and predictions are illustrated in Figures 3 to 7.

### Downstream Tavg @ Peak Busy Hour




- DS Tavg **approaching 2 Mbps** in 2019
  - 2019 YoY (22%) drops over half from 2018 YoY (51%)
- DS Tavg 5-yr **CAGR moves to 38% (~40%)**
  - MSOs' 5-yr CAGRs range from ~27% to ~49%

Figure 3 – Downstream Average Bandwidth Trends



## Upstream Tavg @ Peak Busy Hour




- US Tavg **basically flat at 140 kbps** in 2019
  - Flat 2019 (YoY = 2%) compared to big jump last year (2018 YoY = 46%)
  - 2019 DS growth gains ground on 2019 US growth!
- US Tavg 5-yr **CAGR still 19% (~20%)**

Figure 4 – Upstream Average Bandwidth Trends

# Nielson Law Downstream & Upstream Tmax

**“TRADITIONAL” NIELSEN’S LAW OF INTERNET BANDWIDTH GROWTH**

**(Growth Rate =50%/YEAR ... 10G DS SLA in 2026... 100G DS SLA in 2032)**




- Predictions:
  - DS Tmax CAGR stays at 50%
  - Symmetrical Service w/ US Tmax = 100% of DS Tmax

**Figure 5 – Downstream & Upstream Maximum Bandwidth Trends (Nielson’s Law)**

# Asymmetrical Services Downstream & Upstream Tmax

**“SLOWED” NIELSEN’S LAW OF INTERNET BANDWIDTH GROWTH**

**(Growth Rate =25%/YEAR (2020s)... 15% (after)... 10G DS SLA in 2029)**




- Predictions:
  - DS Tmax CAGR is 25% (in 2020s), 15% (thereafter)
  - Asymmetrical Service w/ US Tmax = 50% of DS Tmax

**Figure 6 – Downstream & Upstream Maximum Bandwidth Trends (“Slowed” Nielson’s Law for Asymmetrical Services)**

# Symmetrical Services Downstream & Upstream Tmax

**“SLOWED” NIELSEN’S LAW OF INTERNET BANDWIDTH GROWTH**

**(Growth Rate =25%/YEAR (2020s)... 15% (after)... 10G DS & US SLA in 2029)**



- Predictions:
  - DS Tmax CAGR is 25% (in 2020s), 15% (thereafter)
  - Symmetrical Service w/ US Tmax = 100% of DS Tmax

**Figure 7 – Downstream & Upstream Maximum Bandwidth Trends (“Slowed” Nielson’s Law for Symmetrical Services)**

All of the above Figures hint at the coming challenges for the future HSD network. There will be higher Average Bandwidths and much higher Maximum Bandwidths that must be supported. In particular, the very challenging Maximum Bandwidths (which will soon exceed 1 Gbps) will likely become the dominant forcing function causing augmentation of the HFC network infrastructure time and time again over the next 2+ decades.

Predicting the future of bandwidth trends is always a challenging task, but it is particularly challenging when focusing on the Maximum Bandwidth (Tmax) trends across a 25-year horizon. These trends describe the Service Level Agreements (SLAs) that MSO Marketing departments are likely to roll out to subscribers on a year-by-year basis. Thus, these SLA trends are decided by a few select people within

each MSO- not by the more-predictable traffic patterns of subscribers. The decisions to increase SLA levels is usually made for either competitive reasons or due to customer demand. In the past, Nielson's Law [NI98] was used to predict these SLA trends, and based on past observations from Nielson, the law assumed that a simple 50% Compound Annual Growth Rate (CAGR) could be applied to both the Upstream and Downstream T<sub>max</sub> values to predict future values. However, in this paper, the authors have decided to use a "Slowed" version of Nielson's Law (shown in Figures 6 and 7 above). The idea of "slowing down" the future CAGRs within Nielson's Law has been proposed by many MSOs and vendors for quite some time, because there are challenges in describing reasonable applications and technologies that would lead to 100 Gbps SLAs being offered by 2030 (which is the value predicted by direct application of Nielson's Law in its traditional form). The particular "Slowed" CAGRs shown in Figures 6 and 7 lead to ~10 Gbps SLAs showing up in the ~2029 time-frame. It is impossible to prove which CAGR predictions for the future are correct, but the "Slowed" CAGRs of Figures 6 and 7 will be utilized within this paper.

To quantify the requirements and the performance of these higher Bandwidth Capacity HFC networks, the authors have found it useful to define several attributes related to the HFC network. These attributes are defined below.

**Number of Subscribers within a Service Group (N<sub>sub</sub>):** Within this paper, a Service Group group will be defined as neighboring subscribers who must share common Bandwidth Capacity on a coax or set of coaxes (assuming that the RF feeds to or from the set of coaxes are combined at a point in the HFC network). The number of subscribers within with a Service Group is sometimes given the label N<sub>sub</sub>.

**Downstream (DS) Maximum Throughput (T<sub>max</sub>):** This attribute is the maximum Downstream bandwidth permitted for a subscriber within the best (highest bandwidth) Service Level Agreement (SLA). It is measured in Gbps and typically increases as time progresses. These increases occur due to market pressures. Based on recent traffic engineering studies outlined in [UL19], this paper will make the following assumption about these increases (unless otherwise specified).

- It is assumed that the DS T<sub>max</sub> value in the year 2020 will be 1 Gbps = 1000 Mbps.
- It is assumed that the DS T<sub>max</sub> values will experience a Compound Annual Growth Rate (CAGR) of 25% in the 2020's, 15% in the 2030's, and 15% in the 2040's. This mimics a gradual reduction that some MSOs and vendors have been expecting (since this CAGR is less than the 50% CAGRs predicted in the past by the Nielson Law [NI98]).
- It is assumed that the DS T<sub>max</sub> values will be "rounded-off" to values representing the following values: 1 Gbps, 2 Gbps, 4 Gbps, 10 Gbps, 20 Gbps, 40 Gbps, and 80 Gbps.
- It is assumed that by 2029, DS T<sub>max</sub> will therefore be 10 Gbps.
- It is assumed that by 2044, DS T<sub>max</sub> will therefore be 80 Gbps.

**Upstream (US) Maximum Throughput (T<sub>max</sub>):** This attribute is the maximum Upstream bandwidth permitted for a subscriber within the best (highest bandwidth) Service Level Agreement (SLA). It is measured in Gbps and typically increases as time progresses. These increases occur due to market pressures. Based on recent traffic engineering studies outlined in [UL19], this paper will make the following assumption about these increases (unless otherwise specified).

- MSO marketing teams appear to be bifurcating into two different groups:
  - Those who plan to offer **Asymmetrical Services** where the US T<sub>max</sub> = 50% of the DS T<sub>max</sub>.
  - Those who plan to offer **Symmetrical Services** where the US T<sub>max</sub> = the DS T<sub>max</sub>.
- For Asymmetrical Services, it is assumed that:

- The US Tmax values will be “rounded-off” to values representing the following values: 0.5 Gbps, 1 Gbps, 2 Gbps, 5 Gbps, 10 Gbps, 20 Gbps, and 40 Gbps.
  - By 2029, US Tmax will therefore be 5 Gbps.
  - By 2044, DS Tmax will therefore be 40 Gbps.
- For Symmetrical Services, it is assumed that:
  - The US Tmax values will be “rounded-off” to values representing the following values: 1 Gbps, 2 Gbps, 4 Gbps, 10 Gbps, 20 Gbps, 40 Gbps, and 80 Gbps.
  - By 2029, US Tmax will therefore be 10 Gbps.
  - By 2044, DS Tmax will therefore be 80 Gbps.
- Offering Symmetrical Services will require more US Bandwidth Capacity than offering Asymmetrical Service.

**Downstream (DS) Average Throughput (Tavg):** This attribute is the average Downstream bandwidth consumed by a single typical subscriber within the busy-hour interval from (say) 9pm-10pm. It is measured in Mbps and typically increases as time progresses. These increases occur due to new Internet applications used by the subscriber. Based on recent traffic engineering studies outlined in [UL19], this paper will make the following assumption about these increases (unless otherwise specified).

- It is assumed that the DS Tavg value in the year 2020 will be 2.3 Mbps.
- It is assumed that the DS Tmax values will experience a Compound Annual Growth Rate (CAGR) of 39% in the 2020’s, 29% in the 2030’s, and 19% in the 2040’s. This mimics a gradual reduction that has been witnessed in the past decade.
- It is assumed that by 2029, DS Tavg will therefore be 44.6 Mbps.
- It is assumed that by 2044, DS Tavg will therefore be 1356.7 Mbps. (Note: At this time, the authors cannot describe the Internet applications that would require this much bandwidth capacity. It can be conjectured that it may be related to holographic multi-media experiences, but there is no proof of that at this time).

**Upstream (US) Average Throughput (Tavg):** This attribute is the average Upstream bandwidth generated by a single typical subscriber within the busy-hour interval from (say) 9pm-10pm. It is measured in Mbps and typically increases as time progresses. These increases occur due to new Internet applications used by the subscriber. Based on recent traffic engineering studies outlined in [UL19], this paper will make the following assumption about these increases (unless otherwise specified).

- It is assumed that the US Tavg value in the year 2020 will be 0.28 Mbps.
- It is assumed that the US Tmax values will experience a Compound Annual Growth Rate (CAGR) of 19% in the 2020’s, 19% in the 2030’s, and 19% in the 2040’s. This mimics the relatively flat CAGR level that has been witnessed in the past decade.
- It is assumed that by 2029, US Tavg will therefore be 1.3 Mbps.
- It is assumed that by 2044, US Tavg will therefore be 18.2 Mbps.

**Downstream (DS) Required BW Capacity per DS Service Group:** This attribute is the Downstream Bandwidth Capacity required to support the subscribers sharing bandwidth within a Downstream Service Group during the busy-hour interval from (say) 9pm-10pm.

- It is assumed that for a Service Group (a group of neighboring subscribers sharing common Bandwidth Capacity) with the actual number of attached subscribers given by the value Nsub, then the average bandwidth consumed by the Service Group will be given by  $N_{sub} \times T_{avg}$ .

- For a Service Group with a number of attached subscribers given by the value  $N_{sub}$  and with the maximum DS SLA bandwidth given by  $T_{max}$ , it is assumed that the following formula accurately describes the amount of High-Speed Data (HSD) Bandwidth Capacity required to keep high Quality of Experience (QoE) levels [CL14]:

$$\text{Required DS HSD Bandwidth Capacity} = N_{sub} * T_{avg} + 1.0 * T_{max} \quad [\text{Eq. 1}]$$

(Note: The second term provides head-room capacity for low-probability bandwidth bursts. Due to smaller  $N_{sub}$  values and much large  $T_{max}$  values in the future, this second term is likely to dominate traffic engineering in the future).

- To simplify the analysis, it is assumed that there will be a constant amount of DS Bandwidth Capacity dedicated to video for all years in this study. That amount will be equal to 336 MHz of spectrum dedicated to SC-QAM Video transport. This could correspond to 56 Annex B 6 MHz channels or 42 Annex A 8 MHz channels. It is assumed that improved video compression techniques will permit more and more video content to be propagated over this spectrum as time progresses.

**Upstream (US) Required BW Capacity per US Service Group:** This attribute is the Upstream Bandwidth Capacity required to support the subscribers sharing bandwidth within a Upstream Service Group during the busy-hour interval from (say) 9pm-10pm.

- It is assumed that for a Service Group (a group of neighboring subscribers sharing common Bandwidth Capacity) with a number of attached subscribers given by the value  $N_{sub}$ , the average bandwidth consumed by the Service Group will be given by  $N_{sub} * T_{avg}$ .
- For a Service Group with a number of attached subscribers given by the value  $N_{sub}$  and with the maximum US SLA bandwidth given by  $T_{max}$ , it is assumed that the following formula accurately describes the amount of High-Speed Data (HSD) Bandwidth Capacity required to keep high Quality of Experience (QoE) levels [CL14]:

$$\text{Required US HSD Bandwidth Capacity} = N_{sub} * T_{avg} + 1.0 * T_{max} \quad [\text{Eq. 2}]$$

(Note: The second term provides head-room capacity for low-probability bandwidth bursts. Due to smaller  $N_{sub}$  values and much large  $T_{max}$  values in the future, this second term is likely to dominate traffic engineering in the future).

**Downstream (DS) HSD Utilization:** This attribute is the Downstream HSD Utilization level, which describes the percentage of the total Required Bandwidth Capacity that is utilized by the Average Bandwidth levels. The formula is given by  $\text{Utilization} = (N_{sub} * T_{avg}) / (N_{sub} * T_{avg} + 1.0 * T_{max})$ .

**Upstream (US) HSD Utilization:** This attribute is the Upstream HSD Utilization level, which describes the percentage of the total Required Bandwidth Capacity that is utilized by the Average Bandwidth levels. The formula is given by  $\text{Utilization} = (N_{sub} * T_{avg}) / (N_{sub} * T_{avg} + 1.0 * T_{max})$ .

**Fiber Depth:** This attribute defines how deep the fiber has been routed into the HFC Network. It is usually described using the Node+X notation, where X specifies the maximum number of serialized amplifiers in a cascade within the coaxial portion of the HFC network. A Node with no amplifiers south of it exists in a Node+0 environment, and a Node with at least one amplifier south of it exists in a Node+Non-Zero environment. A longer coaxial run with X serialized amplifiers usually implies a shorter fiber run, and vice versa. If one assumes that a typical amplifier supports a coaxial length of (say) 1000

feet, then a typical Node+X system might have a maximum aggregated distance of  $\sim(1000 \text{ feet}) \cdot (X+1)$  between the Fiber Node and the most distant subscriber on that coaxial run. The number of subscribers connected to a particular Fiber Node is closely related to the value of X. With each node-split (which reduces the number of subscribers per Fiber Node by roughly a factor of two), there is usually a corresponding reduction in the value of X (since less distance and less serialized amplifiers are usually needed to reach the smaller number of subscriber). Within this paper, we assume a typical MSO has a  $\sim 50\%$  take-rate on their services. While different MSOs assume different numbers for these relationships, this paper will assume the following relationships exists between Fiber Depth and the number of Households Passed (HHP) and the number of subscribers:

**Table 1 – Relationships between Fiber Depth, # of HHP, & # of Subs**

Fiber Depth	# of HHP	# of Subs (Nsub)
Node+0	120	60
Node+1	240	120
Node+2	480	240
Node+3	800	400
Node+4	1200	600
Node+5	1600	800
Node+6	2000	1000

MSOs seem to have bifurcated into two groups. One group has a strong affinity to move to Node+0 as quickly as possible (in an effort to quickly move towards the expected end-game architecture with fiber going directly to the home). The other group has a strong affinity to stay away from Node+0 (in an effort to avoid the increased costs of fiber pulls and numerous Fiber Node deployments that are associated with Node+0 or FTTH). For example, one anonymous MSO has indicated that 71% of the Capex+Opex cost of moving from Node+3 to Node+0 is in the final Node+1 to Node+0 transition, and their annual budgets will not permit them to tackle that large expense for many years to come. As a result, they will opt to utilize Node+Non-Zero architectures (and wider spectral widths yielding higher BW Capacities) for as long as possible. This paper will study both the Node+0 and Node+Non-Zero approaches.

**Top of US Bandwidth Range:** This attribute (also known as the “split”) defines how much of the coaxial spectrum is utilized for US transmissions. It is assumed that the US signals will occupy the spectrum from (say) 5 MHz to the Top of the US Bandwidth Range. Higher required US Bandwidth Capacities will typically require higher US Bandwidth Ranges. Depending on the required frequencies, higher US Bandwidth Ranges may require the development of new chipsets (ex: FPGAs, Hybrids) and new actives (ex: Nodes, Amplifiers, CMs), and new passives (ex: Taps).

**Top of Downstream (DS) Bandwidth Range:** This attribute defines how much of the coaxial spectrum is utilized for DS transmissions. It is assumed that the DS signals will occupy the spectrum from the Bottom of the DS Bandwidth Range to the Top of the DS Bandwidth Range. Higher required DS Bandwidth Capacities will typically require higher DS Bandwidth Ranges. The Bottom of the DS Bandwidth Range is determined by the particular technology used to combine US & DS signals (as will be described below). Depending on the required frequencies, higher DS Bandwidth Ranges may require the development of new chipsets (ex: FPGAs, Hybrids) and new actives (ex: Nodes, Amplifiers, CMs), and new passives (ex: Taps). It is assumed that the Outside Plant (OSP) network will be physically and electrically isolated from subscriber in-home network using portal-based gateways, so augmentation of most in-home passives should not be required.



**Downstream (DS) Equipment Housing Bandwidth:** This attribute is closely related to the Top of the DS Bandwidth Range. However, this attribute recognizes two important facts:

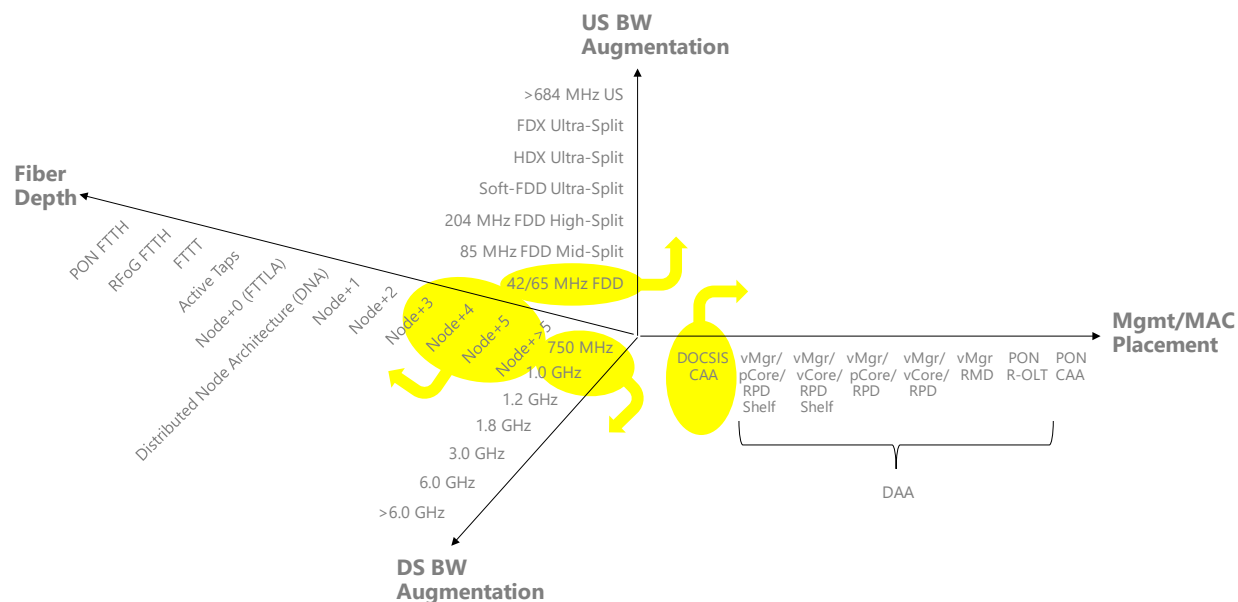
- 1) Outside Plant Equipment (such as Nodes and Amplifiers) are typically expected to be deployed into the field and not replaced for many years (ex: 10-15 years or more).
- 2) If a piece of Outside Plant Equipment is to be deployed in a given year, it should probably include a housing design that can support the expected frequencies of its future life-span.

Thus, the DS Housing Bandwidth represents an attempt to group together several generations of frequency improvements that might be added (via modular plug-ins) into a piece of Outside Plant Equipment, and then it tries to ensure that at least the Housing for that equipment can support those multiple generations and all of the anticipated future frequencies prior to the Housing being replaced. The highest frequency that a particular generation of Housing is expected to support is the DS Housing Bandwidth (in MHz).

**Upstream:Downstream (US:DS) Frequency Band Re-Use Ratio:** A very important attribute is called the US:DS Frequency Band Re-Use. This attribute is measured as a frequency range and uses units of MHz. If a particular solution permits the US and DS frequency spectra to dynamically overlap during normal operation by (say) Z MHz, then that solution is said to have US:DS Frequency Band Re-Use of Z MHz, and Z would be a positive number. If the US and DS frequencies do not overlap but are contiguous, then the Re-Use value Z would be zero. If the US and DS frequencies do not overlap and actually have a frequency domain guard-band between them (for example, for diplex filter rolloffs), then the Re-Use value Z would be negative. In this latter case, a region of spectrum is excluded from being used and represents an undesirable “spectral penalty.” Another related attribute is the US:DS Frequency Band Re-Use Ratio, which is defined to be the Re-Use value Z divided by the top of the useable spectral range. Positive US:DS Frequency Band Re-use Ratios with large absolute values are good because they indicate a lot of spectral overlap. Negative Re-use Ratios with large absolute values tend to be undesirable, because they indicate a lot of wasted spectrum.

## 2. Overview of Some Key Technology Candidates for Future Upstream/Downstream Bandwidth Augmentation

In the Introduction section above, several candidate technologies for the future were listed. Those technologies were sub-divided into several different categories depending on their focus. This has led the authors to visualize the evolution of the future within the four axes of the “4-dimensional” coordinate system shown in Figure 8. This visualization attempts to illustrate that changes will be taking place in at least four different areas, including changes in MAC/Management Placement, Upstream Bandwidth Augmentation, Downstream Bandwidth Augmentation, and Fiber Depth.



**Figure 8 – Changes in Four Different Areas to Support Future Bandwidth Growth**

The yellow regions in Figure 8 indicate roughly where many MSOs are currently at in this evolutionary process on all four axes. The yellow arrows indicate likely directions that the MSOs may choose to take as that move into the future and add more Bandwidth Capacity to their networks. It is likely that different MSOs will choose different hops on different axes at different times. Selecting when to take the “right hop on the right axis at the right time” requires a careful analysis of Bandwidth Capacity requirements (which will be done below).

Over the next 2+ decades, MSOs will be focused on determining ways to augment both Upstream and Downstream Bandwidth Capacities within their HFC networks (prior to transitioning to FTTH PON systems). There are many different technology improvements that can potentially offer this augmentation. MSOs will undoubtedly need to determine which technology to utilize, and they will also need to determine when to transition between different technologies to maximize their Bandwidth Capacity while minimizing their costs. This section will give a brief overview of some of the more promising technology improvements that may be considered in the future. At the end of each sub-section, an assessment of the “Likelihood of Success” for each technology is given. A more detailed analysis of many of these technology improvements (along with suggestions on how to improve their performance) is contained in [AL19].


## 2.1. Traditional-Full-Duplex DOCSIS (Traditional-FDX) for Upstream Augmentation

This well-known technology has been discussed by the Cable Industry for years, and the specification changes associated with Traditional-FDX recently were recently moved into the newly-created DOCSIS 4.0 specification. This exciting technology proposes to have Downstream and Upstream transmissions


occurring in the same frequency band at the same time. In the specification, the overlapping frequency bands can be in any of the following ranges: 108-204 MHz, 108-300 MHz, 108-396 MHz, 108-492 MHz, or 108-684 MHz.

The Traditional-FDX capability offers a powerful and unique benefit that permits Upstream spectrum expansions to occur without causing commensurate reductions in Downstream spectrum widths. One of the key technology enablers leading to Traditional-FDX is the Echo Canceller functionality that is required to cancel high-power, reflected noise in both the Fiber Node's Upstream and the Cable Modem's (CM's) Downstream. Based on experiments and field trials in which the authors have been involved, it is clear that Traditional-FDX should work quite well in Node+0 environments. The basic concept is illustrated in Figure 9.

• **85 MHz FDD**



• **396 MHz Traditional-FDX**




**Figure 9 – Spectrum for 85 MHz FDD Mid-Split & 396 MHz Traditional-FDX (w/ First-Order Guestimates on Bandwidth Capacities)**


Unfortunately, despite many design approaches and analyses, the ability of Traditional-FDX systems to perform in Node+Non-Zero environments (where there exists at least one amplifier south of the Fiber Node) is still questionable and under study. The key component enabling Traditional-FDX operation in a Node+Non-Zero environment is the FDX Amplifier. While FDX Amplifiers of various forms have been proposed and are possible to construct, all have resulted in Time Division Duplex (TDD)/Frequency Division Duplex (FDD) performance levels- not Traditional-FDX performance levels. For example, recent studies on noise amplification and Interference Group formation have shown that a particular issue may develop. In particular, the use of Echo Cancellation and Amplification techniques in FDX Amplifiers may suffer from an Interference Group Elongation problem (see Figure 10) that effectively makes each RF Leg a single large Interference Group operating in a TDD/FDD fashion. While still under study, this Interference Group Elongation problem may preclude Traditional-FDX solutions from operating well in any Node+Non-Zero environments, and it may require MSOs who plan to continue using Node+Non-Zero HFC plants to consider using some of the alternative techniques described in the sub-sections below.

Traditional-FDX solutions may therefore only be useable by MSOs who expend the effort and money to convert their current HFC networks into Node+0 networks. This may create a bifurcation of MSOs into two camps in the future; there may be the Node+0 camp using Traditional-FDX and the Node+Non-Zero camp using other alternative Upstream Bandwidth Augmentation technologies listed below. This paper will study proposals for both camps.

### Ideal FDX Amp Operation



### Observed FDX Amp Operation (to date)



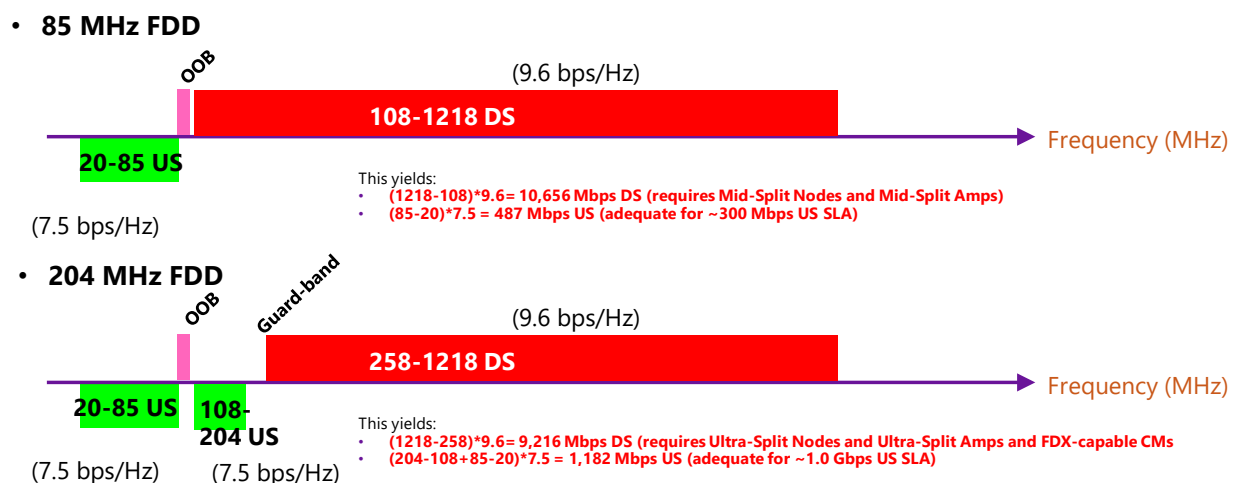
**Figure 10 – Interference Group Elongation Problem That May Occur With FDX Amplifiers (Not A Problem For FDD High-Split or FDD Ultra-Split Approaches)**

For the Node+0 MSO camp planning to use Traditional-FDX, several changes will be required. If the Interference Group Elongation problem remains, then splits to a Node+0 level will likely need to be made. In addition, CCAP Cores, Fiber Nodes, and CMs will all require changes; there will be changes in both the PHY and MAC layers of DOCSIS. The complexities of FDX Echo Cancellation functionality may also force the initial FDX offerings to occur in only RPHY environments and using 1 DSSG x 1 USSG Fiber Nodes, so the impact of these limitations must be analyzed when considering FDX for deployments in the early 2020's.

Likelihood of Success: High probability of success; much development is still required.

## 2.2. 204 MHz Frequency Division Duplex (FDD) High-Splits for Upstream Augmentation

This technology improvement can be utilized by the Node+Non-Zero camp of MSOs. It permits the MSOs to expand their HFC plant's Upstream spectrum from the current 42 or 55 or 65 MHz or 85 MHz spectral widths to a much larger 204 MHz spectral width- resulting in more Upstream Bandwidth Capacity that can be offered to subscribers. This technique proposes to keep things simple by continuing to use Frequency Division Duplex technologies that separate Upstream spectrum from Downstream spectrum- usually with a diplexer filter guard-band in between the disjoint Upstream and Downstream frequency ranges. This is illustrated in Figure 11.



**Figure 11 – Spectrum for 85 MHz FDD Mid-Split & 204 MHz FDD High-Split (w/ First-Order Guestimates on Bandwidth Capacities)**

This transition to 204 MHz FDD High-Split operation typically requires changes to both the existing Nodes and the existing Amplifiers and some existing CMs (only in the high-end subscriber homes) on the HFC plant. Sometimes, plug-in modules may be made available for new duplex filters that can be utilized to simplify this upgrade path. The 204 MHz duplex filters would typically have a roll-off that permits the Downstream spectrum to pass at frequencies higher than ~258 MHz, so the resulting guard-band creates a duplex filter “spectral penalty” of  $\sim(204-258) = -54$  MHz; i.e.- ~54 MHz of spectrum is deemed to be unusable when this technology is deployed without any assistance from Guardband Reduction technologies (using Echo Cancellers or other techniques). As can be seen from the Figure, the move to a 204 MHz FDD High-Split solution also “steals” some of the DS Bandwidth Capacity (in red) and “donates” it to the US Bandwidth Capacity (in green). These effects may or may not represent an issue in the future, as will be described below.

One of the benefits of this FDD approach is that it works well in a Node+X ( $X \geq 0$ ) environment without any complications to the Amplifiers (other than the use of a higher split). It does not require the use of Echo Cancellation techniques in the Amplifiers (as needed for Full Duplex DOCSIS solutions), and

One of the side issues that must be dealt with when 204 MHz FDD High-Split is utilized is the passing of Downstream Out-Of-Band (OOB) signals to some existing Set-Top Boxes. To support this OOB capability vendors are exploring the use of techniques such as the use of high-Q filter modules that can be added to existing or future Amplifiers that would permit these Downstream OOB signals to be passed through the Amplifier even though the Downstream OOB signals are in the Upstream portion of the spectrum.

Likelihood of Success: Very high probability of success; since this is actually DOCSIS 3.1, it is already available.

### 2.3. 204-684 MHz Frequency Division Duplex (FDD) Ultra-Splits for Upstream Augmentation

This technology improvement can also be utilized by the Node+Non-Zero camp of MSOs, and it provides them with similar Upstream Bandwidth augmentation benefits that may be experienced by the Node+0

camp who uses Traditional-FDX- however this FDD Ultra-Split technology also has a commensurate Downstream Bandwidth Capacity degradation that is experienced due to a “zero-sum game” effect).

The technology is very similar to the 204 MHz FDD High Split approach defined above. However, it recognizes that one could push the Upstream Split levels to frequencies well beyond the 204 MHz limit of the original DOCSIS 3.1 specification. In particular, with the expected arrival of FDX-capable Cable Modem (CM) chipsets in the next year or so, it seems clear that one could hypothesize the deployment of FDD Ultra-Splits at any of the Upstream frequencies that are supported by the coming FDX-capable CM chipsets. The resulting FDX-capable CMs would be used in a simple “non-FDX” operating mode, whereby Sounding and fast Resource Block Assignments and other features would not be needed for this simple FDD Ultra-Split operation. The resultant FDD Ultra-Split frequencies may therefore include:


- 204 MHz
- 300 MHz
- 396 MHz
- 492 MHz
- 684 MHz.

The FDD Ultra-Split technique proposes continued use of simple Frequency Division Duplex, separating Upstream spectrum from Downstream spectrum (usually with a diplexer-based frequency guard-band in between) and avoiding the forementioned FDX Interference Group Elongation problem. Unfortunately, the required guard-band with FDD Ultra-Split tends to grow in size with higher Split frequencies (ex: a 396 MHz Ultra-Split might require a 69 MHz guard-band ranging from 396 MHz to 465 MHz). As a result, there are studies under way that could potentially employ Echo Cancellation techniques to reduce or eliminate the need for this guard-band range. Cost and power issues are still being worked out in those studies.


As one might expect, the initiation of this FDD Ultra-Split operation typically requires changes to both the existing Fiber Nodes and the existing Amplifiers on the HFC plant. CMs will also need to be updated within high-end subscriber homes. For some MSOs, it also may require the addition of high-Q filters to support the passage of the Downstream OOB signals within the extended Upstream frequency band.

Like the FDD High-Split solution, the FDD Ultra-Split solution also “steals” Bandwidth Capacity from the DS; but it steals even more DS Bandwidth Capacity (in red) and “donates” it to the US Bandwidth Capacity (in green), as shown in Figure 12 and Table 2. As illustrated in Table 2, there is a clear zero-sum game trade-off; large increases in US Bandwidth Capacity will typically lead to large decreases in DS Bandwidth Capacity. Fixes for this problem may be available if the Cable Industry ultimately decides to support Extended Spectrum DOCSIS concepts (described below) for Downstream augmentation, using them in conjunction with the FDD Ultra-Split solution for the Upstream augmentation.

• 85 MHz FDD



• 396 MHz FDD



**Figure 12 – Spectrum for 85 MHz FDD Mid-Split & 396 MHz FDD Ultra-Split (w/ First-Order Guestimates on Bandwidth Capacities)**

**Table 2 – US & DS BW Capacities for Various FDD Ultra-Split Frequency Bands Illustrating the Zero-Sum Game**

Top of Ultra-Split FDD US Frequency Band (MHz)	Bottom of US Frequency Band (MHz)	Top of DS Frequency Band (MHz)	Bottom of DS Frequency Band (MHz)	US BW Capacity w/ 7.5 bps/Hz (Mbps)	DS BW Capacity w/ 9.6 bps/Hz (Mbps)	Unuseable Guard-Band #1 from 85-108 MHz (MHz)	Unuseable Guard-Band #2 at US:DS Split (MHz)	Total Unuseable Guard-Band (MHz)	US:DS Frequency Band Re-Use Ratio (%)
204	20	1218	258	1208	9216	23	54	77	-6%
300	20	1218	352	1928	8314	23	52	75	-6%
396	20	1218	465	2648	7229	23	69	92	-8%
492	20	1218	578	3368	6144	23	86	109	-9%
588	20	1218	690	4088	5069	23	102	125	-10%
684	20	1218	803	4808	3984	23	119	142	-12%

Because it does not require FDX Amplifiers with Echo Cancellation, FDD Ultra-Split operation does not suffer from the Interference Group Elongation issue. Thus, FDD Ultra-Split operation permits MSOs to continue to work within Node+Non-Zero environments, saving the MSOs the cost of Node+0 Node-Splits and Fiber Deep optical runs . The implications of these effects will be analyzed in more detail below.

Likelihood of Success: High probability of success; development is still required, but it is low risk development.

## 2.4. Static Soft-Full-Duplex DOCSIS (Static Soft-FDX) for Upstream and Downstream Bandwidth Augmentation

If a Node or Amplifier is designed to permit multiple Upstream Splits (as with Ultra-Split) or if a Node or Amplifier is designed to permit multiple Downstream Spectra (which may occur if/when Extended Spectrum DOCSIS is utilized in the future), then one must consider how those Split changes and/or Spectra changes will be initiated.

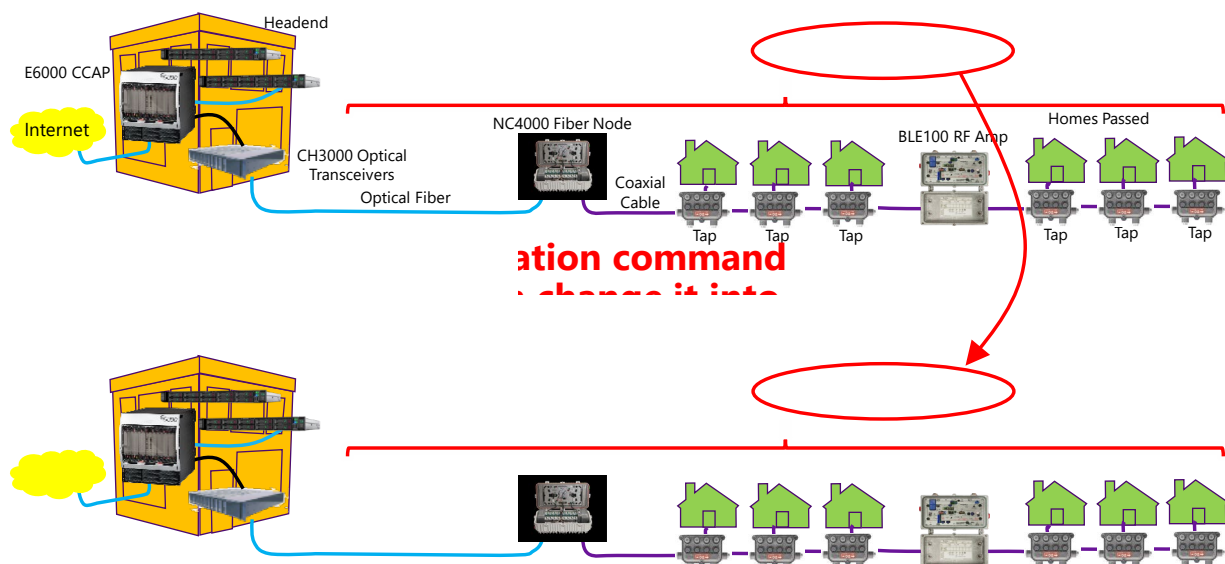
The simplest technique (which was oftentimes used in the past) was to simply require the physical swap-out of plug-in modules to modify the useable frequency band. This approach may be permissible in the future. However, this approach requires a truck roll to the Node or Amplifier being modified, which may be undesirable.

Soft-FDX is an improvement that could permit the MSO to use software configurability from the Head-end to change the Split or Spectra within the distant Node or Amplifier. The Soft-FDX technology requires the addition of a processor and (usually) a CM or receiver to each Amplifier, so the added cost on Amplifiers must be considered by MSOs before deploying this technology.

There are two variants of Soft-FDX; namely Static Soft-FDX and Dynamic Soft-FDX. The Static version is described in this section, and the Dynamic version is described in the next section. Static Soft-FDX is the variant that permits the FDD frequency range splits to be dynamically changed with a Command Line Interface configuration in the head-end. Thus, changes are assumed to occur somewhat infrequently, and the transitions between FDD frequency ranges do not necessarily need to occur very quickly. With Static Soft-FDX, the Upstream Bandwidth and Downstream Bandwidth cannot share the spectrum once the FDD frequency range has been set on a given RF Leg. However, it is possible (although improbable) that an MSO could set the split frequency to be different on different RF Legs, which could lead to some level of FDX operation at the Node Level (with overlapping US & DS operation at a single frequency).


An example of an application of Static Soft-FDX is illustrated in Figure 13, and an example of spectrum changes on a yearly basis is shown in Figure 14 (with time on the y-axis and frequency splits on the x-axis).

Likelihood of Success: High probability of success; development is still required.



**Figure 13 – Static Soft-FDX**





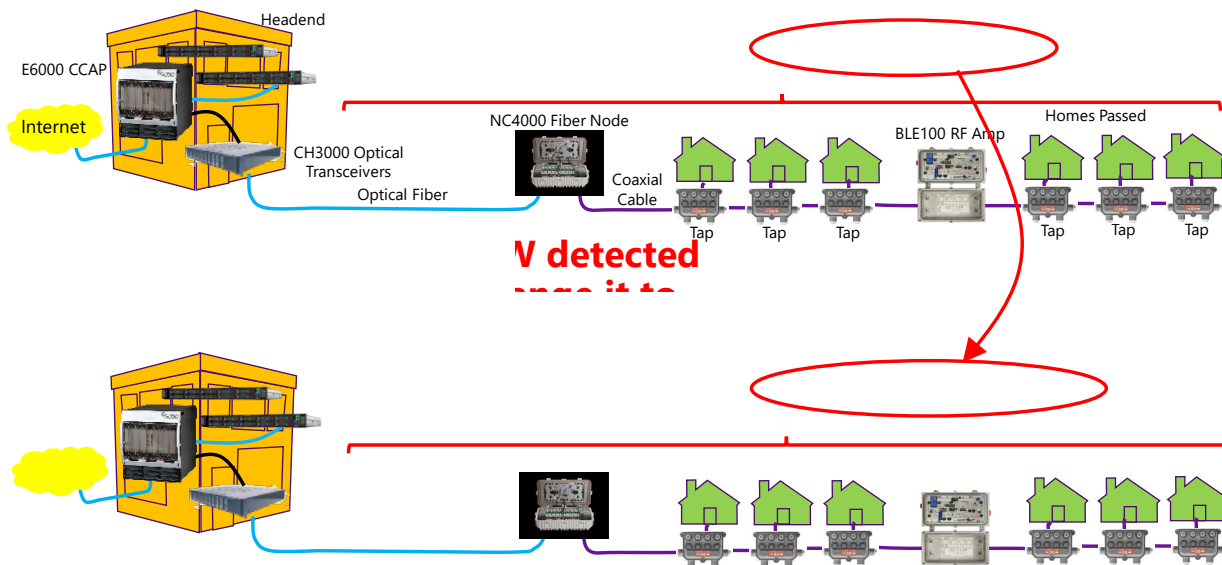
**Figure 14 – Examples of Yearly Spectrum Changes using Static Soft-FDX**

## 2.5. Dynamic Soft-Frequency-Division Duplex (Dynamic Soft-FDX) for Upstream and Downstream Bandwidth Augmentation


Dynamic Soft-FDX is an interesting solution that is half-way between Traditional-FDX and Static Soft-FDX. Like Traditional-FDX, it requires that the changes between different FDD frequency ranges occur frequently and rapidly. The FDD frequency changes within Dynamic Soft-FDX are not typically initiated by Command Line Interface commands (as they are with Static Soft-FDX), but they are instead initiated by the DOCSIS Media Access Control sub-system monitoring the Upstream and Downstream Bandwidth usage in real-time (as is done in Traditional-FDX). Thus, Dynamic Soft-FDX does not actually “steal” Bandwidth Capacity from the Downstream and permanently give it to the Upstream- instead it temporarily “borrows” Bandwidth Capacity from the Downstream and temporarily loans it to the Upstream before quickly returning it back to the Downstream (assuming Upstream bursts are less frequent than Downstream bursts).

Like Static Soft-FDX, Dynamic Soft-FDX relies on simple FDD operations to eliminate the need for the Echo Cancellation in Amplifiers. Thus, it will still work well in Node+Non-Zero environments.

A potential algorithm for Dynamic Soft-FDX might require continual monitoring of Upstream and Downstream bandwidth flows and provide for early recognition of Upstream Bandwidth bursts. Upon recognizing the start of an Upstream Bandwidth Burst, the system could change the frequency ranges on a second-by-second basis to temporarily provide more Upstream Bandwidth Capacity (and less Downstream Bandwidth Capacity) whenever required. An example of an application of Dynamic Soft-FDX is illustrated in Figure 15, and an example of spectrum changes on a second-by-second basis is shown in Figure 16 (with time on the y-axis and frequency splits on the x-axis).



**Figure 15 – Dynamic Soft-FDX**

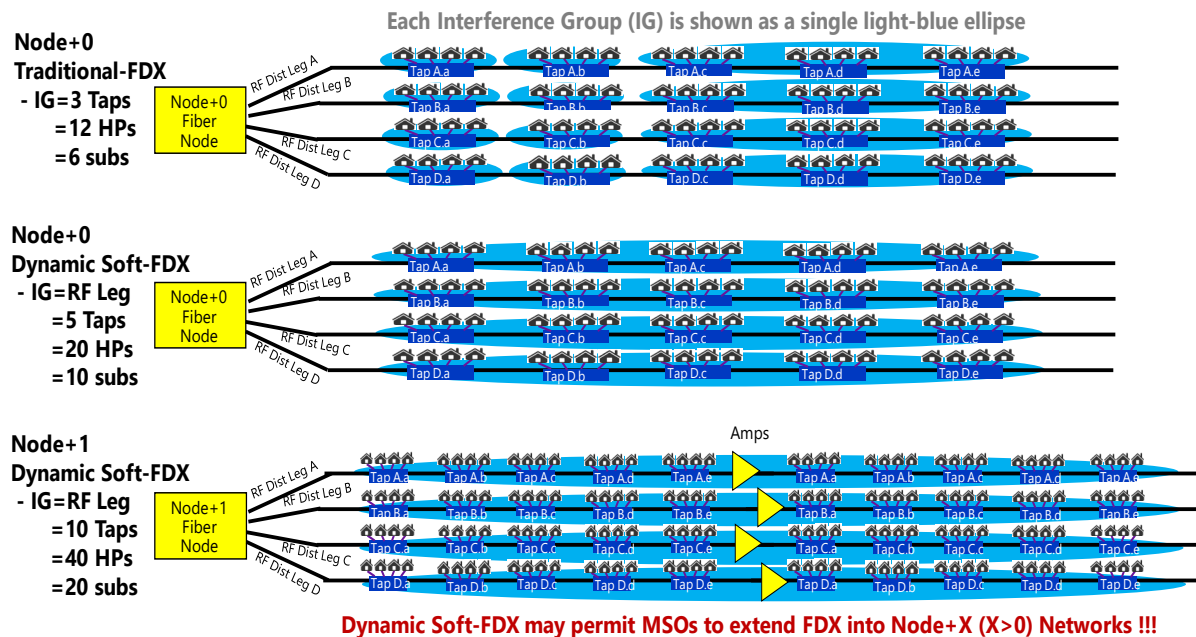


**Figure 16 – Examples of Second-by-Second Spectrum Changes using Dynamic Soft-FDX**

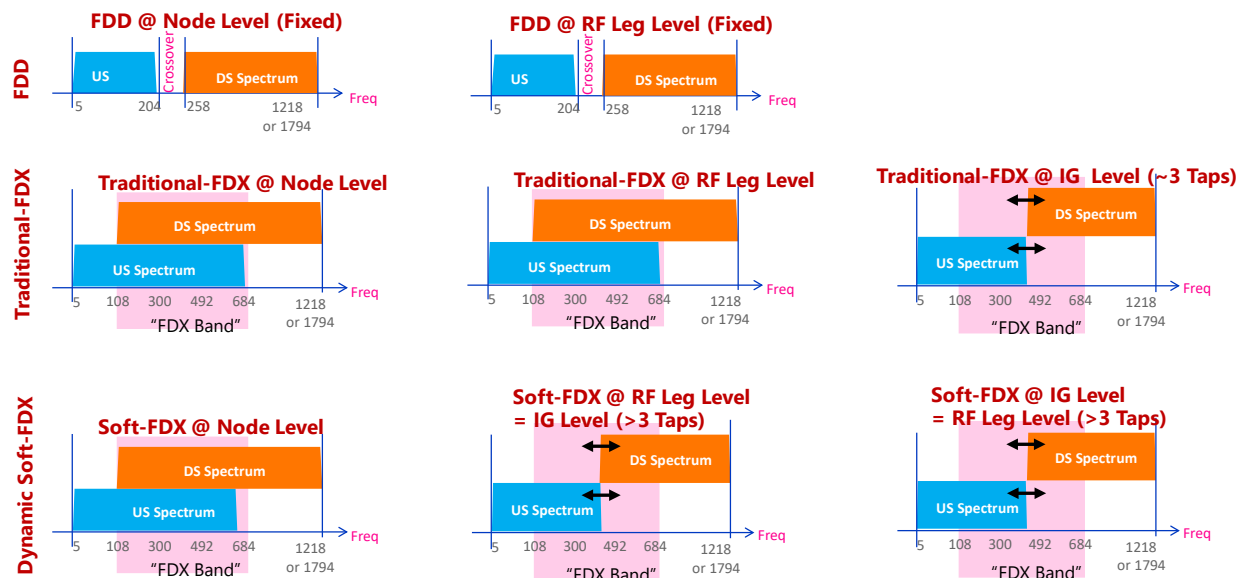
These rapid frequency range changes are quite similar to solutions that have been proposed for Traditional-FDX operation. In fact, Traditional-FDX must perform these functions within each Interference Group of a Service Group. In fact, the operation within a Traditional-FDX Interference Group could be identical to that which is shown in Figure 16.

However, the Dynamic Soft-FDX frequency range modifications would likely be managed at an RF Distribution Leg level (or perhaps at a Node level) instead of at the Interference Group level (as done for Traditional-FDX). Another way to look at this Dynamic Soft-FDX solution is to consider it to be a

simplified Traditional-FDX solution where the Traditional-FDX Interference Group has been enlarged to cover an entire RF Distribution Leg (or perhaps the entire Service Group within a Node) instead of covering small subsets of FDX CMs on neighboring Taps that cause noise to one another. Thus, Dynamic Soft-FDX is indeed a form of FDX, which is apparent when one compares the Interference Groups in Figure 17 and Figure 18.



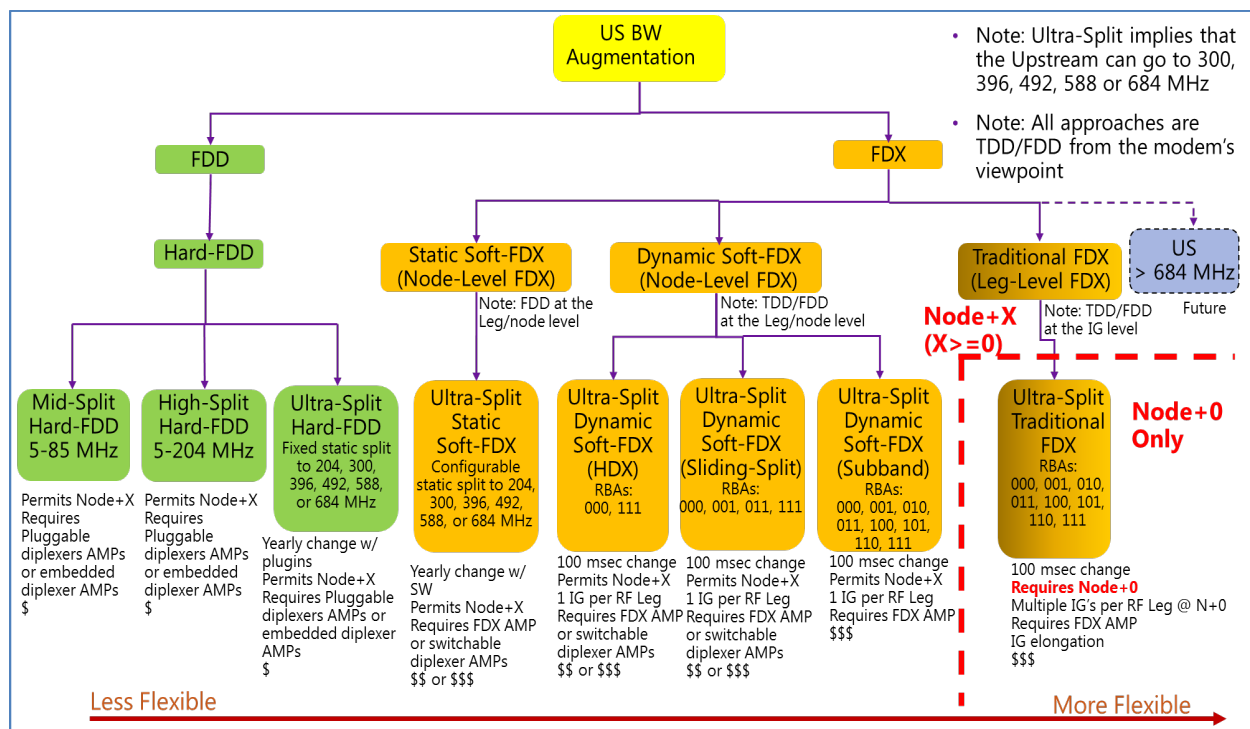
**Figure 17 – Interference Groups in Traditional-FDX and RF Leg-based Dynamic Soft-FDX (for Node+0 and Node+1 Systems)**



**Figure 18 – FDX Operation vs Sliding FDD Operation at Various Levels of FDD Systems, Traditional-FDX Systems, & Dynamic Soft-FDX Systems**

Dynamic Soft-FDX (operating with Interference Groups sized to be entire RF Legs) would still require Sounding algorithms, because each FDX CM must be mapped into the correct Interference Group (RF Leg). However, Sounding algorithms may be able to be simplified when operated at the RF Leg level. In addition, Dynamic Soft-FDX requires that the FDX Resource Block Assignments (RBAs) be changed on a regular basis (like Traditional-FDX), and the Fiber Nodes and Amplifiers and CMs must all support those RBA changes in a coordinated and synchronized fashion. These changes are ultimately managed by the DOCSIS Media Access Control sub-systems that continually monitor the dynamic changes in Upstream and Downstream traffic flows. However, the Dynamic Soft-FDX architecture has embraced large FDX Interference Groups (covering an entire RF Leg) as an acceptable construct. This permits Dynamic Soft-FDX to work well in Node+Non-Zero environments using simpler FDD-based Amplifier designs that do not require Echo Cancellation. The Amplifiers may just require low-cost switchable Filter Banks that could be added to existing Amplifiers already deployed in the field.

In the end, Dynamic Soft-FDX and Traditional-FDX are merely different variations of FDX solutions with different views on acceptable Interference Group sizes. Figure 19 illustrates a taxonomy of Upstream Augmentation solutions, and within that Figure, it can be seen that there are actually a continuum of FDX solutions that use different Interference Group (IG) sizes and different subsets of Resource Block Assignments (RBAs).



**Figure 19 – Taxonomy of Upstream Augmentation Solutions with a Continuum of FDX Solutions (including Tradition-FDX and Dynamic Soft-FDX and Static Soft-FDX)**

Dynamic Soft-FDX can (in some ways) be viewed as a “simplified version” of Traditional-FDX, with slightly larger Interference Groups and more subscribers sharing the FDD/TDD bandwidth. (Note: The same FDD/TDD operation exists inside of Interference Groups of Traditional-FDX). Therefore, Dynamic Soft-FDX should experience slightly lower QoE performance levels than FDX. However, Traffic Engineering studies performed by the authors to discover the magnitude of the QoE impacts resulting from changing the FDX Interference Group sizes from a small subset of CMs to a full RF Distribution leg (or even a full Service Group) have shown promising results.

As an example, the authors ran a convolution-based simulation (using predicted probability density functions of futuristic bandwidth usage based on high-sample-rate data collections from present-day, real-world HFC plants extended into the future). Many scenarios were simulated, but we will focus on just two of those scenarios to illustrate the point. The first is a Node+0 Traditional-FDX environment with 64 subscribers per Node and (worst-case) 8 subscribers per Interference Group. The second is a Node+1 Dynamic Soft-FDX environment with 128 subscribers per Node and 32 subscribers per Interference Group (which represents an entire RF Leg).

In both scenarios, we assume futuristic ~2024 numbers of DS T<sub>max</sub> = 2000 Mbps and US T<sub>max</sub> = 2000 Mbps and DS T<sub>avg</sub> = 13.9 Mbps and US T<sub>avg</sub> = 0.57 Mbps. To create a challenging environment, the “Legacy DOCSIS DS Band” was assumed to carry sixteen 6 MHz 256QAM SC-QAM DOCSIS DS channels plus one 192 MHz OFDM channel (~2016 Mbps in total). The “Legacy DOCSIS US Band” was assumed to carry four 6.4 MHz 64QAM SC-QAM DOCSIS DS channels plus one 43-MHz OFDMA channel operating from 42 to 85 MHz (~422 Mbps on total). The simulation determined the amount of Bandwidth Capacity (in Mbps) required for the shared frequency range to ensure that the Quality of Experience (QoE) levels in 2024 would remain the same as they are today by ensuring that the probability of a bandwidth burst exceeding the capacity of the HFC plant is the same as it is today.

In those simulation results, the total shared FDX BW Capacity required for the 2024 time-frame using the Node+0 Traditional-FDX design (with 8-subscriber Interference Groups) was found to be 2015 Mbps, which (assuming a useable spectral efficiency of 7.5 bps/Hz) required an FDX Band running from ~108-377 MHz.

The total shared FDX BW Capacity required for the 2024 time-frame using the Node+1 Dynamic Soft-FDX design (with 32-subscriber RF Legs = 32-subscriber Interference Groups) was found to be 2164 Mbps, which (assuming a useable spectral efficiency of 7.5 bps/Hz) required an FDX Band running from ~108-397 MHz. (Note: The slight increase in Required Bandwidth Capacity is primarily due to the need to support more subscribers within the Node+1 Service Group).

Thus, both of these solutions come close to fitting nicely within the BW Capacity of a 108-396 MHz FDX Band system. This illustrates that Dynamic Soft-FDX works about as efficiently as Traditional-FDX.

Likelihood of Success: High probability of success; much development is still required.

## **2.6. Extended Spectrum DOCSIS (ESD) for Upstream and Downstream Bandwidth Augmentation**

Extended Spectrum DOCSIS (ESD) is a proposal originally made by the authors at the CableLabs 2015 Summer Conference and again at the NCTA Shows in 2015 [CL15] and 2016 [CL16]. The simple idea proposes to extend the spectrum of DOCSIS 3.1 to Upstream ranges higher than 684 MHz and to Downstream ranges higher than 1218 MHz.

The keys to a successful implementation of this idea include:

- 1) Maintain a Total Composite Power output level from devices (Nodes, Amplifiers, CMs) at levels similar to those of today
- 2) Utilize the same tilted power spectral density (as utilized today) for all SC-QAM signals (ex: Video & pre-DOCSIS 3.1 signals) that are operating at the lower frequencies in the spectrum (which should preclude the need for any re-spacing of Amplifiers in the HFC plant)
- 3) Utilize DOCSIS 3.1 Orthogonal Frequency Division Multiplexing (OFDM) and Low Density Parity Check (LDPC) Forward Error Correction (FEC) for all signals in higher frequency ranges
- 4) Utilize a different power spectral density for the higher frequencies by applying a flat power spectral density in all of the OFDM regions at the high frequency range of the spectrum
- 5) Rely on the benefits of OFDM bit-loading to match the Quadrature Amplitude Modulation (QAM) levels to the particular Signal-to-Noise (SNR) ratios that will undoubtedly decrease at higher frequencies (due to increased attenuation levels at higher frequencies)
- 6) For legacy Set-top Boxes, use high-Q filters or other techniques to pass OOB Downstream video signals within the Upstream frequency range
- 7) Use a 2-port Gateway-style of CM at the portal into the home to electrically isolate the Outside Plant from the in-home network for any home receiving and processing the ESD signal (which is identical to a requirement for Traditional-FDX operation)

There are several variants of Upstream ESD spectral widths and Downstream ESD spectral widths that can be envisioned for the future. Table 3 illustrates some of these proposals. Spectral widths requiring greater than ~8 GHz will require a move to FTTT or FTTH systems. It is also clear that wider spectra usually will display a lower Spectral Efficiency (due to increased attenuation and lowered SNRs at the higher frequencies). The 396 MHz Upstream and 1218 MHz Downstream rows (shown with red numbers

in the Table) are the likely configurations needed to compete with the 8.6 Gbps DS x 2.4 Gbps US of the near-term low-cost Asymmetric XGS-PON competitor shown in Figure 1. An augmentation to the 1218 MHz Upstream ESD row (shown with maroon numbers in the Table) is a likely configuration change needed to compete with the 8.6 Gbps DS x 8.6 Gbps US of the Symmetric XGS-PON competitor shown in Figure 1.

Hypothetical Bandwidth Capacities for each of these systems is shown in the Table, but it should be understood that the equipment to support the higher capacity solutions does not yet exist. As a result, the authors were forced to utilize their best guestimates on the potential performance levels and frequency responses of futuristic, yet-to-be-designed equipment such as Nodes and Amplifiers and Taps and CMs. These performance levels were extrapolations from current designs. Thus, they may or may not be accurate, and the indicated results may or may not be achievable. The actual performance levels will not be known until final designs are completed. However, the performance levels within the Table will be utilized in the analysis below.


**Table 3 – Examples of Extended Spectrum DOCSIS Variants & their Predicted Performance Levels (Zero-Order Models)**

Spectrum Type	Top of Spectrum (MHz)	Bottom of Spectrum (MHz)	Assumed Size of Guard-band Region (85-108?) within the Spectral Range (MHz)	Total Useable Spectral Range (MHz)	Assumed Spectral Efficiency in Node+X System (bps/Hz)	Possible BW Capacity in Node+X System (Mbps)	Assumed Spectral Efficiency in Node+X System w/ Active Taps (bps/Hz)	Possible BW Capacity in Node+X System w/ Active Taps to increase spectral efficiency at high frequencies (Mbps)	Assumed Spectral Efficiency in FTTT System (bps/Hz)	Possible BW Capacity in FTTT System (Mbps)
<b>Downstream:</b>										
1218 MHz Downstream	1218	108	0	1110	9.6	10656	9.6	10656	9.6	10656
1794 MHz Downstream ESD	1794	108	0	1686	9.3	15680	9.3	15680	9.6	16186
3000 MHz Downstream ESD	3000	108	0	2892	7.1	20533	8	23136	9.6	27763
6000 MHz Downstream ESD	6000	108	0	5892	3.2	18854	8	47136	9.6	56563
12000 MHz Downstream ESD	12000	108	0	11892	NA	Not supported w/ TEM propagation on Hard-line	NA	Not supported w/ TEM propagation on Hard-line	9.4	111785
<b>Upstream:</b>										
42 MHz Upstream	42	20	0	22	7.5	165	7.5	165	7.5	165
65 MHz Upstream	65	20	0	45	7.5	338	7.5	338	7.5	338
85 MHz Upstream	85	20	0	65	7.5	488	7.5	488	7.5	488
204 MHz Upstream	204	20	23	161	7.5	1208	7.5	1208	7.5	1208
300 MHz Upstream	300	20	23	257	7.5	1928	7.5	1928	7.5	1928
396 MHz Upstream	396	20	23	353	7.5	2648	7.5	2648	7.5	2648
492 MHz Upstream	492	20	23	449	7.5	3368	7.5	3368	7.5	3368
588 MHz Upstream	588	20	23	545	7.5	4088	7.5	4088	7.5	4088
684 MHz Upstream	684	20	23	641	7.5	4808	7.5	4808	7.5	4808
1218 MHz Upstream ESD	1218	20	23	1175	7.5	8813	7.5	8813	7.5	8813
1794 MHz Upstream ESD	1794	20	23	1751	6.9	12082	7.5	13133	7.5	13133
3000 MHz Upstream ESD	3000	20	23	2957	4.2	12419	7.5	22178	7.5	22178
6000 MHz Upstream ESD	6000	20	23	5957	1.8	10723	7.5	44678	7.5	44678
12000 MHz Upstream ESD	12000	20	23	11957	NA	Not supported w/ TEM propagation on Hard-line	NA	Not supported w/ TEM propagation on Hard-line	6.5	77721


The basic idea behind ESD is to maintain the Total Composite Power output levels emitted from Nodes and Amplifiers and CMs to be at levels which are similar to those which exists today. To illustrate the concept, an example 3000 MHz (3 GHz) system will be described. For the analyses below, the authors assumed -69.5 dBmV as a typical Total Composite Power level for the Nodes and Amplifiers. As described above, the flattening of the Power Spectral Density at higher frequencies will be utilized along with OFDM bit-loading. The resulting Power Spectral Density and Total Composite Power level



(integrated across the frequency band) is shown in Figure 19. The resulting Bit-Loading and Throughput is shown in Figure 20.




**Figure 20 – Launch Power Spectral Density and Total Composite Power for a 3 GHz Extended Spectrum DOCSIS System**





**Figure 21 – OFDM Bit-loading & Corresponding Throughput for a 3 GHz Extended Spectrum DOCSIS Node+0 System**



**Figure 22 – Predicted 1.8 GHz Bandwidth Capacities as a Function of Amp-to-Amp Hard-line Coaxial Length and Amplifier Output Power Levels and Cascade Lengths**

Figure 21 shows interesting predictions on Bandwidth Capacity changes as a function of Amp-to-Amp Hard-line coaxial length, Amplifier output levels, and cascade lengths for a 1.8 GHz solution. All of this data illustrates that the ESD concept does have a reasonable chance of working. In addition, Traditional-FDX or Dynamic Soft-FDX or Static Soft-FDX techniques could all be used to operate within the Upstream ESD spectral ranges and the Downstream ESD spectral ranges described above.

Likelihood of Success: Medium probability of success; much development is still required and Nodes and Amplifiers and Taps and CMs that support these higher frequencies are still being researched.

## 2.7. Distributed Node Architectures and Active Taps for Futuristic Upstream and Downstream Bandwidth Augmentation

In the more distant future, two other advanced technologies may be considered for Bandwidth Capacity augmentation. They are Distributed Node Architectures and Active Taps. These technologies will also be considered in the analysis below, so they will be briefly described.

The Distributed Node Architectures (DNA) [MU16] concept propose a novel idea to use fiber feeds emanating south-bound from a Primary Fiber Node instead of coaxial outputs. The Primary Fiber Node might be a standard Fiber Node or a DAA Fiber Node. These south-bound fiber feeds can then be routed with low loss to other Optical-to-Electrical converter elements (typically at the previous sites of Amplifiers). These Optical-to-Electrical converter elements would be similar to low-cost RF over Glass (RfOG) ONUs, but they would typically reside in Outside Plant housings (such as Amplifier or Tap housings) and their output RF power levels on the south-bound coax would be capable of driving many homes connected to the south-bound coax. The DNA solution offers several potential benefits. First, it could greatly increase the area served and the number of subscribers served by the Primary Fiber Node, and this could help reduce the Cost per HHP for the deployments of the Primary Fiber Nodes. Secondly, it could provide a technique that permits future network evolutions such as the deployment of Fiber-To-The-Tap (FTTT) technologies. In these futuristic instantiations, the Primary Fiber Node could be a DAA Node supporting (for example) 24 Service Groups. The South-bound fiber feeds could employ Wavelength Division Multiplexing (WDM) techniques to deliver unique signals to each of 24 Taps subtending from the Primary Fiber Node. Each Tap would then receive a unique signal from the fiber feed and convert that unique signal into an RF signal that is then split and amplified and fed to the four or eight drop coaxes to each home. Using short ~150 foot drop coaxes from the Taps (as the only coaxial runs in the system), a previous study showed that very high bandwidth capacities exceeding 100 Gbps could be delivered via DOCSIS to the homes in these DNA FTTT environments. [CL15] [CL16]

The Active Tap concept is another futuristic idea that can help increase Bandwidth Capacity to subscribers. It places small bidirectional amplifiers in all or some of the Taps within the coaxial run. The amplification of the Downstream and Upstream RF signals in the upper frequency ranges can help to increase Signal-to-Noise ratios and ensure that Extended Spectrum DOCSIS systems can operate without having to resort to low bit-loading levels in the upper frequency range. This approach can help to increase the average Spectral Efficiency and the total Bandwidth Capacity of the system, as illustrated in the blue column of Table 3.

Likelihood of Success: Medium probability of success; much development is still required and most of these concepts are still being researched.

### **3. Analysis of Migration Paths for Different Architectures**

With growing competitive threats and growing Bandwidth Capacity requirements, MSOs are planning to make HFC plant changes as the industry enters the 2020 decade. With the large array of technologies (outlined in the previous section) promising a myriad of solutions for their use in the future, many MSOs are asking important questions about which technologies to utilize and when to utilize them. In addition, different constraints are forcing MSOs to consider differing paths that are quite divergent.

Analyzing the many different plans that have been considered by MSOs in the past few years is not possible within a single paper. As a result, this paper will only select a subset of the many MSO architectural paths that are currently being considered- paths with very divergent approaches will be included to contrast the attributes of the different solutions. While only a few paths will be included, it is nevertheless hoped that the small number of MSO architectures analyzed within this paper will give some hints on the pros and cons of the different paths for the Cable Industry in the future.

Eighteen different architectures will be analyzed. The architectures can be divided into two different groupings based on the Upstream to Downstream Bandwidth ratio attribute. Some MSOs plan to offer their subscribers Asymmetric SLAs (with Upstream T<sub>max</sub> = ~50% of Downstream T<sub>max</sub>), while other

MSOs plan to offer their subscribers Symmetric SLAs (with Upstream T<sub>max</sub> = Downstream T<sub>max</sub>). Nine of the architectures will use Asymmetric SLAs, and nine of the architectures will use Symmetric SLAs.

The nine architectures (labeled 1 through 9) in both of the sets will each have a unique set of attributes that have been considered by MSOs. The attributes of the nine architectures are outlined below.

- **Architecture 1 (Traditional-FDX):** A Traditional-FDX solution is assumed with overlapping Upstream and Downstream frequency ranges. The Downstream T<sub>max</sub> experiences a 25% CAGR in the 2020's, a 15% CAGR in the 2030's, and a 15% CAGR in the 2040's. The Downstream T<sub>avg</sub> experiences a 39% CAGR in the 2020's, a 29% CAGR in the 2030's, and a 19% CAGR in the 2040's. The Upstream T<sub>avg</sub> experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to begin with 1.2 GHz Node+0 operation in 2020, and then the MSO attempts to stay with 1.2 GHz Node+0 and ESD until Bandwidth Capacity requirements force a transition to ESD spectra. Architecture 1a supports Asymmetric SLAs, and Architecture 1b supports Symmetric SLAs.
- **Architecture 2 (Static Soft-FDX Base-line w/ 1.2 GHz Affinity):** A “base-line” Static Soft-FDX solution is assumed with non-overlapping Upstream and Downstream frequency ranges and a guard-band between those frequency ranges. The Downstream T<sub>max</sub> experiences a 25% CAGR in the 2020's, a 15% CAGR in the 2030's, and a 15% CAGR in the 2040's. The Downstream T<sub>avg</sub> experiences a 39% CAGR in the 2020's, a 29% CAGR in the 2030's, and a 19% CAGR in the 2040's. The Upstream T<sub>avg</sub> experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to start with 1.2 GHz Node+3 operation in 2020, and then the MSO moves to Node+2 and Node+1 when required (avoiding ESD as long as possible and avoiding the costs of Node+0); this continues until Bandwidth Capacity requirements force a transition to ESD spectra. The MSO is assumed to try to keep 1.2 GHz operation for as long as possible, giving preference to Node-splits over ESD operation during this time. Architecture 2a supports Asymmetric SLAs, and Architecture 2b supports Symmetric SLAs.
- **Architecture 3 (Static Soft-FDX w/ Node+3 Affinity):** A Static Soft-FDX solution is assumed with non-overlapping Upstream and Downstream frequency ranges and a guard-band between those frequency ranges. The Downstream T<sub>max</sub> experiences a 25% CAGR in the 2020's, a 15% CAGR in the 2030's, and a 15% CAGR in the 2040's. The Downstream T<sub>avg</sub> experiences a 39% CAGR in the 2020's, a 29% CAGR in the 2030's, and a 19% CAGR in the 2040's. The Upstream T<sub>avg</sub> experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to start with 1.2 GHz Node+3 operation in 2020. This continues until Bandwidth Capacity requirements force a transition to ESD spectra. The MSO is assumed to try to keep Node+3 operation for as long as possible, giving preference to ESD operation over Node-splits during this time. Architecture 3a supports Asymmetric SLAs, and Architecture 3b supports Symmetric SLAs.
- **Architecture 4 (Static Soft-FDX w/ 15% DS T<sub>max</sub> CAGR):** A Static Soft-FDX solution similar to the “base-line” is assumed with non-overlapping Upstream and Downstream frequency ranges and a guard-band between those frequency ranges, but the Downstream T<sub>max</sub> experiences a 15% CAGR in the 2020's (instead of 25%), a 15% CAGR in the 2030's, and a 15% CAGR in the 2040's. The Downstream T<sub>avg</sub> experiences a 39% CAGR in the 2020's, a 29% CAGR in the 2030's, and a 19% CAGR in the 2040's. The Upstream T<sub>avg</sub> experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to start with 1.2 GHz Node+3 operation in 2020, and then the MSO moves to Node+2 and Node+1 when required (avoiding the costs of Node+0); this continues until Bandwidth Capacity requirements force a transition to ESD spectra.

The MSO is assumed to try to keep 1.2 GHz operation for as long as possible, giving preference to Node-splits over ESD operation during this time. Architecture 4a supports Asymmetric SLAs, and Architecture 4b supports Symmetric SLAs.

- Architecture 5 (Static Soft-FDX w/ Reduced US Tmax):** A Static Soft-FDX solution similar to the “base-line” is assumed with non-overlapping Upstream and Downstream frequency ranges and a guard-band between those frequency ranges. The Downstream Tmax experiences a 25% CAGR in the 2020’s, a 15% CAGR in the 2030’s, and a 15% CAGR in the 2040’s. The Upstream Tmax is slightly reduced from the “base-line”. The Downstream Tavg experiences a 39% CAGR in the 2020’s, a 29% CAGR in the 2030’s, and a 19% CAGR in the 2040’s. The Upstream Tavg experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to start with 1.2 GHz Node+3 operation in 2020, and then the MSO moves to Node+2 and Node+1 when required (avoiding the costs of Node+0); this continues until Bandwidth Capacity requirements force a transition to ESD spectra. The MSO is assumed to try to keep 1.2 GHz operation for as long as possible, giving preference to Node-splits over ESD operation during this time. Architecture 5a supports Asymmetric SLAs, and Architecture 5b supports Symmetric SLAs.
- Architecture 6 (Static Soft-FDX w/ Selective Subscriber Migration):** A Static Soft-FDX solution similar to the “base-line” is assumed with non-overlapping Upstream and Downstream frequency ranges and a guard-band between those frequency ranges. The Downstream Tmax experiences a 25% CAGR in the 2020’s, a 15% CAGR in the 2030’s, and a 15% CAGR in the 2040’s, but the subscribers with the top tier SLA are assumed to always be moved to an alternative infrastructure, yielding Tmax’s that are ½ of their normal values in the “base-line.” The Downstream Tavg experiences a 39% CAGR in the 2020’s, a 29% CAGR in the 2030’s, and a 19% CAGR in the 2040’s. The Upstream Tavg experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to start with 1.2 GHz Node+3 operation in 2020, and then the MSO moves to Node+2 and Node+1 when required (avoiding the costs of Node+0); this continues until Bandwidth Capacity requirements force a transition to ESD spectra. The MSO is assumed to try to keep 1.2 GHz operation for as long as possible, giving preference to Node-splits over ESD operation during this time. Architecture 6a supports Asymmetric SLAs, and Architecture 6b supports Symmetric SLAs.
- Architecture 7 (Static Soft-FDX w/ Guard-band Elimination):** A Static Soft-FDX solution similar to the “base-line” is assumed with non-overlapping Upstream and Downstream frequency ranges, but no guard-band between those frequency ranges. The Downstream Tmax experiences a 25% CAGR in the 2020’s, a 15% CAGR in the 2030’s, and a 15% CAGR in the 2040’s. The Downstream Tavg experiences a 39% CAGR in the 2020’s, a 29% CAGR in the 2030’s, and a 19% CAGR in the 2040’s. The Upstream Tavg experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to start with 1.2 GHz Node+3 operation in 2020, and then the MSO moves to Node+2 and Node+1 when required (avoiding the costs of Node+0); this continues until Bandwidth Capacity requirements force a transition to ESD spectra. The MSO is assumed to try to keep 1.2 GHz operation for as long as possible, giving preference to Node-splits over ESD operation during this time. Architecture 7a supports Asymmetric SLAs, and Architecture 7b supports Symmetric SLAs.
- Architecture 8 (Dynamic Soft-FDX Base-line):** A “base-line” Dynamic Soft-FDX solution is assumed with non-overlapping Upstream and Downstream frequency ranges within each RF Leg (but overlapping Upstream and Downstream frequency ranges within a Node) and a guard-band between those frequency ranges. Upstream and Downstream bandwidth monitoring that can rapidly change the split in any RF Leg is also assumed. The Downstream Tmax experiences a

25% CAGR in the 2020's, a 15% CAGR in the 2030's, and a 15% CAGR in the 2040's. The Downstream Tavg experiences a 39% CAGR in the 2020's, a 29% CAGR in the 2030's, and a 19% CAGR in the 2040's. The Upstream Tavg experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to start with 1.2 GHz Node+3 operation in 2020, and then the MSO moves to Node+2 and Node+1 when required (avoiding the costs of Node+0); this continues until Bandwidth Capacity requirements force a transition to ESD spectra. The MSO is assumed to try to keep 1.2 GHz operation for as long as possible, giving preference to Node-splits over ESD operation during this time. Architecture 8a supports Asymmetric SLAs, and Architecture 8b supports Symmetric SLAs.

- Architecture 9 (Static Soft-FDX w/ Active Taps):** A Static Soft-FDX solution similar to the “base-line” is assumed with non-overlapping Upstream and Downstream frequency ranges and a guard-band between those frequency ranges, but Active Taps are included in the path of the coax to improve Singal-to-Noise ratios and Spectral Efficiencies at higher frequencies. The Downstream Tmax experiences a 25% CAGR in the 2020's, a 15% CAGR in the 2030's, and a 15% CAGR in the 2040's. The Downstream Tavg experiences a 39% CAGR in the 2020's, a 29% CAGR in the 2030's, and a 19% CAGR in the 2040's. The Upstream Tavg experiences a 19% CAGR for all time. 2-Service Group Nodes are used until 2028, and then 4-Service Group Nodes are used until FTTT solutions are required. The MSO is assumed to start with 1.2 GHz Node+3 operation in 2020, and then the MSO moves to Node+2 and Node+1 when required (avoiding the costs of Node+0); this continues until Bandwidth Capacity requirements force a transition to ESD spectra. The MSO is assumed to try to keep 1.2 GHz operation for as long as possible, giving preference to Node-splits over ESD operation during this time. Architecture 2a supports Asymmetric SLAs, and Architecture 2b supports Symmetric SLAs.

The authors attempted to predict the yearly decisions that might be made by an MSO working with each of the nine Architectures and for both Asymmetric SLAs (in Figures 22-30) and Symmetric SLAs (in Figures 31-39), and the resulting changes and migration paths (for the 25 years from 2020 to 2044) for all eighteen of the resulting Architectures are displayed in Figures 22-39 below. The orange region gives a description of the HFC Plant, the yellow region describes the Upstream Bandwidth requirements, and the green section describes the Downstream Bandwidth requirements.

While the Figures display only even-numbered years (for brevity), many important decisions needed to be made on a yearly basis within the predictive analysis of each of the Figures. The need for change within each year of each Figure was predominantly driven by expected yearly increases in US Tavg, US Tmax, DS Tavg, and DS Tmax. The questions answered within each year of each Figure included:

- 1) What is the year's US Tavg & US Tmax & DS Tavg & DS Tmax & Nsub value?
- 2) What is the Required Upstream HSD Bandwidth Capacity given by  $N_{sub} \cdot T_{avg} + 1.0 \cdot T_{max}$ ?
- 3) What is the Required Downstream HSD Bandwidth Capacity given by  $N_{sub} \cdot T_{avg} + 1.0 \cdot T_{max}$ ?
- 4) Should a Node-split or a move to FTTT or a move to Selective Subscriber Migration be performed?
- 5) What is the year's new US Tavg & US Tmax & DS Tavg & DS Tmax & Nsub value?
- 6) What is the new Required Upstream HSD Bandwidth Capacity given by  $N_{sub} \cdot T_{avg} + 1.0 \cdot T_{max}$ ?
- 7) What is the new Required Downstream HSD Bandwidth Capacity given by  $N_{sub} \cdot T_{avg} + 1.0 \cdot T_{max}$ ?
- 8) Should the Bottom & Top of the US DOCSIS Spectrum be moved to increase US Capacity?
- 9) What SNR levels and bit-loading levels and US Spectral Efficiency can be supported in the resulting US DOCSIS Spectrum (assuming launch power levels are fixed)?

- 10) Does that US DOCSIS Spectrum and US Spectral Efficiency support the Required Upstream HSD Bandwidth Capacity requirement?
- 11) If not, return to step (8)
- 12) Should the Bottom & Top of the DS DOCSIS Spectrum be moved to increase DS Capacity?
- 13) What SNR levels and bit-loading levels and DS Spectral Efficiency can be supported in the resulting DS DOCSIS Spectrum (assuming launch power levels are fixed)?
- 14) Does that DS DOCSIS Spectrum and DS Spectral Efficiency support the Required Downstream HSD Bandwidth Capacity requirement?
- 15) If not, return to step (12)
- 16) Is the solution acceptable?
- 17) If not, return to step (4)

The results of all of these decisions made on a yearly basis are illustrated in the Migration Paths within each of the Figures below. It should be understood that these particular Migration Paths are not being define as “the best” Migration Paths. With so many decisions to be made on a yearly basis, it is clear that different MSOs will likely make different decisions leading to many different and desirable Migration Paths. As a result, the Migration Paths depicted below should only be used as examples to guide our analysis. Other equally valid Migration Paths are also possible.

FDX Asymmetric Tmax (US = 50% of DS) with 0 Guardbands (EC)	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+0	Node+0	Node+0	Node+0	Node+0	Node+0	Node+0	Node+0	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	120	120	120	120	120	120	120	120	120	120	120	120	120
Nsub in Service Group	30	30	30	30	30	15	15	15	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	417	417	417	417	417	417	417	417	417	417	417	417	417
US Tavg (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
US Tmax_max (Mbps)	500	1000	1000	2000	2000	5000	5000	5000	10000	10000	20000	20000	40000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	508	1012	1017	2024	2034	5024	5034	5048	10018	10026	20036	20051	40073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	1.65%	1.18%	1.66%	1.18%	1.66%	0.48%	0.67%	0.95%	0.18%	0.26%	0.18%	0.26%	0.18%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	73	140	141	275	276	675	676	678	1341	1342	2677	2679	5348
Top of US Band (MHz)	85	204	204	300	300	684	684	684	1794	1794	3000	3000	6000
DS Tavg (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1069	2133	2258	4498	4962	10862	11435	12387	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	6.45%	6.25%	11.41%	11.06%	19.38%	7.94%	12.55%	19.27%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	5	5	5	5	5	5	5	5	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	9.3	9.3	9.3	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-23	96	96	192	192	576	576	576	1686	1686	2892	2892	5892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-2%	8%	8%	16%	16%	32%	32%	32%	56%	56%	48%	48%	49%
Bottom of DS DOCSIS Spectrum (MHz)	108	108	108	108	108	108	108	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	555	666	679	913	961	1612	1674	1776	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1218	1218	1794	1794	1794	3000	3000	6000	6000	12000

**Figure 23 – Migration Path for Architecture 1a (Traditional-FDX- Asymmetric SLA)**

Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+2	Node+1	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	480	240	240	240	240	120	120	120	120	120
Nsub in Service Group	200	200	200	120	60	30	30	30	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	105	209	209	209	209	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
US_Tmax_max (Mbps)	500	1000	1000	2000	2000	5000	5000	5000	10000	10000	20000	20000	40000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	556	1079	1112	2095	2068	5048	5068	5096	10018	10026	20036	20051	40073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	10.07%	7.35%	10.10%	4.55%	3.27%	0.95%	1.34%	1.88%	0.18%	0.26%	0.18%	0.26%	0.18%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	79	149	153	284	281	678	681	684	1341	1342	2677	2679	5348
Top of US Band (MHz)	85	204	204	300	300	684	684	684	1794	1794	3000	3000	6000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	5991	5923	11724	12869	14775	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	33.23%	32.47%	14.71%	22.29%	32.32%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	1	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	7.1	7.1	7.1	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-15	-36	-36	-53	-53	-120	-120	-120	1686	1686	2892	2892	5892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-1%	-3%	-3%	-4%	-4%	-4%	-4%	-4%	56%	56%	48%	48%	49%
Bottom of DS DOCSIS Spectrum (MHz)	100	240	240	353	353	804	804	804	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	588	877	963	1313	1305	2791	2952	3221	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1218	1218	3000	3000	3000	3000	3000	6000	6000	12000

**Figure 24 – Migration Path for Architecture 2a (Static Soft-FDX Baseline w/ 1.2 GHz Affinity- Asymmetric SLA)**

Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+3	Node+3	Node+2	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	800	800	480	240	120	120	120	120	120
Nsub in Service Group	200	200	200	200	200	100	60	30	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	63	63	105	209	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
US_Tmax_max (Mbps)	500	1000	1000	2000	2000	5000	5000	5000	10000	10000	20000	20000	40000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	556	1079	1112	2159	2225	5159	5135	5096	10018	10026	20036	20051	40073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	10.07%	7.35%	10.10%	7.37%	10.12%	3.09%	2.64%	1.88%	0.18%	0.26%	0.18%	0.26%	0.18%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	79	149	153	293	302	693	690	684	1341	1342	2677	2679	5348
Top of US Band (MHz)	85	204	204	300	300	684	684	684	1794	1794	3000	3000	6000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	7318	10410	15747	15738	14775	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	45.34%	61.58%	36.50%	36.46%	32.32%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	1	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.3	9.3	7.1	3.2	3.2	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-15	-36	-36	-53	-53	-120	-120	-120	1686	1686	2892	2892	5892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-1%	-3%	-3%	-3%	-3%	-4%	-2%	-2%	56%	56%	48%	48%	49%
Bottom of DS DOCSIS Spectrum (MHz)	100	240	240	353	353	804	804	804	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	588	877	963	1475	1808	3358	6058	5757	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1794	1794	3000	6000	6000	3000	3000	6000	6000	12000

**Figure 25 – Migration Path for Architecture 3a (Static Soft-FDX w/ Node+3 Affinity - Asymmetric SLA)**

Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+1	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	240	240	240	240	120	120	120	120	120
Nsub in Service Group	200	200	200	200	60	30	30	30	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	209	209	209	209	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
US_Tmax_max (Mbps)	500	500	1000	1000	2000	2000	2000	5000	5000	5000	10000	10000	10000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	556	579	1112	1159	2068	2048	2068	5096	5018	5026	10036	10051	10073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	10.07%	13.69%	10.10%	13.72%	3.27%	2.34%	3.28%	1.88%	0.36%	0.51%	0.36%	0.51%	0.72%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	79	82	153	160	281	278	281	684	674	675	1343	1345	1348
Top of US Band (MHz)	85	85	204	204	300	300	300	684	684	675	1794	1794	1794
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 15% in 2020s, 15% in 2030s, 15% in 2040s	1000	1323	1749	2313	3059	4046	5350	7076	9358	12375	16367	21645	28625
"Rounded-off" DS Tmax_max (Mbps)	1000	1000	2000	2000	4000	4000	4000	10000	10000	10000	20000	20000	20000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	1889	3717	5318	5923	5724	6869	14775	11059	11763	22706	23832	25427
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	47.06%	46.20%	62.39%	32.47%	30.12%	41.77%	32.32%	9.58%	14.99%	11.92%	16.08%	21.34%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	1	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	9.6	9.3	7.1	9.6	9.6	9.6	9.6	9.6
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-15	-15	-36	-36	-53	-53	-53	-120	576	576	1686	1686	1686
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-1%	-1%	-3%	-3%	-4%	-4%	-3%	-4%	32%	32%	56%	56%	56%
Bottom of DS DOCSIS Spectrum (MHz)	100	100	240	240	353	353	353	804	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	588	633	963	1130	1305	1285	1427	3221	1596	1669	2809	2927	3093
Tap BW (MHz)	1218	1218	1218	1218	1218	1218	1794	3000	1794	1794	3000	3000	3000

**Figure 26 – Migration Path for Architecture 4a (Static Soft-FDX w/ 15% DS Tmax CAGR-Asymmetric SLA)**

Static SoftFDD Asymmetric Tmax (US = 40% of DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+2	Node+1	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	480	240	240	240	240	120	120	120	120	120
Nsub in Service Group	200	200	200	120	60	30	30	30	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	105	209	209	209	209	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40
US_Tmax_max (Mbps)	400	800	800	1600	1600	4000	4000	4000	8000	8000	16000	16000	32000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	456	879	912	1695	1668	4048	4068	4096	8018	8026	16036	16051	32073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	12.28%	9.02%	12.31%	5.63%	4.05%	1.18%	1.67%	2.34%	0.23%	0.32%	0.23%	0.32%	0.23%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	66	122	127	231	227	545	547	551	1074	1075	2143	2145	4281
Top of US Band (MHz)	85	204	204	204	204	492	492	492	1794	1794	3000	3000	6000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	5991	5923	11724	12869	14775	21059	21763	42706	43832	85427
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	33.23%	32.47%	14.71%	22.29%	32.32%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	1	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	7.1	7.1	7.1	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-15	-36	-36	-36	-36	-86	-86	-86	1686	1686	2892	2892	5892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-1%	-3%	-3%	-3%	-3%	-3%	-3%	-3%	56%	56%	48%	48%	49%
Bottom of DS DOCSIS Spectrum (MHz)	100	240	240	240	240	578	578	578	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	588	877	963	1200	1193	2565	2727	2995	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1218	1218	3000	3000	3000	3000	3000	6000	6000	12000

**Figure 27 – Migration Path for Architecture 5a (Static Soft-FDX w/ Reduced US Tmax-Asymmetric SLA)**



Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband & Sel Sub Mig	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+1	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	240	240	240	240	120	120	120	120	120
Nsub in Service Group	200	200	200	200	60	30	30	30	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	209	209	209	209	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
US Tmax_max (Mbps)	500	500	500	1000	1000	2000	2000	5000	5000	5000	10000	10000	20000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	556	579	612	1159	1068	2048	2068	5096	5018	5026	10036	10051	20073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	10.07%	13.69%	18.34%	13.72%	6.33%	2.34%	3.28%	1.88%	0.36%	0.51%	0.36%	0.51%	0.36%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	79	82	87	160	147	278	281	684	674	675	1343	1345	2681
Top of US Band (MHz)	85	85	85	204	204	300	300	684	684	675	1794	1794	3000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	500	781	1221	1907	2980	4284	5666	7493	9909	13105	17332	22921	30313
"Rounded-off" DS Tmax_max (Mbps)	1000	1000	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	1889	2717	5318	3923	5724	6869	14775	11059	11763	22706	23832	45427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	47.06%	63.20%	62.39%	49.02%	30.12%	41.77%	32.32%	9.58%	14.99%	11.92%	16.08%	11.95%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	1	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	9.6	9.3	7.1	9.6	9.6	9.6	9.6	9.6
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-15	-15	-15	-36	-36	-53	-53	-120	576	576	1686	1686	2892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-1%	-1%	-1%	-3%	-3%	-4%	-3%	-4%	32%	32%	56%	56%	48%
Bottom of DS DOCSIS Spectrum (MHz)	100	100	100	240	240	353	353	804	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	588	633	719	1130	984	1285	1427	3221	1596	1669	2809	2927	5176
Tap BW (MHz)	1218	1218	1218	1218	1218	1218	1794	3000	1794	1794	3000	3000	6000

**Figure 28 – Migration Path for Architecture 6a (Static Soft-FDX w/ Selective Subscriber Migration- Asymmetric SLA)**

Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 0 Guardbands (EC)	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+1	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	240	240	240	240	120	120	120	120	120
Nsub in Service Group	200	200	200	200	60	30	30	30	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	209	209	209	209	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
US Tmax_max (Mbps)	500	1000	1000	2000	2000	5000	5000	5000	10000	10000	20000	20000	40000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	556	1079	1112	2159	2068	5048	5068	5096	10018	10026	20036	20051	40073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	10.07%	7.35%	10.10%	7.37%	3.27%	0.95%	1.34%	1.88%	0.18%	0.26%	0.18%	0.26%	0.18%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	79	149	153	293	281	678	681	684	1341	1342	2677	2679	5348
Top of US Band (MHz)	85	204	204	300	300	684	684	684	1794	1794	3000	3000	6000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	7318	5923	11724	12869	14775	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	45.34%	32.47%	14.71%	22.29%	32.32%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	2	2	2	2	2	2	2	2	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	7.1	7.1	7.1	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	0	0	0	0	0	0	0	0	1686	1686	2892	2892	5892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	0%	0%	0%	0%	0%	0%	0%	0%	56%	56%	48%	48%	49%
Bottom of DS DOCSIS Spectrum (MHz)	85	204	204	300	300	684	684	684	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	573	841	927	1398	1253	2671	2833	3101	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1218	1218	3000	3000	3000	3000	3000	6000	6000	12000

**Figure 29 – Migration Path for Architecture 7a (Static Soft-FDX w/ Guard-band Elimination- Asymmetric SLA)**

Dynamic SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+2	Node+2	Node+2	Node+2	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	480	480	480	480	120	120	120	120	120
Nsub in Service Group	200	200	200	200	120	60	60	60	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	105	105	105	105	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
US_Tmax_max (Mbps)	500	1000	1000	2000	2000	5000	5000	5000	10000	10000	20000	20000	40000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	556	1079	1112	2159	2135	5096	5135	5192	10018	10026	20036	20051	40073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	10.07%	7.35%	10.10%	7.37%	6.33%	1.88%	2.64%	3.70%	0.18%	0.26%	0.18%	0.26%	0.18%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	79	149	153	293	290	684	690	697	1341	1342	2677	2679	5348
Top of US Band (MHz)	85	204	204	300	300	684	684	684	1794	1794	3000	3000	6000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	7318	7846	13448	15738	19549	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	45.34%	49.02%	25.64%	36.46%	48.85%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	3	3	3	3	3	3	3	3	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	9.3	7.1	7.1	9.6	9.6	9.6	9.6	9.6
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-b	-23	96	96	192	192	576	576	1686	1686	2892	2892	2892	5892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-2%	8%	8%	16%	16%	32%	19%	19%	56%	56%	48%	48%	49%
Bottom of DS DOCSIS Spectrum (MHz)	108	108	108	108	108	108	108	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	596	781	867	1259	1314	2010	2780	3317	2952	3025	5418	5535	10393
Tap BW (MHz)	1218	1218	1218	1218	1218	1794	3000	3000	3000	3000	6000	6000	12000

**Figure 30 – Migration Path for Architecture 8a (Dynamic Soft-FDX Baseline- Asymmetric SLA)**

Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	4	4	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+2	Node+1	Node+1	Node+1	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	480	240	240	240	240	240	240	120	120	120
Nsub in Service Group	200	200	200	120	60	30	30	30	30	30	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	105	209	209	209	209	209	209	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
US_Tmax_max (Mbps)	500	1000	1000	2000	2000	5000	5000	5000	10000	10000	20000	20000	40000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	556	1079	1112	2095	2068	5048	5068	5096	10136	10192	20036	20051	40073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	10.07%	7.35%	10.10%	4.55%	3.27%	0.95%	1.34%	1.88%	1.34%	1.89%	0.18%	0.26%	0.18%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	79	149	153	284	281	678	681	684	1356	1364	2677	2679	5348
Top of US Band (MHz)	85	204	204	300	300	684	684	684	1794	1794	3000	3000	6000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	5991	5923	11724	12869	14775	27945	33222	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	33.23%	32.47%	14.71%	22.29%	32.32%	28.43%	39.80%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	1	1	1	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	8.0	8.0	8.0	8.0	8.0	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-15	-36	-36	-53	-53	-120	-120	-120	-314	-314	2892	2892	5892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-1%	-3%	-3%	-4%	-4%	-4%	-4%	-4%	-5%	-5%	48%	48%	49%
Bottom of DS DOCSIS Spectrum (MHz)	100	240	240	353	353	804	804	804	2108	2108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	588	877	963	1313	1305	2605	2748	2987	5937	6597	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1218	1218	3000	3000	3000	6000	6000	6000	6000	12000

**Figure 31 – Migration Path for Architecture 9a (Static Soft-FDX w/ Active Taps- Asymmetric SLA)**

FDX Symmetric Tmax (US = DS) with 0 Guardbands (EC)	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+0	Node+0	Node+0	Node+0	Node+0	Node+0	Node+0	Node+0	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	120	120	120	120	120	120	120	120	120	120	120	120	120
Nsub in Service Group	30	30	30	30	30	15	15	15	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	417	417	417	417	417	417	417	417	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
US Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	1008	2012	2017	4024	4034	10024	10034	10048	20018	20026	40036	40051	80073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	0.83%	0.59%	0.84%	0.59%	0.84%	0.24%	0.34%	0.48%	0.09%	0.13%	0.09%	0.13%	0.09%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	6.9	6.9	6.9	7.5	7.5	7.5	7.5	6.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	139	273	274	542	543	1458	1459	1461	2674	2675	5343	5345	12324
Top of US Band (MHz)	204	300	300	492	492	1794	1794	1794	3000	3000	6000	6000	12000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1069	2133	2258	4498	4962	10862	11435	12387	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	6.45%	6.25%	11.41%	11.06%	19.38%	7.94%	12.55%	19.27%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	5	5	5	5	5	5	5	5	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	9.3	9.3	9.3	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	96	192	192	384	384	1686	1686	1686	2892	2892	5892	5892	11892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	8%	16%	16%	32%	32%	94%	94%	94%	96%	96%	98%	98%	99%
Bottom of DS DOCSIS Spectrum (MHz)	108	108	108	108	108	108	108	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	555	666	679	913	961	1612	1674	1776	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1218	1218	1794	1794	1794	3000	3000	6000	6000	12000

**Figure 32 – Migration Path for Architecture 1b (Traditional-FDX- Symmetric SLA)**

Static SoftFDD Symmetric Tmax (US = DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	24	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+2	Node+2	Node+2	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	480	480	480	240	120	120	120	120	120	120
Nsub in Service Group	200	200	200	120	120	60	30	4	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	105	105	105	209	417	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
US Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	1056	2079	2112	4095	4135	10096	10068	10013	20018	20026	40036	40051	80073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	5.30%	3.81%	5.32%	2.33%	3.27%	0.95%	0.67%	0.13%	0.09%	0.13%	0.09%	0.13%	0.09%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	6.9	6.9	7.5	7.5	7.5	7.5	7.5	6.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	146	282	287	551	556	1468	1464	1340	2674	2675	5343	5345	12324
Top of US Band (MHz)	204	300	300	492	492	1794	1794	1794	3000	3000	6000	6000	12000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	5991	7846	13448	12869	10637	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	33.23%	49.02%	25.64%	22.29%	5.99%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	5	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.3	9.3	3.2	3.2	9.6	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-36	-53	-53	-86	-86	-314	-314	1686	2892	2892	5892	5892	11892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-3%	-4%	-4%	-5%	-5%	-5%	-5%	94%	96%	96%	98%	98%	99%
Bottom of DS DOCSIS Spectrum (MHz)	240	353	353	578	578	2108	2108	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	728	989	1076	1558	1758	6647	6466	1552	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1794	1794	6000	6000	1794	3000	3000	6000	6000	12000

**Figure 33 – Migration Path for Architecture 2b (Static Soft-FDX Baseline w/ 1.2 GHz Affinity- Symmetric SLA)**

Static SoftFDD Asymmetric Tmax (US = DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	24	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+3	Node+2	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	800	480	240	120	120	120	120	120	120
Nsub in Service Group	200	200	200	200	200	60	30	4	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	63	105	209	417	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
US_Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	1056	2079	2112	4159	4225	10096	10068	10013	20018	20026	40036	40051	80073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	5.30%	3.81%	5.32%	3.82%	5.33%	0.95%	0.67%	0.13%	0.09%	0.13%	0.09%	0.13%	0.09%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	6.9	6.9	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	146	282	287	560	568	1468	1464	1340	2674	2675	5343	5345	10681
Top of US Band (MHz)	204	300	300	492	684	1794	1794	1794	3000	3000	6000	6000	12000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	7318	10410	13448	12869	10637	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	45.34%	61.58%	25.64%	22.29%	5.99%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	5	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.3	7.1	3.2	3.2	9.6	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-36	-53	-53	-86	-120	-314	-314	1686	2892	2892	5892	5892	11892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-3%	-4%	-4%	-5%	-4%	-5%	-5%	94%	96%	96%	98%	98%	99%
Bottom of DS DOCSIS Spectrum (MHz)	240	353	353	578	804	2108	2108	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	728	989	1076	1701	2606	6647	6466	1552	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1794	3000	6000	6000	1794	3000	3000	6000	6000	12000

**Figure 34 – Migration Path for Architecture 3b (Static Soft-FDX w/ Node+3 Affinity - Symmetric SLA)**

Static SoftFDD Symmetric Tmax (US = DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	24	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	240	240	240	120	120	120	120	120	120
Nsub in Service Group	200	200	200	200	60	30	30	4	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	209	209	209	417	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
US_Tmax_max (Mbps)	1000	1000	2000	2000	4000	4000	4000	10000	10000	10000	20000	20000	20000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	1056	1079	2112	2159	4068	4048	4068	10013	10018	10026	20036	20051	20073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	5.30%	7.35%	5.32%	7.37%	1.66%	1.18%	1.67%	0.13%	0.18%	0.26%	0.18%	0.26%	0.36%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	146	149	287	293	547	545	547	1340	1341	1342	2677	2679	2681
Top of US Band (MHz)	204	204	300	300	492	492	492	1794	1794	1794	3000	3000	3000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 15% in 2020s, 15% in 2030s, 15% in 2040s	1000	1323	1749	2313	3059	4046	5350	7076	9358	12375	16367	21645	28625
"Rounded-off" DS Tmax_max (Mbps)	1000	1000	2000	2000	4000	4000	4000	10000	10000	10000	20000	20000	20000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	1889	3717	5318	5923	5724	6869	10637	11059	11763	22706	23832	25427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	47.06%	46.20%	62.39%	32.47%	30.12%	41.77%	5.99%	9.58%	14.99%	11.92%	16.08%	21.34%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	5	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	9.6	9.3	9.6	9.6	9.6	9.6	9.6	9.6
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-36	-36	-53	-53	-86	-86	-86	1686	1686	1686	2892	2892	2892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-3%	-3%	-4%	-4%	-5%	-5%	-5%	94%	94%	94%	96%	96%	96%
Bottom of DS DOCSIS Spectrum (MHz)	240	240	353	353	578	578	578	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	728	772	1076	1242	1531	1510	1653	1552	1596	1669	2809	2927	3093
Tap BW (MHz)	1218	1218	1218	1218	1794	1794	1794	1794	1794	1794	3000	3000	3000

**Figure 35 – Migration Path for Architecture 4b (Static Soft-FDX w/ 15% DS Tmax CAGR- Symmetric SLA)**

Static SoftFDD ~Symmetric Tmax (US = 90% of DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	24	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+2	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	480	240	240	240	120	120	120	120	120	120
Nsub in Service Group	200	200	200	120	60	30	30	4	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	105	209	209	209	417	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90
US Tmax_max (Mbps)	900	1800	1800	3600	3600	9000	9000	9000	18000	18000	36000	36000	72000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	956	1879	1912	3695	3668	9048	9068	9013	18018	18026	36036	36051	72073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	5.86%	4.22%	5.87%	2.58%	1.84%	0.53%	0.75%	0.14%	0.10%	0.14%	0.10%	0.14%	0.10%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	6.9	6.9	7.5	7.5	7.5	7.5	7.5	6.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	132	256	260	498	494	1316	1319	1207	2407	2408	4810	4812	11093
Top of US Band (MHz)	204	300	300	492	492	1794	1794	1794	3000	3000	6000	6000	12000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	5991	5923	11724	12869	10637	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	33.23%	32.47%	14.71%	22.29%	5.99%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	5	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.3	9.3	3.2	3.2	9.6	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-36	-53	-53	-86	-86	-314	-314	1686	2892	2892	5892	5892	11892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-3%	-4%	-4%	-5%	-5%	-5%	-5%	94%	96%	96%	98%	98%	99%
Bottom of DS DOCSIS Spectrum (MHz)	240	353	353	578	578	2108	2108	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	728	989	1076	1558	1551	6108	6466	1552	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1794	1794	6000	6000	1794	3000	3000	6000	6000	12000

**Figure 36 – Migration Path for Architecture 5b (Static Soft-FDX w/ Reduced US Tmax-Symmetric SLA)**

Static SoftFDD Symmetric Tmax (US = DS) with 1 Guardband & Sel Sub Mig	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	24	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	240	240	240	120	120	120	120	120	120
Nsub in Service Group	200	200	200	200	60	30	30	4	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	209	209	209	417	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
US Tmax_max (Mbps)	1000	1000	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	1056	1079	1112	2159	2068	4048	4068	10013	10018	10026	20036	20051	40073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	5.30%	7.35%	10.10%	7.37%	3.27%	1.18%	1.67%	0.13%	0.18%	0.26%	0.18%	0.26%	0.18%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	146	149	153	293	281	545	547	1340	1341	1342	2677	2679	5348
Top of US Band (MHz)	204	204	204	300	300	492	492	1794	1794	1794	3000	3000	6000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	500	781	1221	1907	2980	4284	5666	7493	9909	13105	17332	22921	30313
"Rounded-off" DS Tmax_max (Mbps)	1000	1000	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	1889	2717	5318	3923	5724	6869	10637	11059	11763	22706	23832	45427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	47.06%	63.20%	62.39%	49.02%	30.12%	41.77%	5.99%	9.58%	14.99%	11.92%	16.08%	11.95%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	5	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	9.3	9.3	9.6	9.6	9.6	9.6	9.6	9.6
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	-36	-36	-36	-53	-53	-86	-86	1686	1686	1686	2892	2892	5892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	-3%	-3%	-3%	-4%	-4%	-5%	-5%	94%	94%	94%	96%	96%	98%
Bottom of DS DOCSIS Spectrum (MHz)	240	240	240	353	353	578	578	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	728	772	859	1242	1097	1530	1653	1552	1596	1669	2809	2927	5176
Tap BW (MHz)	1218	1218	1218	1218	1218	1794	1794	1794	1794	1794	3000	3000	6000

**Figure 37 – Migration Path for Architecture 6b (Static Soft-FDX w/ Selective Subscriber Migration- Symmetric SLA)**

Static SoftFDD Symmetric Tmax (US = DS) with 0 Guardbands (EC)	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+1	Node+1	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	240	240	240	240	240	120	120	120	120	120
Nsub in Service Group	200	200	200	60	60	30	30	30	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	209	209	209	209	209	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
US Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	1056	2079	2112	4048	4068	10048	10068	10096	20018	20026	40036	40051	80073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	5.30%	3.81%	5.32%	1.18%	1.66%	0.48%	0.67%	0.95%	0.09%	0.13%	0.09%	0.13%	0.09%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	146	282	287	545	547	1345	1347	1351	2674	2675	5343	5345	10681
Top of US Band (MHz)	204	300	300	492	492	1794	1794	1794	3000	3000	6000	6000	12000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	4995	5923	11724	12869	14775	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	19.93%	32.47%	14.71%	22.29%	32.32%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	2	2	2	2	2	2	2	2	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.3	9.6	7.1	7.1	7.1	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	0	0	0	0	0	0	0	0	2892	2892	5892	5892	11892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	0%	0%	0%	0%	0%	0%	0%	0%	96%	96%	98%	98%	99%
Bottom of DS DOCSIS Spectrum (MHz)	204	300	300	492	492	1794	1794	1794	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	692	937	1023	1365	1445	3781	3943	4211	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1218	1794	6000	6000	6000	3000	3000	6000	6000	12000

**Figure 38 – Migration Path for Architecture 7b (Static Soft-FDX w/ Guard-band Elimination- Symmetric SLA)**

Dynamic SoftFDD Symmetric Tmax (US = DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	24	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+3	Node+2	Node+2	Node+2	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	800	480	480	480	120	120	120	120	120	120
Nsub in Service Group	200	200	200	200	120	60	60	4	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	63	105	105	105	417	417	417	417	417	417
US TavG (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
US Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	1056	2079	2112	4159	4135	10096	10135	10013	20018	20026	40036	40051	80073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	5.30%	3.81%	5.32%	3.82%	3.27%	0.95%	1.34%	0.13%	0.09%	0.13%	0.09%	0.13%	0.09%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	6.9	6.9	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	146	282	287	560	556	1468	1474	1340	2674	2675	5343	5345	10681
Top of US Band (MHz)	204	300	300	492	492	1794	1794	1794	3000	3000	6000	6000	12000
DS TavG (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	7318	7846	13448	15738	10637	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	45.34%	49.02%	25.64%	36.46%	5.99%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	3	3	3	3	3	3	3	5	5	5	5	5	5
DS Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.6	9.6	7.1	7.1	9.6	9.6	9.6	9.6	9.6	9.6
US:DS Frequency Band Re-Use BW (MHz)... Positive = Re-Use & Negative = Guard-band	96	192	192	384	384	1686	1686	1686	2892	2892	5892	5892	11892
Re-Use Ratio (Re-Use BW/Top of Spectrum)	8%	16%	16%	32%	32%	56%	56%	94%	96%	96%	98%	98%	99%
Bottom of DS DOCSIS Spectrum (MHz)	108	108	108	108	108	108	108	108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + 56 Annex B Video QAMs (MHz)	632	797	884	1292	1347	2652	2975	1866	3163	3236	5943	6060	11443
Tap BW (MHz)	1218	1218	1218	1218	1218	3000	3000	1794	3000	3000	6000	6000	12000

**Figure 39 – Migration Path for Architecture 8b (Dynamic Soft-FDX Baseline- Symmetric SLA)**

Static SoftFDD Symmetric Tmax (US = DS) with 1 Guardband	2020	2022	2024	2026	2028	2030	2032	2034	2036	2038	2040	2042	2044
Node Type (# SGs/Node)	2	2	2	2	2	4	4	4	24	24	24	24	24
Fiber Depth	Node+3	Node+3	Node+3	Node+2	Node+1	Node+1	Node+1	Node+1	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap	Fiber2Tap
# Homes Passed in Node	800	800	800	480	240	240	240	240	120	120	120	120	120
Nsub in Service Group	200	200	200	120	60	30	30	30	4	4	4	4	4
# Nodes to Support a 50,000 HHP Market	63	63	63	105	209	209	209	209	417	417	417	417	417
US Tavq (Mbps) w/ 19% CAGR	0.3	0.4	0.6	0.8	1.1	1.6	2.3	3.2	4.5	6.4	9.1	12.9	18.2
US_Tmax_max:DS_Tmax_max Ratio	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
US_Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd US HSD BW Capacity w/ K=1.0 (Mbps)	1056	2079	2112	4095	4068	10048	10068	10096	20018	20026	40036	40051	80073
US HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	5.30%	3.81%	5.32%	2.33%	1.66%	0.48%	0.67%	0.95%	0.09%	0.13%	0.09%	0.13%	0.09%
US Spectral Efficiency (bps/Hz)	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5
Bottom of US DOCSIS Spectrum (MHz)	5	5	5	5	5	5	5	5	5	5	5	5	5
Top of US DOCSIS Spectrum w/ 7.5 bps/Hz (MHz)	146	282	287	551	547	1345	1347	1351	2674	2675	5343	5345	10681
Top of US Band (MHz)	204	300	300	492	492	1794	1794	1794	3000	3000	6000	6000	12000
DS Tavq (Mbps) w/ CAGR of 39% in 2020s, 29% in 2030s, 19% in 2040s	2.3	4.4	8.6	16.6	32.1	57.5	95.6	159.2	264.8	440.7	676.6	958.1	1356.7
Natural DS Tmax_max (Mbps) w/ CAGR 25% in 2020s, 15% in 2030s, 15% in 2040s	1000	1563	2441	3815	5960	8568	11331	14986	19819	26210	34663	45842	60626
"Rounded-off" DS Tmax_max (Mbps)	1000	2000	2000	4000	4000	10000	10000	10000	20000	20000	40000	40000	80000
Req'd DS HSD BW Capacity w/ K=1.0 (Mbps)	1460	2889	3717	5991	5923	11724	12869	14775	21059	21763	42706	43832	85427
DS HSD Utilization (Nsub*Tavg)/(Nsub*Tavg+1.0*Tmax)	31.51%	30.77%	46.20%	33.23%	32.47%	14.71%	22.29%	32.32%	5.03%	8.10%	6.34%	8.74%	6.35%
1=Stat_FDD_GB 2=Stat_FDD_EC 3=Dyn_FDD_GB 4=Dyn_FDD_EC 5=FDX?	1	1	1	1	1	1	1	1	5	5	5	5	5
US Spectral Efficiency (bps/Hz)	9.6	9.6	9.6	9.3	9.3	8.0	8.0	8.0	9.6	9.6	9.6	9.6	9.4
US:DS Frequency Band Re-Use (MHz)... Positive = Re-Use & Negative = Guard-band	-36	-53	-53	-86	-86	-314	-314	-314	2892	2892	5892	5892	11892
Re-Use Ratio (Re-Use/Top of Spectrum)	-3%	-4%	-4%	-5%	-5%	-5%	-5%	-5%	96%	96%	98%	98%	99%
Bottom of DS DOCSIS Spectrum (MHz)	240	353	353	578	578	2108	2108	2108	108	108	108	108	108
Top of DS DOCSIS Spectrum w/ 9.0 bps/Hz + S6 Annex B Video QAMs (MHz)	728	989	1076	1558	1551	3909	4053	4291	2638	2711	4893	5010	9532
Tap BW (MHz)	1218	1218	1218	1794	1794	6000	6000	6000	3000	3000	6000	6000	12000

**Figure 40 – Migration Path for Architecture 9b (Static Soft-FDX w/ Active Taps-Symmetric SLA)**

## Conclusions

After constructing the eighteen Migration Paths within the previous section and after studying the general trends, several observations and conclusions could be developed.

With such a large amount of data available to the authors in the various Migration Paths above, some analytics tools that filtered out un-interesting years were developed to help pull insights from the data. Some of the outputs from those tools are illustrated in Figures 40-57 (Asymmetric SLAs in Figures 40-48, and Symmetric SLAs in Figures 49-57). Only the important years with “Big Changes” are high-lighted for each of the eighteen Architectures.

1a) FDX Asymmetric Tmax (US = 50% of DS) with 0 Guardbands (EC)	2020	2022	2025	2029	2035	2039	2044
US:DS Transmission Technology	FDX						
Life-span of US:DS Transmission Technology	25						
Fiber Depth	Node+0			Fiber2Tap			
Life-span of Fiber Depth (Years)	15			10			
US BW Transitions (MHz)	85	204	300	684	1794	3000	6000
DS BW Transitions (MHz)	1218			1794	3000	6000	12000
Snapped DS Housing Transitions (MHz)	1218			6000			12000
Life-span of Housing (Years)	9			15			1

**Figure 41 – “Big Changes” for Architecture 1a (Traditional-FDX- Asymmetric SLA)**

2a) Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2022	2025	2026	2028	2029	2035	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX				
Life-span of US:DS Transmission Technology	15				10				
Fiber Depth	Node+3		Node+2		Node+1		Fiber2Tap		
Life-span of Fiber Depth (Years)	6		2		7		10		
US BW Transitions (MHz)	85	204	300			684	1794	3000	6000
DS BW Transitions (MHz)	1218					3000		6000	12000
Snapped DS Housing Transitions (MHz)	1218					6000			12000
Life-span of Housing (Years)	9					15			1

**Figure 42 – “Big Changes” for Architecture 2a (Static Soft-FDX Baseline w/ 1.2 GHz Affinity- Asymmetric SLA)**

3a) Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2022	2025	2026	2029	2031	2032	2034	2035	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX						
Life-span of US:DS Transmission Technology	15				10						
Fiber Depth	Node+3			Node+2			Node+1		Fiber2Tap		
Life-span of Fiber Depth (Years)	12			2			1		10		
US BW Transitions (MHz)	85	204	300		684				1794	3000	6000
DS BW Transitions (MHz)	1218			1794	3000	6000			3000	6000	12000
Snapped DS Housing Transitions (MHz)	1218			6000							12000
Life-span of Housing (Years)	6			18							1

**Figure 43 – “Big Changes” for Architecture 3a (Static Soft-FDX w/ Node+3 Affinity - Asymmetric SLA)**

4a) Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2023	2028	2032	2033	2036	2040
US:DS Transmission Technology	St FDD_GB				FDX		
Life-span of US:DS Transmission Technology	16				9		
Fiber Depth	Node+3		Node+1		Fiber2Tap		
Life-span of Fiber Depth (Years)	8		8		9		
US BW Transitions (MHz)	85	204	300		684		1794
DS BW Transitions (MHz)	1218				1794	3000	1794
Snapped DS Housing Transitions (MHz)	1218				3000		3000
Life-span of Housing (Years)	12				13		

**Figure 44 – “Big Changes” for Architecture 4a (Static Soft-FDX w/ 15% DS Tmax CAGR- Asymmetric SLA)**

5a) Static SoftFDD Asymmetric Tmax (US = 40% of DS) with 1 Guardband	2020	2022	2025	2026	2028	2029	2035	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX				
Life-span of US:DS Transmission Technology	15				10				
Fiber Depth	Node+3		Node+2		Node+1		Fiber2Tap		
Life-span of Fiber Depth (Years)	6		2		7		10		
US BW Transitions (MHz)	85	204	300	204		492	1794	3000	6000
DS BW Transitions (MHz)	1218					3000		6000	12000
Snapped DS Housing Transitions (MHz)	1218					6000			12000
Life-span of Housing (Years)	9					15			1

**Figure 45 – “Big Changes” for Architecture 5a (Static Soft-FDX w/ Reduced US Tmax- Asymmetric SLA)**



6a) Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband & Sel Sub Mig	2020	2025	2028	2029	2032	2033	2036	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX				
Life-span of US:DS Transmission Technology	16				9				
Fiber Depth	Node+3		Node+1			Fiber2Tap			
Life-span of Fiber Depth (Years)	8		8			9			
US BW Transitions (MHz)	85	204		300		684		1794	3000
DS BW Transitions (MHz)	1218				1794	3000	1794	3000	6000
Snapped DS Housing Transitions (MHz)	1218				3000				6000
Life-span of Housing (Years)	12				12				1

**Figure 46 – “Big Changes” for Architecture 6a (Static Soft-FDX w/ Selective Subscriber Migration- Asymmetric SLA)**

7a) Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 0 Guardbands (EC)	2020	2022	2025	2027	2029	2035	2039	2044
US:DS Transmission Technology	St FDD_EC				FDX			
Life-span of US:DS Transmission Technology	15				10			
Fiber Depth	Node+3			Node+1		Fiber2Tap		
Life-span of Fiber Depth (Years)	7			8		10		
US BW Transitions (MHz)	85	204	300		684	1794	3000	6000
DS BW Transitions (MHz)	1218				3000		6000	12000
Snapped DS Housing Transitions (MHz)	1218				6000			12000
Life-span of Housing (Years)	9				15			1

**Figure 47 – “Big Changes” for Architecture 7a (Static Soft-FDX w/ Guard-band Elimination- Asymmetric SLA)**

8a) Dynamic SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband	2020	2022	2025	2028	2029	2031	2035	2039	2044
US:DS Transmission Technology	Dyn FDD_GB				FDX				
Life-span of US:DS Transmission Technology	15				10				
Fiber Depth	Node+3			Node+2		Fiber2Tap			
Life-span of Fiber Depth (Years)	8			7		10			
US BW Transitions (MHz)	85	204	300		684		1794	3000	6000
DS BW Transitions (MHz)	1218				1794	3000		6000	12000
Snapped DS Housing Transitions (MHz)	1218				6000				12000
Life-span of Housing (Years)	9				15				1

**Figure 48 – “Big Changes” for Architecture 8a (Dynamic Soft-FDX Baseline- Asymmetric SLA)**

9a) Static SoftFDD Asymmetric Tmax (US = 50% of DS) with 1 Guardband & Active Taps	2020	2022	2025	2028	2029	2035	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX			
Life-span of US:DS Transmission Technology	19				6			
Fiber Depth	Node+3		Node+2		Node+1	Fiber2Tap		
Life-span of Fiber Depth (Years)	5		3		11	6		
US BW Transitions (MHz)	85	204	300		684	1794	3000	6000
DS BW Transitions (MHz)	1218				3000	6000		12000
Snapped DS Housing Transitions (MHz)	1218				6000			12000
Life-span of Housing (Years)	9				15			1

**Figure 49 – “Big Changes” for Architecture 9a (Static Soft-FDX w/ Active Taps- Asymmetric SLA)**

1b) FDX Symmetric Tmax (US = DS) with 0 Guardbands (EC)	2020	2022	2025	2029	2035	2039	2044
US:DS Transmission Technology	FDX						
Life-span of US:DS Transmission Technology	25						
Fiber Depth	Node+0			Fiber2Tap			
Life-span of Fiber Depth (Years)	15			10			
US BW Transitions (MHz)	204	300	492	1794	3000	6000	12000
DS BW Transitions (MHz)	1218			1794	3000	6000	12000
Snapped DS Housing Transitions (MHz)	1218			6000			12000
Life-span of Housing (Years)	9			15			1

**Figure 50 – “Big Changes” for Architecture 1b (Traditional FDX- Symmetric SLA)**

2b) Static SoftFDD Symmetric Tmax (US = DS) with 1 Guardband	2020	2022	2025	2029	2031	2034	2035	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX				
Life-span of US:DS Transmission Technology	14				11				
Fiber Depth	Node+3		Node+2		Node+1		Fiber2Tap		
Life-span of Fiber Depth (Years)	5		6		3		11		
US BW Transitions (MHz)	204	300	492	1794			3000	6000	12000
DS BW Transitions (MHz)	1218			1794	6000		1794	3000	6000
Snapped DS Housing Transitions (MHz)	1218			6000					12000
Life-span of Housing (Years)	5			19					1

**Figure 51 – “Big Changes” for Architecture 2b (Static Soft-FDX Baseline w/ 1.2 GHz Affinity- Symmetric SLA)**

3b) Static SoftFDD Asymmetric Tmax (US = DS) with 1 Guardband	2020	2022	2025	2028	2029	2031	2034	2035	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX					
Life-span of US:DS Transmission Technology	14				11					
Fiber Depth	Node+3			Node+2		Node+1	Fiber2Tap			
Life-span of Fiber Depth (Years)	9			2		3	11			
US BW Transitions (MHz)	204	300	492	684	1794			3000	6000	12000
DS BW Transitions (MHz)	1218		1794	3000	6000		1794	3000	6000	12000
Snapped DS Housing Transitions (MHz)	1218		6000							12000
Life-span of Housing (Years)	5		19							1

**Figure 52 – “Big Changes” for Architecture 3b (Static Soft-FDX w/ Node+3 Affinity - Symmetric SLA)**

4b) Static SoftFDD Symmetric Tmax (US = DS) with 1 Guardband	2020	2023	2028	2033	2040
US:DS Transmission Technology	St FDD_GB		FDX		
Life-span of US:DS Transmission Technology	13		12		
Fiber Depth	Node+3		Node+1		Fiber2Tap
Life-span of Fiber Depth (Years)	8		5		12
US BW Transitions (MHz)	204	300	492	1794	3000
DS BW Transitions (MHz)	1218		1794		3000
Snapped DS Housing Transitions (MHz)	1218		3000		
Life-span of Housing (Years)	8		17		

**Figure 53 – “Big Changes” for Architecture 4b (Static Soft-FDX w/ 15% DS Tmax CAGR- Symmetric SLA)**

5b) Static SoftFDD ~Symmetric Tmax (US = 90% of DS) with 1 Guardband	2020	2022	2025	2028	2029	2034	2035	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX				
Life-span of US:DS Transmission Technology	14				11				
Fiber Depth	Node+3		Node+2		Node+1	Fiber2Tap			
Life-span of Fiber Depth (Years)	5		3		6	11			
US BW Transitions (MHz)	204	300	492		1794		3000	6000	12000
DS BW Transitions (MHz)	1218		1794		6000	1794	3000	6000	12000
Snapped DS Housing Transitions (MHz)	1218		6000						12000
Life-span of Housing (Years)	5		19						1

**Figure 54 – “Big Changes” for Architecture 5b (Static Soft-FDX w/ Reduced US Tmax-Symmetric SLA)**

6b) Static SoftFDD Symmetric Tmax (US = DS) with 1 Guardband & Sel Sub Mig	2020	2025	2028	2029	2033	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX		
Life-span of US:DS Transmission Technology	13				12		
Fiber Depth	Node+3		Node+1		Fiber2Tap		
Life-span of Fiber Depth (Years)	8		5		12		
US BW Transitions (MHz)	204	300		492	1794	3000	6000
DS BW Transitions (MHz)	1218			1794		3000	6000
Snapped DS Housing Transitions (MHz)	1218			3000			6000
Life-span of Housing (Years)	9			15			1

**Figure 55 – “Big Changes” for Architecture 6b (Static Soft-FDX w/ Selective Subscriber Migration- Symmetric SLA)**

7b) Static SoftFDD Symmetric Tmax (US = DS) with 0 Guardbands (EC)	2020	2022	2025	2028	2029	2035	2039	2044
US:DS Transmission Technology	St FDD_EC				FDX			
Life-span of US:DS Transmission Technology	15				10			
Fiber Depth	Node+3		Node+1		Fiber2Tap			
Life-span of Fiber Depth (Years)	5		10		10			
US BW Transitions (MHz)	204	300	492		1794	3000	6000	12000
DS BW Transitions (MHz)	1218			1794	6000	3000	6000	12000
Snapped DS Housing Transitions (MHz)	1218			6000				12000
Life-span of Housing (Years)	8			16				1

**Figure 56 – “Big Changes” for Architecture 7b (Static Soft-FDX w/ Guard-band Elimination- Symmetric SLA)**

8b) Dynamic SoftFDD Symmetric Tmax (US = DS) with 1 Guardband	2020	2022	2025	2027	2029	2034	2035	2039	2044
US:DS Transmission Technology	Dyn FDD_GB				FDX				
Life-span of US:DS Transmission Technology	14				11				
Fiber Depth	Node+3		Node+2		Fiber2Tap				
Life-span of Fiber Depth (Years)	7		7		11				
US BW Transitions (MHz)	204	300	492		1794		3000	6000	12000
DS BW Transitions (MHz)	1218				3000	1794	3000	6000	12000
Snapped DS Housing Transitions (MHz)	1218				6000				12000
Life-span of Housing (Years)	9				15				1

**Figure 57 – “Big Changes” for Architecture 8b (Dynamic Soft-FDX Baseline- Symmetric SLA)**

9b) Static SoftFDD Symmetric Tmax (US = DS) with 1 Guardband & Active Taps	2020	2022	2025	2028	2029	2035	2039	2044
US:DS Transmission Technology	St FDD_GB				FDX			
Life-span of US:DS Transmission Technology	15				10			
Fiber Depth	Node+3		Node+2		Node+1		Fiber2Tap	
Life-span of Fiber Depth (Years)	5		3		7		10	
US BW Transitions (MHz)	204	300	492		1794	3000	6000	12000
DS BW Transitions (MHz)	1218		1794		6000	3000	6000	12000
Snapped DS Housing Transitions (MHz)	1218		6000					12000
Life-span of Housing (Years)	5		19		1			

**Figure 58 – “Big Changes” for Architecture 9b (Static Soft-FDX w/ Active Taps-Symmetric SLA)**

Several key observations can be deduced from the results and are outlined below. These observations assume that the listed assumptions for each Architecture remain valid until the future 2044 time-frame and that the required technologies can be developed within the required time-frames.

- Node+Non-Zero and Node+0 Life-spans:
  - FDX solutions can support the MSO Bandwidth Capacity requirements using traditional Node+0 HFC networks until the ~2035 time-frame (requiring FTTT or FTTH solutions after that time-frame)
  - Static Soft-FDX solutions can support the MSO Bandwidth Capacity requirements using traditional Node+Non-Zero HFC networks until the ~2034-2035 time-frame (requiring FTTT or FTTH solutions after that time-frame)
  - Dynamic Soft-FDX solutions can support the MSO Bandwidth Capacity requirements using traditional Node+Non-Zero HFC networks until the ~2034-2035 time-frame (requiring FTTT or FTTH solutions after that time-frame)
  - FTTT or FTTH solutions are probably not required until the mid-2030 time-frame; greater Tmax CAGRs will cause the transition to FTTT or FTTH to occur sooner; smaller Tmax CAGRs will cause the transition to FTTT or FTTH to occur later; some MSOs may opt to move to FTTH sooner than required
  - Active Taps can extend the life-span of the Node+X HFC solutions and delay the deployment of FTTT/FTTH solutions by 1-4 years (depending on traffic statistics)
- Frequency Requirements:
  - Frequency Band changes are closely correlated to Tmax changes
  - For all Architectures (excluding the 15% Tmax CAGR & Selective Subscriber Migration solutions), Ultra-Split US frequencies (300+ MHz) will be required by 2025
  - For all Architectures (excluding the 15% Tmax CAGR & Selective Subscriber Migration solutions), Extended Spectrum DOCSIS frequencies (1794, 3000, 6000, and 12000 MHz) will be required at various times for both the US & DS
  - For Traditional-FDX Architectures, Asymmetric SLAs permit 1218 MHz DS operation until 2029, and Symmetric SLAs also permit 1218 MHz DS operation until 2029 (the benefit of overlapped DS & US spectra)
  - For Static Soft-FDX Baseline Architectures, Asymmetric SLAs permit 1218 MHz DS operation until 2029, but Symmetric SLAs end the life-span of 1218 MHz DS operation by as early 2025
  - For Dynamic Soft-FDX Architectures, Asymmetric SLAs permit 1218 MHz DS operation until 2029, and Symmetric SLAs also permit 1218 MHz DS operation until 2029 (the benefit of overlapped DS & US spectra)

- A 15% Tmax CAGR in the 2020's (instead of a 25% Tmax CAGR) permits 1218 MHz DS operation to work for ~3 extra years (2032 for Asymmetric SLAs and 2028 for Symmetric SLAs)
- Selective Subscriber Migration (eliminating the highest Tmax values) permits 1218 MHz DS operation to work for ~3-4 extra years (2032 for Asymmetric SLAs and 2029 for Symmetric SLAs)
- Guard-band Elimination can extend the life-span of 1218 MHz DS operation by 0-3 years (depending on traffic statistics)
- Staying in a Node+3 Architecture (and avoiding transitions to Node+2 or Node+1 or Node+0) forces the HFC network to transition to Ultra-Split US Extended Spectrum DOCSIS DS frequencies much more rapidly
- Other Interesting Findings:
  - Dynamic Soft-FDX keeps the total Node counts lower for the longest period of time (due to later required Node-splits)
  - Static Soft-FDX keeps the total Node counts lower for a medium period of time
  - Traditional-FDX keeps the total Node counts lower for the shortest period of time (due to earlier required Node-splits)
  - For Static Soft-FDX, the percentage of total DS spectrum that is unuseable Guard-band spectrum is quite small ( $\leq 5\%$  of the total) due to the fact that large Guard-bands are only needed when Extended Spectrum DOCSIS DSs are used
  - DS HSD Utilization levels are always less than or equal to 65%
  - US HSD Utilization levels are always less than 18%

Future work will continue to analyze the various network Architectures described above (plus other new ideas). However, these initial results and observations indicate that MSOs should be able to find one or more solutions that permit them to operate on Node+Non-Zero HFC networks or Node+0 HFC networks deep into the future. If the assumptions above are valid, then transitions to FTTT or FTTH architectures could be delayed until the mid-2030's if the Cable Industry decides to embrace Ultra-Split US and Extended Spectrum DOCSIS US and Extended Spectrum DOCSIS DS frequency ranges (and if their associated technologies can be developed).

For most MSOs, an important decision will likely need to be made between following a Node+0 Migration Path or following a Node+Non-Zero Migration Path. Both are fine paths, but the resulting investments and required technologies and upgrade steps on the two paths are quite different. The two paths do share some key technologies along the way; for example, both paths can make use of FDX-capable CMs, using them in slightly different operating modes. And both paths can make use of Extended Spectrum DOCSIS capabilities to support the extremely high Tmax values of the future. But the Node+0 Migration Path is focused on the use of Traditional-FDX, whereas the Node+Non-Zero Migration Path is focused on the use of either Static Soft-FDX or Dynamic Soft-FDX. The underlying goal behind the Node+Non-Zero approach is to delay the need to move to Node+0 and delay the associated costs.

As described above, Dynamic Soft-FDX is an interesting blend between the simplicity of Static Soft-FDX and the bandwidth savings (due to overlapped US & DS spectral) of Traditional-FDX; it is a simplified form of FDX that has the same benefits of re-using over-lapping frequency ranges for both Upstream and Downstream transmissions, but it does not require Echo Cancellation (which permits it to work in a Node+Non-Zero environment with Amplifiers). Thus, both Traditional-FDX and Dynamic Soft-FDX offer similarly efficiencies, but Traditional-FDX may require a Node+0 environment whereas Dynamic Soft-FDX may permit operation within a Node+Non-Zero environment.

It should be clearly stated and understood that many of the advanced technologies described in this forward-looking paper do not yet exist, and their ultimate performance levels are still conjecture. This includes Traditional-FDX, Static Soft-FDX, Dynamic Soft-FDX, Ultra-Split Upstreams, Extended Spectrum DOCSIS Upstreams, Extended Spectrum DOCSIS Downstreams, Active Taps, FTTT, and Distributed Node Architectures. Active research is still on-going in all of these key technology areas. If road-blocks are encountered that preclude some of the technologies, then MSO's may need to move towards FTTT or FTTH solutions sooner than predicted within this paper.

In the opinion of the authors, all three of the different Architectures (Static Soft-FDX, Dynamic Soft-FDX, and Traditional-FDX) will likely find applications in the future evolution of the MSO's HFC plant. However, since it is the "new kid" on the block with some very interesting attributes, the authors recommend that many Node+Non-Zero MSOs should at least consider the benefits of Dynamic Soft-FDX architectures (using ESD) as a way to greatly extend the life-span of their Node+Non-Zero HFC plants while maintaining simpler Amplifier solutions. Later upgrades to Node+0 Traditional-FDX architectures and/or Active Tap solutions and/or DNA-based FTTT solutions and/or FTTH solutions may make sense as subsequent steps.

# Abbreviations

bps	bits per second
BW	bandwidth
CAGR	Compound Annual Growth Rate
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Customer Premise Equipment
DAA	Distributed Access Architecture
dB	decibel
dBmV	decibel (relative to a millivolt)
DNA	Distributed Node Architecture
DOCSIS	Data Over Cable System Interface Specification
DS	Downstream
ESD	Extended Spectrum DOCSIS
FDD	Frequency Division Duplex
FDX	Full Duplex DOCSIS
FEC	Forward Error Correction
FPGA	Field Programmable Gate Array
FTTH	Fiber-To-The-Home
FTTLA	Fiber-To-The-Last-Active
FTTN	Fiber-To-The-Node
FTTT	Fiber-To-The-Tap
HDX	Half Duplex
HFC	Hybrid Fiber-coax
HHP	Households Passed
HSD	High Speed Data
Hz	hertz
IG	Interference Group
ISBE	International Society of Broadband Experts
K	K value (describing the QoE Coefficient)
LDPC	Low Density Parity Check
MAC	Media Access Control
MSO	Multiple System Operator
Nsub	Number of subscribers
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
OOB	Out Of Band
OSP	Outside Plant
pCore	Physical Core
PHY	Physical Layer
PON	Passive Optical Network
PSD	Power Spectral Density
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
RBA	Resource Block Assignment
RFoG	RF over Glass
RMC	Remote MAC Core

RMD	Remote MACPHY Device
R-OLT	Remote Optical Line Termination
RPD	Remote PHY Device
RxD	Remote MACPHY Device or Remote PHY Device
SC-QAM	Single Carrier- Quadrature Amplitude Modulation
SCTE	Society of Cable Telecommunications Engineers
SDN	Software Defined Network
SG	Service Group
SLA	Service Level Agreement
SNR	Signal-to-Noise Ratio
sub	subscriber
TCP	Total Composite Power
TDD	Time Division Duplex
Tavg	Average Throughput
Tmax	Maximum Throughput
US	Upstream
vCore	Virtualized Core
W	watt
WDM	Wavelength Division Multiplexing



# Bibliography & References

- [AL19] A. Al-Banna et. al., “Operational Considerations and Configurations for FDX & Soft FDD,” SCTE Cable-Tec 2019, SCTE
- [CL14] T. J. Cloonan et. al., “Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements,” SCTE Cable-Tec 2014, SCTE
- [CL15] T. J. Cloonan et. al., “Lessons from Telco and Wireless Providers: Extending the Life of the HFC Plant with New Technologies,” NCTA Spring Technical Forum 2015, NCTA
- [CL16] 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
- [MU16] V. Mutalik et. al., “Cable’s Success Is In Its DNA,” SCTE Cable-Tec 2016, SCTE
- [NI98] J. Nielson, “Nielson’s Law of Internet Bandwidth,” <https://www.nngroup.com/articles/law-of-bandwidth/>
- [UL19] J. Ulm et. al., “The Broadband Network Evolution continues – How do we get to Cable 10G?,” SCTE Cable-Tec 2019, SCTE

# **Proactive Customer Maintenance**

## **Going Beyond Proactive Network Maintenance With Data, Analytics, and a Laser Focus on Customer Experience**

A Technical Paper prepared for SCTE•ISBE by

**Andrew Joseph Milley**  
Sr. Manager, Technology Analytics  
Cox Communications Inc.  
6305 Peachtree Dunwoody Rd. Sandy Springs GA 30342  
andrew.milley@cox.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents .....	2
Introduction .....	3
Background .....	4
The ROCK (Resetting Obvious Channel Knockouts).....	6
Project NEO (Nightly Equipment Optimization).....	9
The BOOMSTICK (Kills Zombies, 'Nuff Said).....	12
Conclusion .....	14
Abbreviations.....	15
Bibliography & References .....	15

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - The ROCK Control Loop Diagram.....	7
Figure 2 - The ROCK Fixing a Modem With a Single Channel Impairment .....	7
Figure 3 - Regional implementation of the ROCK.....	8
Figure 4 - NEO Control Loop Diagram .....	10
Figure 5 - NEO Improvement to Downstream Receive Power Out of Spec .....	11
Figure 6 - NEO Improvement to Downstream SNR Out of Spec .....	11
Figure 7 - BOOMSTICK Control Loop Diagram.....	13
Figure 8 - BOOMSTICK Identification Model Comparing Zombies to a Control.....	13

# Introduction

The year is 2019 and service providers now have more data on every aspect of our business than ever before. Specifically, every piece of equipment in our network and customers' homes can be measured for performance and often issues can be solved or mitigated through remote action. It is not only interesting but is becoming necessary for service providers to leverage this data in new and novel ways to identify problems, make decisions and act proactively to solve problems. This capability to solve customer impacting equipment issues in real-time, or "proactive customer maintenance," is rapidly moving from a competitive advantage to business as usual.

Data analytics plays a key part in this process, which can include advanced modelling techniques like machine learning and artificial intelligence. However, the most effective solutions are often born from a combination of deep technological insight, rapid data enablement, and simple but powerful logic to empower the right action at the right time. Whether the output is a remote command, work recommendation, or even a proactive customer notification, speed to solution is the single most important aspect of a value-add system for proactive customer maintenance.

This paper will cover the general strategy of delivering proactive customer maintenance, that is, proactive and real-time efforts focused on providing ongoing excellent performance for every customer in our network, leveraging data and automation. It will cover the background and history of such efforts at Cox Communications Inc. (CCI) as well as current efforts. Finally, it will discuss the part that data analytics and emerging technologies plays in these efforts. Where noted some descriptions of these efforts has been left intentionally vague due to IP concerns, but concepts and results will be shown.

# Background

This section will provide some context on a few key concepts which will become important in the future sections of the paper.

Historically, and especially in the analog days, service providers had to rely on the customer to notify them of issues with their service. Multi-customer issues could be correlated from multiple customer contacts in common areas (i.e. outages), often via “buzz” in the operations center. Information would propagate through this system until the right fix agent was identified and sent. This process is high cost, incredibly inefficient, and difficult on customers. The desires to consolidate operations, reduce costs, and improve the customer experience have all driven the need for proactive monitoring and more timely notification of issues. The ultimate state for both the service provider and customer would be to identify every issue as, or even before, it causes an impairment and fix it proactively. In the cases where a physical technician is needed the desire is to identify the right person, the first time, at a time that works for the customer.

A control loop is a fundamental mechanism for maintaining a critical process value at a desired set point. A closed loop system is one that requires no human or manual intervention to maintain the process value, rather it is done automatically. There are three main components of a control loop: The sensor, the controller, and the actuator. Often a closed loop control system exists within a very small area, such as between equipment on a manufacturing floor. But the elements of a closed loop system can exist physically very far from one another and still meet the definition if their operation depends on each other and they operate automatically without human intervention. Many of the examples in this paper are examples of the latter, where the three elements of the control system all exist in different parts of the network but can still work together to maintain a critical process value.

Partial service mode is a DOCSIS resiliency measure where a modem is operating without the full set of channels defined in the Receive Channel Set (RCS) and/or Transmit Channel Set (TCS). A channel may be missing in this way either due to the modem being unable to acquire the channel during bonding, or because communication on that channel was lost during operation. This can either be due to a physical network impairment (like an RF trap) that must be removed by a technician or due to a temporary impairment caused by (for example) noise ingress. Partial service mode is considered a resiliency measure because the modem will continue to operate in this state but with a limited total bandwidth compared to what the modem can normally use (whether it is provisioned to use that bandwidth is a different matter altogether). Generally, when a channel is unusable, the modem and Cable Modem Termination System (CMTS) negotiate to not use that channel; however, in practice there have been observed states where the CMTS is sending frames to a downstream channel that a modem believes is not enabled, thus causing frame errors. In these cases, partial service mode is no longer a resiliency feature but is actively causing degraded service.

Modem performance is defined as the regularly measured operating state of several key performance indicators of a modem compared to their engineering specifications. Some examples of key performance indicators are DOCSIS Channel Downstream Receive Power (DS RX) and DOCSIS Channel Downstream Signal-to-Noise Ratio (DS SNR). DS RX is a measure of the received power in dB on each bonded downstream DOCSIS channel while DS SNR is a measure of the clarity of the signal as received by the modem. Both measurements have specifications as defined in the DOCSIS spec as well as (potentially tighter) specs defined by the outside plant and/or engineering teams. The intention of any proactive maintenance activity aimed at improving modem performance is to reduce the number of out of specification measurements in the network.

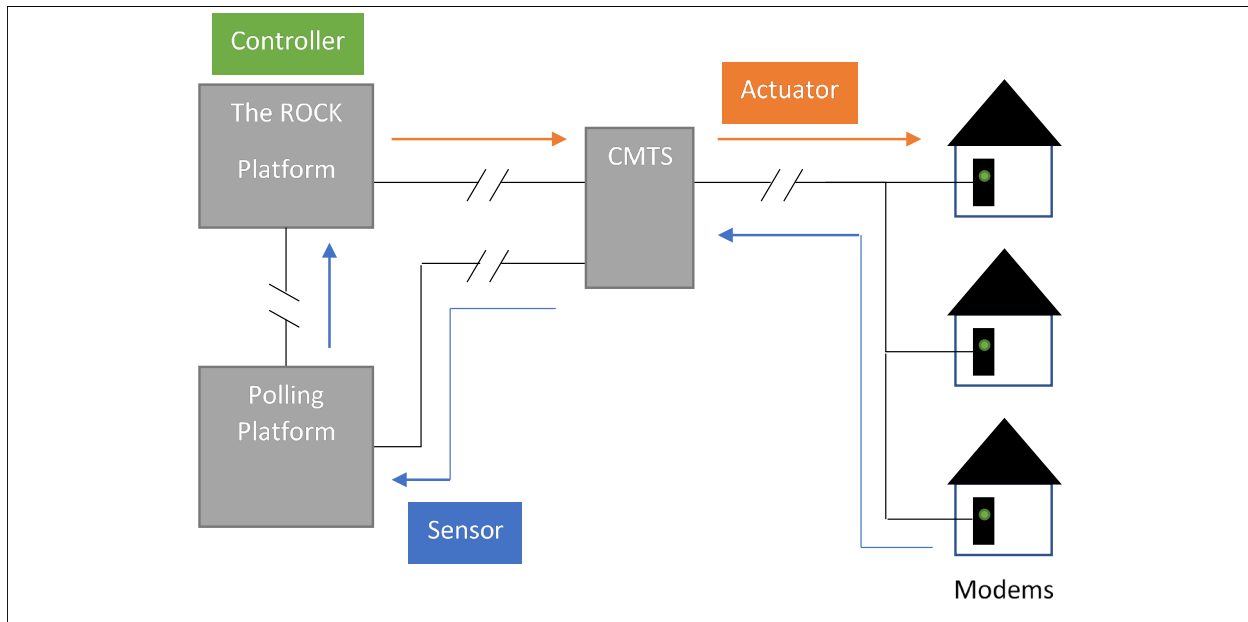
A zombie modem is a modem which is connected to the CMTS on layer 2 (MAC layer) but has lost connection on layer 3 (IP layer). To the CMTS, a modem in this state will appear to be connected but will lose all internet connectivity, including the ability to be polled over Simple Network Management Protocol (SNMP). To the customer, this will appear as a total loss of service. Depending on the cause of the layer 3 loss, connection may be able to be recovered with a hard reboot (for example, due to a software error or encryption key corruption in the modem) or it may not (for example, due to a mis-handled protocol change request from the CMTS to the modem). In both cases a zombie modem is very likely to cause customer impact, and in the latter case without proper identification will likely result in similar troubleshooting efforts to a totally inoperable modem (necessitating replacement).

# The ROCK (Resetting Obvious Channel Knockouts)

As noted in the background section, partial service mode is a resiliency feature of DOCSIS that in practice does not always behave as expected. In many cases, especially those involving channel knockouts due to noise ingress, missing channels do not get re-bonded after the event is cleared. This behavior may cause a service impact if the CMTS and modem have not properly negotiated the loss of that channel. Of principle concern are cases where the CMTS will continue to send frames to the modem over channels that the modem believes have been turned off.

The baseline rate of modems in partial service mode on CCI's network prior to conversion to Converged Cable Access Platform (CCAP) and with 32 downstream DOCSIS channels was between 1% and 1.5%. After conversion to the CCAP platform and expanding to 48 downstream DOCSIS channels, the rate jumped to between 3% and 4%. Experimentation has shown that rebooting these devices has no significant affect, however a series of re-initialization commands can be run on the CMTS that have a 60-80% chance of restoring a modem to its full channel bonding state. At scale this capability can be leveraged to bring the network back down near the baseline rate of 1-1.5% of modems in partial service.

The ROCK is a closed-loop control system that CCI uses to find and fix modems in partial service mode. The sensor part of the system is taking scheduled polling data from all devices at a 2-hour interval and identifying modems which are missing at least one channel. The control function waits until the most convenient time for the customer to experience a small service impact (this is usually at 4AM local time which is both in the maintenance window and the measured time of least usage/activity across CCI's network). The ROCK then iterates through all identified modems, issuing the necessary series of commands to the CMTS to re-initialize the modem. Any modems that do not recover are put in a holding list so that they are not continually affected every morning. This holdout list generally represents the baseline, are likely due to physical channel impairments, and can help prioritize proactive labor in the cable plant. Figure 1 below shows the control loop diagram for the ROCK system. Figure 2 shows a timeline of channel behavior on a single modem before and after being fixed. Channel 855, which is increasing frame errors (uncorrectables) at a steady rate goes from an impaired state to normal operation after action is taken.




**Figure 1 - The ROCK Control Loop Diagram**



**Figure 2 - The ROCK Fixing a Modem With a Single Channel Impairment**



The ROCK program was implemented over five weeks in April 2019 using a staggered regional roll out schedule. Figure 3 shows the immediate impact in each region after the program was started. The average regional improvement (excluding California, which had previously been piloted and started at a lower baseline) was 52%, with the total improvement by the end of the rollout being 50% or roughly 120,000 modems. The intention of the ROCK is to maintain the baseline of impaired modems; in parallel a data analytics effort is underway to understand the root cause of the fixable impairments and feed that information back to the device vendor(s) for a hopeful bug fix or feature add.



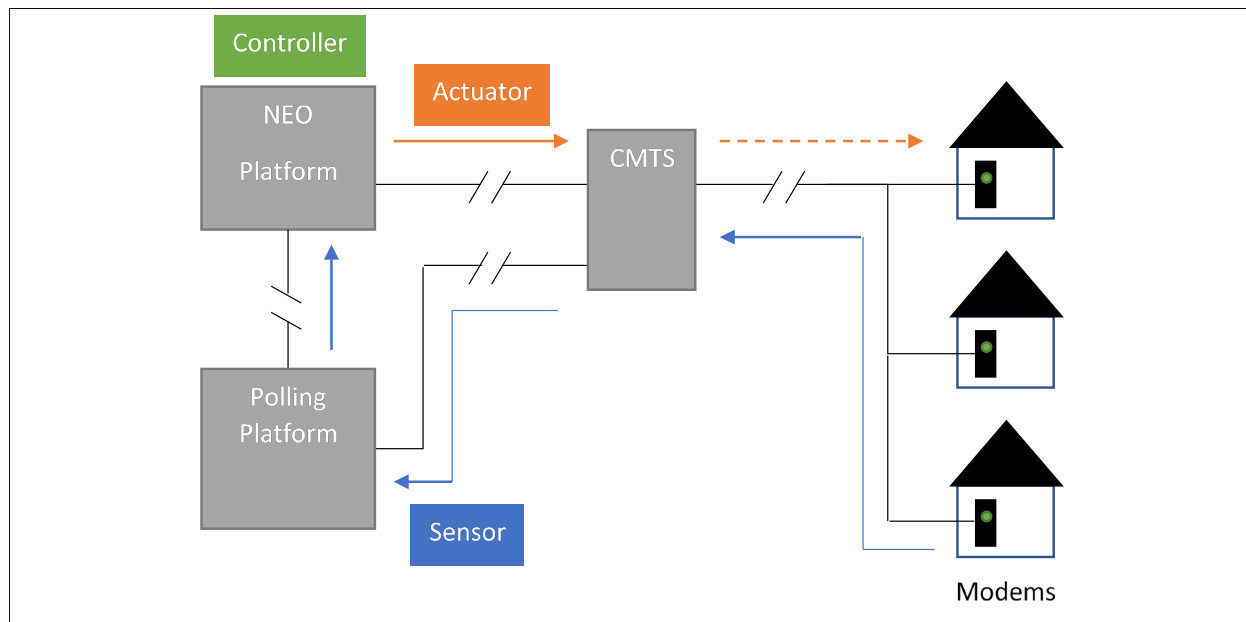
**Figure 3 - Regional implementation of the ROCK**

# Project NEO (Nightly Equipment Optimization)

As noted in the background section, a proven key performance measure of modem performance is DOCSIS Channel Downstream Receive Power (DS RX). Every bonded downstream channel must fall within a certain power range as given by the DOCSIS spec. Essentially, the higher the received power (up to a reasonable level) the better the modem will perform. DS RX at the modem is a function of several factors, but it is strongly correlated with transmit power at the edge quadrature amplitude modulator (edge QAM), which is configured through the CMTS. The CMTS does not have the ability to measure the receive power of any given modem let alone an aggregate, but it can be measured and aggregated from the modems themselves through scheduled polling. Due to non-ideal elements in the cable plant which change properties with temperature, drift over time, and generally can move the performance away from the designed center point, modem channel receive power will have a certain amount of variance over the course of a year and across the radio frequency (RF) spectrum.


Using scheduled DS RX polling data, the optimal transmit power for every channel can be calculated with the intention of bringing as many modems into specification as possible. Classically this would have to be tuned in the plant by a field technician which is difficult and prohibitively expensive. This adjustment can be made on a very fast basis, as quickly as DS RX data can be polled, or at a slower pace (monthly for example). The calculated adjustments are such that the total error of every channel is minimized across all modems actively bonded to that channel while also meeting some key requirements of the cable plant. The first is that no two adjacent channels can be more than 3dB of transmit power apart, and the second is that an individual channel cannot be adjusted more than an upper limit.

The result is an optimized cable plant without requiring physical adjustments to physical elements. This results in less variance over the spectrum and over the course of the year (due to temperature swings) which will result in a net improvement to customer experience and lower care and field services costs. Finally, it bubbles actual plant issues to the top which can then be proactively addressed, further optimizing the plant and allowing proactive work to be the most effective. Figure 4 shows the simplified control loop diagram of NEO.




**Figure 4 - NEO Control Loop Diagram**

Experimental results are shown below in Figures 5 and 6 where a small area of the network (roughly 1600 modems) was optimized using NEO. Prior to the change these modems were reporting roughly 52k out of spec DS RX measurements and 18k out of spec DS SNR measurements per poll. After the adjustment the same modems were reporting 45k and 15k out of spec measurements, respectively. This represented an improvement of 13% and 16% out of spec measurements in DS RX and DS SNR respectively. A hidden added benefit of the process is that by optimizing DS RX, there is also a benefit to DS SNR performance.



**Figure 5 - NEO Improvement to Downstream Receive Power Out of Spec**



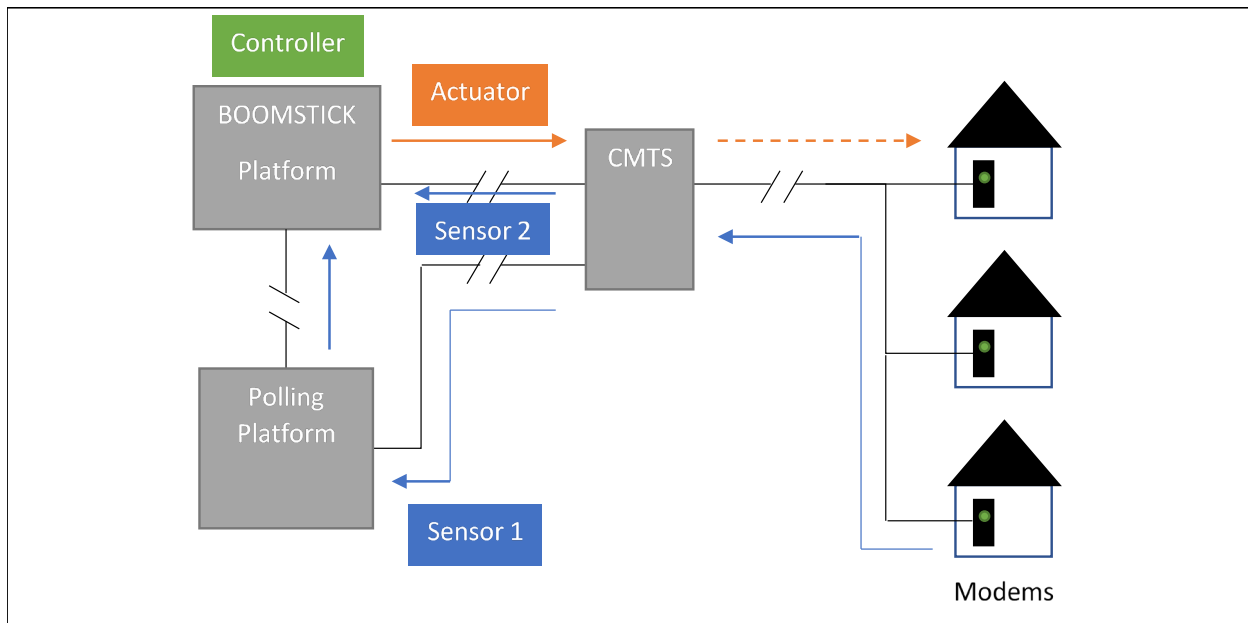
**Figure 6 - NEO Improvement to Downstream SNR Out of Spec**

# The BOOMSTICK (Kills Zombies, ‘Nuff Said)


As noted in the background section, a zombie modem is a device that has entered a state where it is still powered, is still connected to the CMTS on the MAC layer (layer 2) but has stopped responding on the IP layer (layer 3) due to an internal modem software crash or missed profile change from the CMTS. A modem reboot may solve the issue in some cases if the state is identified, but because the modem has stopped responding to IP requests it cannot be proactively reset over SNMP. However, there is action that can be taken on the CMTS side which will take the modem out of this state once identified. The added benefit of performing this action is that it will solve zombie states that are not fixable with a modem reboot (such as ones due to a missed protocol change request from the CMTS). Because the zombie state represents an immediate service impact, the desire is to identify and fix these modems as near real-time as possible to prevent or limit customer impact.

Identification of zombie state is key and cannot be done with a simple check. Conceptually, a modem not responding to an SNMP request but which still shows in a connected state on the CMTS is an obvious zombie, but SNMP polling at a high enough rate to detect these issues isn't feasible. Instead a combination of existing polled data can be used to significantly limit the list of suspected zombies which can then be polled for definitive proof before action is taken. The combination of data that can help identify these modems includes layer 2 connectivity and service flow usage data, both of which are polled from the CMTS.

Generally, a near real-time model can be created which checks for significant changes in both upstream and downstream usage for a modem while also confirming layer 2 connectivity. If the modem shows a significant change in upstream and downstream usage compared to nominal and is still showing connection on layer 2, the modem will be put on a list for a secondary check of layer 3 connectivity using a real-time SNMP poll. If this final check fails, the modem can be fixed using a series of commands on the CMTS which will re-initialize the modem into its proper operating state. Figure 7 shows the control loop diagram while Figure 8 shows an example of three identified zombies using this methodology when compared to a control (non-zombie) modem. A key calculation from the usage data will be shown in red for the three zombie modems, while the control modem never enters this condition.



**Figure 7 - BOOMSTICK Control Loop Diagram**



**Figure 8 - BOOMSTICK Identification Model Comparing Zombies to a Control**

# Conclusion

In conclusion, several examples of proactive customer maintenance systems have been introduced. These methods are intended to improve modem performance, and thus customer experience, by measuring key performance indicators on a regular or real-time basis, feeding those into models of varying complexity, and ultimately enabling or taking an action to fix the issue. These are proven methods that have shown significant improvement in the CCI network.

The impact of data engineering, advanced analytical, and modeling techniques on these efforts cannot be understated. However, the genesis of these systems has broadly been a methodology of identifying an opportunity through close partnership between the business and field, rapidly determining feasibility and what action to take, and then utilizing as much existing infrastructure and capability as possible to quickly stand up a solution. Often the modeling complexity is in understanding the problem, and less so in solving it.

## Abbreviations

CCI	Cox Communications Inc.
CCAP	Converged Cable Access Platform
CMTS	Cable Modem Termination System
DS RX	DOCSIS Downstream Channel Receive Power
DS SNR	DOCSIS Downstream Channel Signal-to-Noise Ratio
Edge QAM	edge quadrature amplitude modulator
NEO	nightly equipment optimization
RF	radio frequency
RCS	Receive Channel Set
ROCK	resetting obvious channel knockouts
SNMP	Simple Network Management Protocol
TCS	Transmit Channel Set

## Bibliography & References

B. Volpe, W. Miller, “Advanced Troubleshooting in a DOCSIS 3.0 Plant,” Cable-Tec Expo, Nov. 14, 2011. Accessed on: Jun. 24, 2019. [Online]. Available: [https://volpefirm.com/wp-content/uploads/2012/01/VM\\_Expo2011\\_v1-blog.pdf](https://volpefirm.com/wp-content/uploads/2012/01/VM_Expo2011_v1-blog.pdf)



# **Kickstarting Proactive Network Maintenance with the Proactive Operations Platform and Example Application**

**An easy way to start your own PNM journey**

A Technical Paper prepared for SCTE•ISBE by

**Jason Rupe, Ph.D.**

Principal Architect

CableLabs®

858 Coal Creek Circle

303.661.3332

j.rupe@cablelabs.com

**Jingjie Zhu**

Senior Engineer

CableLabs®

858 Coal Creek Circle

303.661.3312

j.zhu@cablelabs.com

# Table of Contents

Title	Page Number
Table of Contents .....	2
Introduction .....	3
The Proactive Operations Platform .....	4
1. Control & Schedule Worker .....	6
2. Work Queues .....	7
3. Workers and Workflow Description .....	7
4. Config Files .....	9
5. Data Stores .....	9
6. Scheduling Considerations .....	10
7. Additional Notes about ProOps .....	10
PNM Example Application for ProOps .....	11
1. Base data polling (Pollers) .....	11
2. Triggers (Analyzers) .....	11
3. Actions .....	12
4. Modules .....	12
Configuring an Application .....	13
Envisioned Use Cases .....	22
Conclusion .....	24
Abbreviations .....	25
Bibliography & References .....	25

## List of Figures

Title	Page Number
Figure 1 – A depiction of the ProOps platform on top of CCF and a network .....	4
Figure 2 – The depiction of ProOps from Figure 1, with OODA overlaid and arrows showing the process flow .....	6
Figure 3 – Depiction of the example application that comes with ProOps today .....	11
Figure 4 – Configuration management server web GUI index page (task queue configurations) .....	14
Figure 5 – Configuration management server – defined workers .....	14
Figure 6 – Job scheduling .....	15
Figure 7 – Configuring job scheduling .....	16
Figure 8 – Configuring for T3 and T4 device event log errors .....	17
Figure 9 – Configuring the RxMER statistics worker .....	17
Figure 10 – Configuring a data collection worker .....	18
Figure 11 – Channel estimation statistics worker .....	19
Figure 12 – Top-level system monitoring .....	19
Figure 13 – Detailed worker system monitoring .....	20
Figure 14 – System monitoring graph .....	21
Figure 15 – System monitoring graph of data collection .....	22

# Introduction

As part of its long standing proactive network maintenance (PNM) project, CableLabs® has been assessing the needs of operators and vendors in the area of PNM with the goal of reducing adoption friction for members and vendors. We identified a few main issues, an important one of which is addressed by the work presented in this paper: the Proactive Operations (ProOps) platform.

ProOps is a platform (environment, framework) for turning data into operations action. That includes proaction, when the data allow it. PNM data enables proaction, so we built an example application that comes with ProOps, which serves multiple purposes: as an example to show how to use ProOps, as a starting point for trying basic network data-driven PNM and reactive operations, and as a launch point for implementing and sharing PNM best practices.

To turn data into action, most operators rely on engineering and technician expertise. It is common to simply gather and plot the data, then look at the output. CableLabs built the cable modem validation application (CMVA) for that latter purpose (as well as for cable modem (CM) certification test automation and sharing). But without a human expert sifting through the data, not much can be done with it. CMVA is great for developing PNM ideas, but it still requires experts to do the next step, and developers to build solutions to try. That requires investment risk that we surmise is a roadblock to implementation of PNM. But for many operators, and some vendors, there just aren't enough available experts to do the work manually. The industry needs help getting over the hurdle of turning the data we exposed into action we can take with confidence.

ProOps was built to facilitate the automation of turning data into operations action. Generally, we identified the steps to accomplish that task as 1) data extraction (observation), 2) analysis across time and network elements (orient), 3) correlating problems and measuring severity (decide), and 4) defining work items that are worthy of attention (act). The steps can be labeled as observe, orient, decide, and act (OODA) to roughly follow the OODA process or OODA loop, which is a cyclic process developed by US Air Force Colonel John Boyd [1,2]. Combat operations resembles network operations more than we care to admit perhaps, so the labels fit. Boyd systematized the combat operations process as a rapid cycle of the OODA loop. Likewise, network operations follows a similar process, and the concept helps explain how ProOps works.

In the future, we expect to release applications and modules to enhance existing applications, in the ProOps environment. Once ProOps is installed, new applications will work like updates to ProOps, making the operations impact even less. New applications and modules will interwork with existing applications in the same deployment of ProOps, or parallel deployments utilizing the same common collection framework (CCF) instance are possible too. Multiple deployment models are available today, and more to come as vendor and operator members request.

ProOps is currently available for use by CableLabs members and vendors under nondisclosure agreement (NDA) and intellectual property rights (IPR), with the additional common code collection (C3) community agreement.

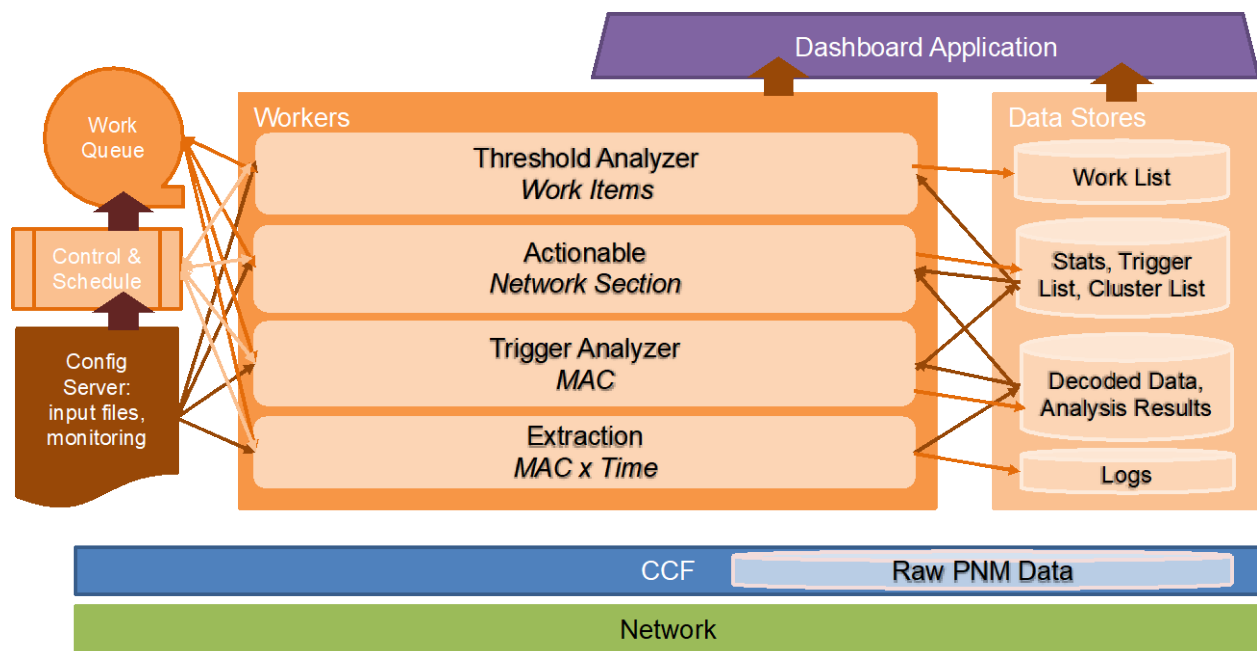
In the rest of this paper, we explain in-depth the structure and function of ProOps, and the example application that comes with it today. This will lead us to explain some of the basic ways you can configure and build your own solutions in ProOps to support your operations improvement experiments and inventions. We also cover use cases for ProOps that we envision, which hopefully will interest you or spur your own imagination to invent use cases we haven't thought of yet.

# The Proactive Operations Platform

The ProOps platform is a framework for constructing applications that turn network data into action. ProOps is simpler to understand in the context of an application, so we later explain it in the context of the simple but complete example application that comes with ProOps. As a foundation, we first explain the ProOps application environment as it stands alone, before an application has been constructed. We also explain briefly here the elements of the environment, covering the details later in the application context.

An important class of components of ProOps is the worker, which is a module of code which can be tasked by the system. Some workers are a part of and come with ProOps, and some workers function for the application. A worker is a module of code which is scheduled to execute according to the configuration instructions, taking inputs as directed, and sending its output where directed. In software development terms, it is like a microservice, but subscribed to a task queue instead of waiting for an application programming interface (API) call. An application in ProOps is configured from workers that execute the needed instructions.

Figure 1 shows the elements of the PNM application environment as a raw environment, without hosting a working application, or any workers; the workers box represents just the worker organization without the workers that make up the application.



**Figure 1 – A depiction of the ProOps platform on top of CCF and a network.**

We begin at the bottom of ProOps and work our way up. The descriptions of the network or CCF are outside the scope of this paper. For the latter, we refer the reader to the CCF architecture document [3]. The remaining elements are all a part of ProOps, so are described briefly here.

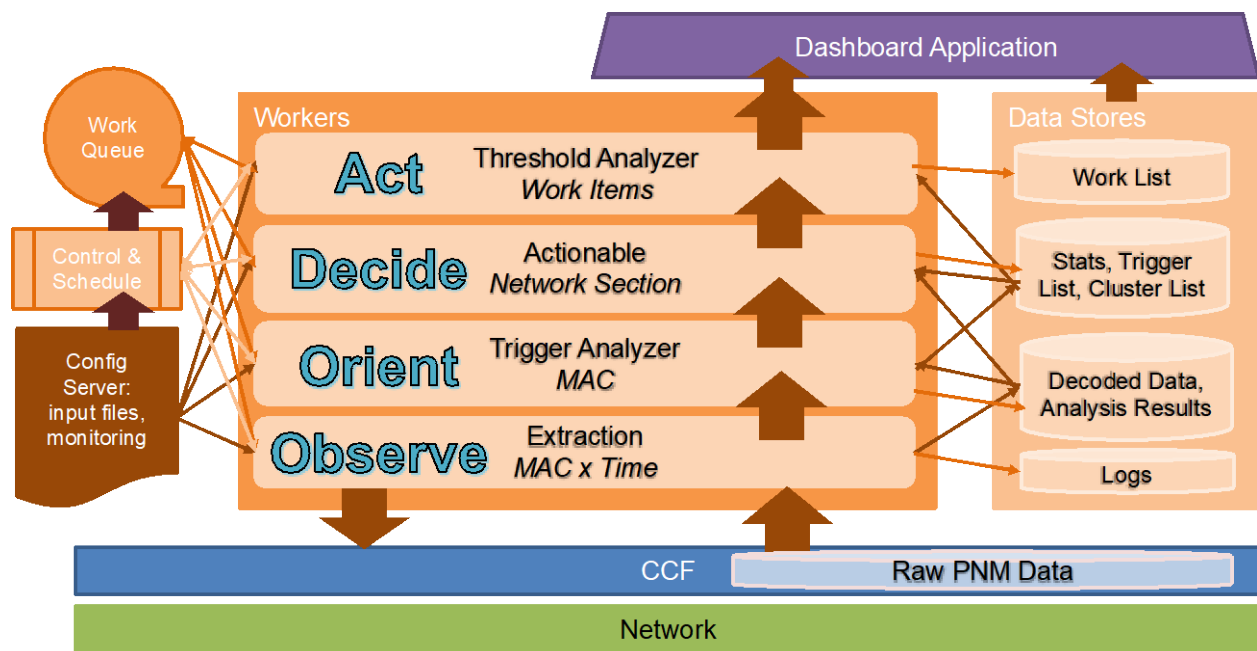
- The configuration server (config server) hosts the configuration data which define the messaging mechanisms and necessary additional configurations for workers. It describes the workers to be connected, and any changes that affect the work they do. The user can control applications through this interface for the most part, as the configuration is what defines the application from a

pool of workers. When workers start up, they request configurations from the configuration server using RESTful APIs.

- The control and schedule section handles the scheduling of workers, and controls task creation and ending as needed to handle the processing according to the configuration server. The tasks describe the data that each worker is fed, the data that each worker outputs, where the outputs are held, where the inputs are obtained from, and any additional information for specific workers such as thresholds, Internet protocol (IP) addresses of cable modem termination systems (CMTSs) and CCFs, etc.
- The work queues are a number of first-in, first-out (FIFO) queues of tasks that need to be done by the workers according to the schedule dictated by the configuration server.
- The workers are held in the worker environment in the center of the figure. These workers simply exist in a pool and are organized as dictated by the configuration server, but we recognize that for PNM we will need the data from the network to be translated into actionable work. Therefore, we have identified reasonable steps we expect are sufficient for most any application envisioned. These steps are part of the design advantage of ProOps, but not a constraint to the solution possibilities. Those steps are briefly described as follows. Note how they align to the well known OODA loop or OODA process.
  - **Extraction** takes place periodically as dictated by the initial schedule outlined in the configuration. Because the same media access control (MAC) addresses are collected multiple times on a(n) (approximate) schedule, we note that this step conducts data collection of MAC by time (**MAC x Time**). The output at this layer goes to the log file, but also decoded data and analysis results are held, too. Because this layer is focused on data extraction, we align this layer to the **Observe** step of the OODA process.
  - **Trigger Analyzer** will analyze the data from the extraction and determine whether a change of action is triggered, such as obtaining new data elements, taking data elements on a different schedule, collecting data from related network elements, and placing some of these MAC addresses on a list that identifies further action. The analysis on this layer likely evaluates data on a single MAC over time. For this reason, we note that this step conducts data analysis by **MAC**. The output from this step consists of statistics, and a determination of which MACs are triggered. Because this layer establishes some context to the information, we align this layer to the **Orient** step of the OODA process.
  - The **Actionable** step takes the MAC addresses that are flagged by the trigger analyzer and determines the network section relevant to the identified MACs. For example, clustering of MACs is possible here. But to make these actionable, we also have to introduce a measure of performance and comparison for these MACs or clusters of MACs, so that decisions can be made to bring the network section to actionable work. Therefore, we note that this step conducts analysis on the **network sections**. The output from this step can be clusters of MAC addresses, or single MAC addresses representing network points or CMs, but will definitely be additional statistics to support decisions. Because this layer makes our oriented view of the data something we can act on, it is really the layer in which decisions are made. Therefore, we align this layer to the **Decide** step of the OODA process.
  - The **Threshold Analyzer** will look at the measure of performance and comparison, introduce further analysis of other MACs in the cluster, look at historical information, and decide which of the identified actionable network sections are good choices to act on based on information provided in the configuration server. This step therefore conducts analysis on **work items**, and provides output to the work list. Because this layer considers the decision made that there is something that can be acted upon, and either selects or allows the selection of work that will be acted on, we align this layer to the **Act** step of the OODA process.

- The data stores hold the data elements needed for the application as dictated by the configuration server. The following elements are expected to be needed in many applications.
  - Logs will hold output from the extraction layer.
  - Decoded data and analysis results will be held in a data lake or database depending on the needs of the application.
  - Statistics, triggered lists, and clustering lists are captured in a database or tables as needed by the application.
  - The work list is likewise captured in a list or database as needed.
- Finally, a dashboard application is fed by the entire set of identified data, so that geographic information system (GIS) maps, graphs, and detailed work packages can be assembled as needed by whatever system this will feed. This approach allows most any operator to take the output from the various steps of an application in this environment and tie it easily into its existing ticketing systems and other operations needs.

The previous Figure 1 is revised now to show the flow of information and how the layers reflect the OODA steps, in Figure 2.



**Figure 2 – The depiction of ProOps from Figure 1, with OODA overlaid and arrows showing the process flow.**

Next, we discuss some of the detail in the sections of the ProOps architecture. In the explanation of the architecture, we put **workers** in bold, *data* in italics, and configuration underlined to help keep it all straight.

## 1. Control & Schedule Worker

**Control & schedule** is a special worker that interacts at the system level to control the other workers, follow the application and flow defined in the configuration file, and schedule work according to the task description or conditions discovered from the other workers or their data output. The tasks are posted periodically by the schedule worker to the task queue(s) and assigned to workers by subscription. If a

configuration instructs data collection to happen every  $x$  hours, then this worker will schedule work for the queue based on this schedule. Further, if a result from a data poll and analysis indicates triggering a rule given in the config file, then this worker will update the schedule it follows for those impacted CMs accordingly, and therefore change the schedule that dictates how it places work in the work queue. It must therefore hold state, but can always regain a lost state by reading the config data and the information in the data stores, as those are sufficient to define the schedule.

Note that if this worker is responsible for monitoring a trigger condition to instantiate another worker, such as a **ranker worker**, then it must have access to the data or worker output that triggers the follow up work. If instead it is to act on a schedule only, then maintaining the time against that schedule is important, but not catastrophic if state is lost. The *log file* could be used to capture time and a backup of **control & schedule** state to avoid a loss of state at failure.

This special worker also regulates the system so to not tax the network too much, and to poll and archive data based on scheduling requirements from the application configuration. **The control & schedule worker** reads the config data, builds a schedule for how other workers are started for data collection, then creates the needed polling workers at the appropriate time.

Note that the **control & schedule worker** is the one worker type that keeps track of timing; hence, it needs to maintain state information.

## 2. Work Queues

This entity is a number of queues of tasks to be handled by different workers in the system. The **control & schedule** worker will place tasks on the queues for other workers, and likewise another worker can generate a task for the **control & schedule** worker to handle. The work queues manage tasks and assignments, holding the tasks to do like a pipeline connecting tasks to workers.

## 3. Workers and Workflow Description

In this architecture, workers are created to handle tasks, and the work to be handled arrives to a queue. A special **control & schedule** worker handles all timing for scheduling, and assigns work to the other workers. Based on this concept, the following workers need to be defined, as illustrated in Figure 2. The previous figures only show in some cases the worker classes, so more detail is described here than can be shown in the figures. Further, other types of workers can be defined at will, and nothing precludes the addition of layers, or the extension of function of identified layers, in this architecture. The layering is simply a framework that facilitates action to support a network and services. The configuration file, described later, is what enforces the construct of the overall system, and therefore the workflows or applications built in the environment.

- Observe = Extraction [**polling worker, translator, linear calculator**]: This is a pool of workers which form a linear processing of data from the CCF. The new **polling workers** interact with the CCF to poll the data. The data received are then translated by **translator workers** if and when necessary, and finally processed by **linear calculator workers**. The linear calculator worker is a special calculator worker that is always applied to data right off of the CCF. The outputs are cataloged into the *decoded data and analysis results data lake or database* by a **catalog worker** (if needed).
- Orient = Trigger Analyzer [**calculator workers, anomaly detector workers, machine learning workers, trigger analyzer worker**]: This is a pool of workers who evaluate data as it shows up in the *decoded data and analysis results*, against the analyses invoked and configured by the config file, placing statistics and a list of triggered MAC addresses into the *Stats, Trigger List*,

*Cluster list database*. The **control & schedule worker** will monitor the database for new data which triggers the need for a specific worker type, or monitor a clock to trigger, based on the config file. Then it will schedule the appropriate workers to complete the configured tasks.

**Calculator workers** will simply calculate statistics from the raw data and place the statistics into the database. **Anomaly detector workers** will search a series of data (over frequency, time, etc.), indicate the data included in a detected anomaly, and in some cases indicate the class of anomaly detected. **Machine learning workers** will apply other machine learning techniques to data and statistics. Both the latter worker types may contain or use calculator workers for example, too. A **trigger analyzer** worker will assess these outputs against the triggers indicated in the config file. The output from these analyses are tasks that can be to sample new data, or sample some data sources for specific MACs more frequently, or to re-analyze older data, for example.

A note is warranted here about the role of the config file versus the changes decided: The config file holds information about the frequency of data collection, but only for the default data frequency; it also holds information about the triggers and actions from triggers regarding data polling frequency, but only as conditions, not as system state. Changes in system state are held in memory; or, in the case of MAC addresses subject to frequent polling, in the *stats*, *trigger list*, *cluster list* database.

This means workers can relate to all other workers in a nesting, series, parallel, or other configuration relationship. More worker types can be defined to fit into this pool as needed by specific applications.

- Decide = Actionable Layer [**severity calculator worker**, **ranker worker**, **cluster worker**]: This worker class likely only contains a few worker types (depending on the nature of the application configured), tasked with translating the trigger list and collected statistics and data for triggered MAC addresses, and providing the necessary performance measures with which to rank the MAC addresses. It can work exclusively with the *stats*, *trigger list database* as it reads the information, processes, then adds an updated measure of performance to the data, and provides an updated ranking. That assumes the configuration file is set up to provide the needed data to supply the measures needed for action. The work here can be done periodically, or triggered by a condition such as a number of data updates on the triggered list, or a number of new entries to the triggered list. The **control & schedule worker** will trigger action as directed by the config file, and trigger a **severity calculator worker** based on the information in the config file which specifies what triggers the calculation. Once the calculation is updated, the **ranker worker** will evaluate the updated measures of performance and rank the MAC addresses by severity, placing the performance measures and the ranked list back into the database. The **cluster worker** will evaluate the information relating to the MAC addresses and attempt to cluster them by similarities in the statistics and performance measures. Workers on this layer may need access to the raw data or decoded data or analysis results in some cases too. Either on a schedule or as triggered by conditions, as specified in the config file, the **cluster worker** is started to evaluate triggered MAC addresses and their performance and other needed information, conduct clustering based on multiple dimensions of data, and output clusters of MAC addresses that are likely experiencing the same problem.
- Act = Threshold Analyzer: The threshold analyzer will examine the trigger clusters (which are assumed to include clusters defined by single MAC addresses in many cases) and decide which of the clusters are actionable. This can be done through the config file which holds the rules that trigger the action based on an external financial model, operations rules, work load, number of potential jobs on the list, etc. The **control & schedule worker** will read the configuration file to determine when or under what condition to trigger a sorting of the opportunities in the *trigger cluster database* (and potentially the *trigger list*). The sorting is done by the **threshold analyzer**



**sorter worker** which will read the config file to determine what measures to use in the sorting, and then gather the information it needs, and finally sort the opportunities into a *dispatch list*. An output of this group of workers is also a measure of performance that is used for ranking, but potentially also additional measures for ranking or consideration including but not limited to an estimate of the benefit achieved by completing the associated work package. This analyzer layer can also call on external information such as available technicians in an area, day of week, likelihood of access to the network elements necessary, etc.

## 4. Config Files

One or more configuration files (config files) are necessary to define the application from the ProOps elements, and to control how the application functions day to day. There are several types of configuration information needed for this workflow, and that information could be stored in a single file, or multiple files, or in a database, for example. It could be edited in a file editor, or through a graphical user interface (GUI) defined for the application. The files could be separate or together or organized in any way. As developed currently, there is an interface with organization to configure workers and general parts of an application. We assume in this description that configuration will be read-in from a file, but will consider each configuration type as a separate file for the sake of explanation.

- Extraction Config File: This config file specifies the frequency of polling for each data element, and which data elements are polled as part of the application. It may also contain instructions for how to handle missing data, timeouts, and other problems that may be encountered in data polling.
- Trigger Analyzer Config File: This config file contains instructions of which analysis workers are included in the application or workflow, and how they are started (such as when data appears in the *decoded data and analysis results* database or data lake). It also includes the rules by which CMs are triggered for action to the next level, and flagged for ranking, plus what workers need to do to support the action layer.
- Action Config File: This config file contains the details of the model for the measure of performance for CMs based on their statistics and trigger state. The model may be complex or simple, depending on how data over time or missing data are treated, and whether a nonlinear model is needed. It also specifies what to use for ranking. This config file also may provide information for how to cluster CMs, essentially controlling the model for clustering, but also specifying the worker that is used.
- Threshold Analyzer Config File: This config file provides the rules by which opportunities (clusters or individual CMs) are ranked and sorted, included or excluded, on the dispatch list. It also specifies how often the list is revisited and updated, which means how often the clusters and individual CMs are analyzed.

## 5. Data Stores

Several different data stores are needed, though they need not be distinct. However, as the nature of what they store is different, it is likely advantageous to keep them in solutions suited best to their forms, and therefore desirable to be distinct in at least some way.

- Raw PNM Data: The PNM data can be kept in the file system of the CCF, or in the trivial file transfer protocol (TFTP) server used by the operator, or in a separate data lake. As the data held here are disparate, it is not likely that a database is an appropriate solution. Therefore we refer to the (Raw) PNM Data store as a data lake.

- *Decoded data, analysis results:* The decoded data and analysis results can go into a data lake or database, as long as the relations are held between MAC, time, and data to analysis. The data here can be contained in a relational database as it can contain decoded data and statistics of multiple types associated with individual CMs by type, time of day, etc.
- *Stats, trigger list, cluster list:* The data here can be contained in a relational database as it can contain statistics of multiple types associated with individual CMs by type, time of day, etc., and a separate related list of the CMs that meet the threshold for triggering based on the rules set in the config file. The cluster information can be held in a list, a database, or other simple form as it only needs to be a list of how the CMs are clustered in common groups for PNM work opportunities. But all these separate elements are related, so a relational database is a good candidate.
- *Work (dispatch) list:* This data store is simply a list of CMs or CM clusters which are opportunities for work. Each item on the list has a performance measure that indicates how critical the problem is, or how large the proactive opportunity is, so may even be an expected benefit in dollars, for example. There should be links back to the supporting evidence for the decision to add this work to a dispatch list, and to help with troubleshooting.
- *Logs:* The system will keep a list of system logs for overall system health. This may be a method to maintain state under failure too.

## 6. Scheduling Considerations

- A CM can usually only respond to one request at a time, or very few. But if a CM doesn't respond at the time data were requested, and there is reason to expect it should, then a retry may be in order.
- CCF jobs have three states: accepted, complete, and failed. Some complete jobs may have incorrect, errored, or incomplete data responses as well, so a complete job may still be a failed job.
- The parsing of the data output may result in a new fourth state of complete but errored.
- Thus, the state of a completed job from the CCF will create subsequent requests to be scheduled based on the final job status.
  - If the job status is complete, and parsing yields a correct output, then schedule the next job according to the instructions in the config file.
  - If the job status is complete but errored, or failed, then schedule a new request immediately, or after a short amount of time as specified by the config file.
  - The config file should also state whether subsequent jobs are scheduled based on the completion time of the last job, or based on a system clock. In other words, if a data type is requested six times a day, and the first one is delayed by an hour, then does the next data request happen in three hours or four hours? Both may be possible, but configurable in the config file.
- The **control & schedule worker** will need to sort jobs on a common system clock. The control worker handles this task.

## 7. Additional Notes about ProOps

- The config file must account for missing data in the way we analyze and calculate performance measures. A non-responsive CM may end up on a separate list.
- RabbitMQ is our choice for the first version of ProOps for managing the worker queue.
- We stagger polling to get the data we need and manage the poller resources.

# PNM Example Application for ProOps

Here we define the example application that comes with ProOps. Out of the box, very little configuration is necessary to get the system running with the example application. However, the utility of the example application is highly dependent on the user's ability to configure it for their use case specifically.

The first version of the example application that comes with ProOps does not do clustering; the identified CMs will go directly to the threshold analyzer as single CMs. This can be reflected as each CM is its own cluster, and the cluster algorithm doing no work, but still following the diagram in Figure 3.

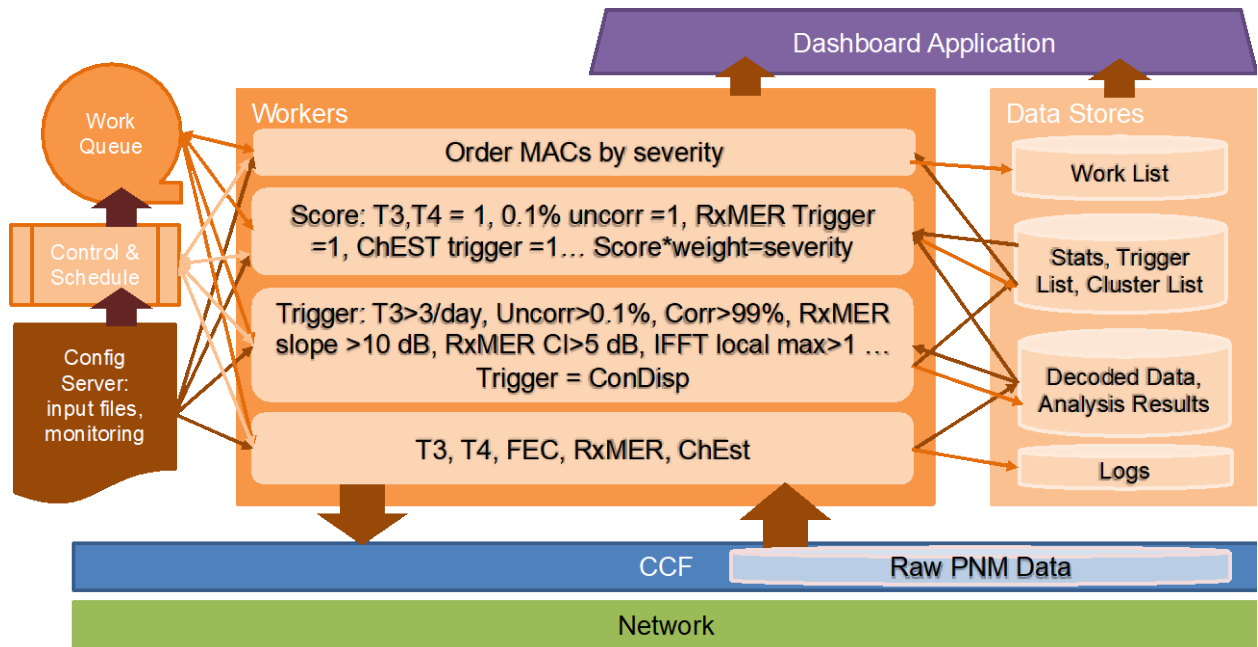


Figure 3 – Depiction of the example application that comes with ProOps today.

## 1. Base data polling (Pollers)

- ☐ Pull T3 and T4 errors from CM logs every hour (configurable).
- ☐ Pull forward error correction (FEC) stats every hour (configurable).
- ☐ Pull receive modulation error ratio (RxMER) every four hours (configurable).
- ☐ Pull channel estimation every four hours (configurable) at same time.

## 2. Triggers (Analyzers)

- ☐ Repeating errors at a configurable frequency, say more than three per day, configurable.
- ☐ Uncorrectables higher than 0.1%, correctables at 99% or higher for over an hour, all configurable.
- ☐ RxMER and SpecAn results with a slope greater than 10 dB, configurable. Also, trigger any RxMER with a 95% confidence interval greater than 5 dB (calculated from percentile responses),

configurable. Also look for ripples in either or both by inverse fast Fourier transform (IFFT) with an energy spike. Take the largest continuous segment of data from either and transform (IFFT) it into time domain, then search for local maximum. Each result with more than one local maximum is triggered (this removes the high frequency main energy). Transform to distance and severity of echo tunnel, trigger anything higher than 10 dB, configurable.

- ChEst indicating a range greater than 10 dB, configurable. Also look for ripples in either or both using same method as above. Transform to distance and severity of echo tunnel, trigger anything higher than 10 dB, configurable.
- Constellation display data would be extracted after other issues were found, such as FEC, T3, T4, RxMER issues. For now, just plot and calculate stats of mean, variance, and third moment.

### 3. Actions

- A CM must be triggered for an action to apply.
- We apply a weight to each trigger type (management information base (MIB), and trigger, so that a given measure could have more than one trigger).
- Score: Create a score for each measure based on the following rules.
  - 1 point for each error or boot issue
  - 1 point for each percent uncorrectable
  - 1 point for each trigger (RxMER, ConDisp), each time
  - 1 point for each trigger in ChEst
- Each day, multiply the score by the weight (score\*weight), and sum the results for each CM each day. Use an exponentially weighted moving average (EWMA) scheme to get the CMs their final scores over time. The EWMA approach is a worker that can be exchanged with other approaches.
- Each CM triggered will have its frequency of data collection doubled, configurable for each measurement. Only after having a 0 score\*weight for three full days, configurable, will a CM be removed from the action list. The frequency of data collection will be configurable by trigger type, and may cross measures. For example, triggering FEC could result in increasing data collection on FEC by five times the rate, and ConDisp each time FEC is collected, or RxMER at twice the rate.
- All CMs on the action list are prioritized by their score\*weight, summed daily over the duration they are on the list (each day's score\*weight summed over the contiguous days each CM is on the list, for each CM), or using the EWMA approach as we will for this example application.
- The list is reported for the CMTS, and the CMTS is scored accordingly by summing the measures for the CMs on the list.

### 4. Modules

The example application that comes with ProOps has several modules, following the structure of ProOps explained previously.

- Control & schedule: A polling scheduler to schedule the data requests in a manner to help the CMTS not get overloaded.
- A data store of results collected, statistics, and scores for each CM, over time. Age off the raw data.
- Translator workers: A translator for each data type pulled.
- Linear calculators: Rapid versions of other calculators or anomaly detectors or other entities which are always applied to data pulled, by data type, before holding in a data store.
- Calculator workers: A calculator or more for each data type pulled.

- ❑ Anomaly detector: An anomaly detector for the spectrum and RxMER data. This is just a simple statistical calculator for the first release. A true anomaly detector will be made available later.
- ❑ Trigger analyzer: An evaluator to decide on the triggering of a CM, and its weight.
- ❑ Severity calculator Worker: A calculator that calculates the severity of the opportunities.
- ❑ Cluster Worker: A clustering algorithm that handles work assignments. This is defined for a later release.
- ❑ Ranker worker: A calculator and ranker for clusters and CMs together to calculate the measure of severity and rank the opportunities.
- ❑ Threshold analyzer sorter worker: A threshold analyzer to indicate dispatchable work, based on rules configured.
- ❑ A configuration file to allow changing all the various default values of the application.
- ❑ A mechanism (GUI later, via a management interface or the dashboard application) for accessing the list of CMs (and clusters) to act on, and a way to compare their severity.
- ❑ Work queue: The queue of work to be performed, organized in a FIFO manner.

The example application is untested in live plant, so the initial configuration described may be far off from a useful combination of settings. The settings will allow verification of function, and be instructive for configuration and developing more functionality, as we now explain.

## Configuring an Application

Applications are defined by how the configuration connects the workers. Thus you can define a new application from an existing one by changing the configuration, or adding new workers to the worker environment and configuring their use, or a combination of these actions. Adding a new worker is outside the scope of this technical report, but we will explain how to configure an application using existing workers as it further explains how to make use of ProOps for PNM uses.

ProOps has several base workers that must be configured for an application to work, as opposed to optional workers that can be configured or not depending on need.

Each PNM measurement can be configured on multiple dimensions, aside from the obvious parameters such as network scope for the deployment and other factors controlled by the context of the deployment, or settings for specific PNM measurements like sampling period for FEC statistics. Some of the settings specific to ProOps that should be evaluated for suitability of default settings include the following.

- Time between data requests to a network element for a given data type.
- Time to wait for a time out on a data request.
- Triggering based on range, standard deviation, linear or curve fitting regression parameters, maximum, minimum, mean, median, mode, third or fourth moments, or other user defined statistics including machine learning or artificial intelligence analysis outcomes.
- Weight given to a performance measure as the result of a trigger, such as the range of a measure times a constant factor, or the square of the sum of the uncorrectable errors in the last hour.
- Which worker or driver to use for a given request, and where to send the output of a worker.
- Various queue management settings to control workers assigned to work.

Cablelabs
PNM Management Server
Config Management
Job Scheduling
System Monitor

Edit Config
1

Required Fields

Worker Identifier (Unique Key)

Task Queue Name

Task Queue Username

Task Queue Password

Task Queue IP

Task Queue Exchange

Task Queue Routing Key

CREATE/UPDATE

**Figure 4 – Configuration management server web GUI index page (task queue configurations).**

See Figure 4. On the index page of the configuration server, the task queue information can be configured. The task queue information is essential for all internal communications. When workers start, they request task queue configuration along with other configurations from the configuration server through RESTful APIs.

Show 10 entries
Search:

Worker Identifier (Unique Key)	Task Queue Name	Task Queue Username	Task Queue IP	Task Queue Exchange	Task Queue Routing Key	Remove Config
result_db_worker	pro_ops_results	testing	10.70.35.120:5672	pro_ops_results	pro_ops_results	REMOVE
pro_ops_job_scheduling	pro_ops_jobs	testing	10.70.35.120:5672	pro_ops_jobs.all	pro_ops_jobs	
data_collector_3	pro_ops_cm_data_collection_jobs	testing	10.70.35.120:5672	pro_ops_cm_data_collection	pro_ops_cm_data_collection_jobs	REMOVE
data_collector	pro_ops_cm_data_collection_jobs	testing	10.70.35.120:5672	pro_ops_cm_data_collection	pro_ops_cm_data_collection_jobs	REMOVE
cm_xmer_stats	pro_ops_xmer_stats_jobs	testing	10.70.35.120:5672	pro_ops_xmer_stats	pro_ops_xmer_stats_jobs	REMOVE
cm_result_threshold_trigger	pro_ops_result_trigger_jobs	testing	10.70.35.120:5672	pro_ops_result_trigger	pro_ops_result_trigger_jobs	REMOVE
cm_condisp_stats	pro_ops_condisp_stats_jobs	testing	10.70.35.120:5672	pro_ops_condisp_stats	pro_ops_condisp_stats_jobs	REMOVE
cm_channel_estimation_stats	pro_ops_channel_estimation_stats_jobs	testing	10.70.35.120:5672	pro_ops_channel_estimation_stats	pro_ops_channel_estimation_stats_jobs	REMOVE

Showing 1 to 8 of 8 entries
Previous
1
Next

**Figure 5 – Configuration management server – defined workers.**

Defined workers are listed at the configuration server's index page, shown in Figure 5. When a new worker starts up, the configuration server uses the worker identifier to determine which configuration the worker uses and replies the RESTful API call with the worker's configuration. The worker identifier is a unique key or name that each worker owns. The worker configuration can be removed by clicking the "REMOVE" button.

Cablelabs
PNM Management Server
Config Management
Job Scheduling
System Monitor
testing

Required Fields
Worker Name (Unique Key)
Worker Identifier
data\_collector
Job Interval (Seconds)
CREATE/UPDATE

Edit Job Message

Show 10 entries
Search:

Worker Name (Unique Key)	Worker Identifier	Job Interval (Seconds)	Last Run	Status Control	Remove Schedule
trigger	cm_resul_threshold_trigger	300	2019-07-16 01:16:09	START	REMOVE
rxmer_stats	cm_rxmer_stats	600	2019-07-16 01:16:09	STOP	REMOVE
data_collector_04	data_collector	30		START	REMOVE
data_collector_03	data_collector	30		START	REMOVE
data_collector_02	data_collector	30	2019-07-16 01:18:35	STOP	REMOVE
data_collector_01	data_collector	600	2019-07-16 01:09:38	START	REMOVE
condisp_stats	cm_condisp_stats	300	2019-07-16 01:17:44	STOP	REMOVE
chest_stats	cm_channel_estimation_stats	300	2019-05-26 22:16:33	START	REMOVE

Showing 1 to 8 of 8 entries
Previous 1 Next

© PNM Management Server Version 0.1 - 2019 Cablelabs

**Figure 6 – Job scheduling.**

The job scheduling page, shown in Figure 6, displays all defined job schedules. When a new scheduling is started, the configuration server sends a message to the task queue which the job scheduling worker is subscribed to. The job scheduling worker will then add the job to its scheduling queue and send task messages to the task queue periodically based on the job interval.

Cablelabs
PNM Management Server
Config Management
Job Scheduling
System Monitor
testing

Required Fields
Worker Name (Unique Key)
trigger
Worker Identifier
cm\_result\_threshold\_trigger
Job Interval (Seconds)
300
CREATE/UPDATE

Edit Job Message

```

1 {
2   "cm_mdc": [],
3   "data_service_address": "http://10.70.35.120:10001",
4   "data_service_store_api": "/cmTriggerResults/",
5   "cm_rxmer_stats": {
6     "data_service_api": "/cmRxmerStats/",
7     "data_service_address": "http://10.70.35.120:10001",
8     "duration": 86400,
9     "trigger_field_config": [
10      {
11        "field": "range",
12        "threshold": {
13          "greater_than": 2
14        }
15      },
16      {
17        "field": "std",
18        "operation": "OR",
19        "threshold": {
20          "greater_than": 0
21        }
22      },
23      {
24        "field": "ordinary_linear_regression_m",
25        "operation": "OR",
26        "threshold": {
27          "out_range": [-0.0001, 0.0001]
28        }
29      }
30    ]
31  },
32  "cm_event": {
33    "data_service_api": "/getCMEventsData/",
34    "data_service_address": "http://10.70.35.120:10001",
35    "duration": 86400,
36    "trigger_field_config": [
37      {
38        "field": "T3_T4_count",
39        "threshold": {
40          "greater_than": 0

```

Show 10 entries
Search:

Worker Name (Unique Key)	Worker Identifier	Job Interval (Seconds)	Last Run	Status Control	Remove Schedule
trigger	cm_result_threshold_trigger	300	2019-07-16 01:16:09	START	REMOVE
rxmer_stats	cm_rxmer_stats	600	2019-07-16 01:16:09	STOP	REMOVE
data_collector_04	data_collector	30		START	REMOVE
data_collector_03	data_collector	30		START	REMOVE
data_collector_02	data_collector	30	2019-07-16 01:16:35	STOP	REMOVE
data_collector_01	data_collector	600	2019-07-16 01:09:38	START	REMOVE
condisp_stats	cm_condisp_stats	300	2019-07-16 01:17:44	STOP	REMOVE
chest_stats	cm_channel_estimation_stats	300	2019-06-26 22:16:33	START	REMOVE

Showing 1 to 8 of 8 entries
Previous 1 Next

© PNM Management Server Version 0.1 - 2019 Cablelabs

**Figure 7 – Configuring job scheduling.**

Here in Figure 7 is an example of what can be added to the job scheduling configuration. Other than job interval configuration, we can also add configurations such as data service address and data store API to let the worker know where to retrieve data and where to store results. For example, by configuring the “cm\_rxmer\_stats” configuration field, the worker will retrieve data from <http://10.70.35.120:10001/cmRxmerStats> with a backlog duration of 86400 seconds (a day), and find all CMs that have an RxMER range (max - min) greater than 2 dB, or standard deviation greater than 0 (which is set artificially small just for demonstration purposes), or an *m* value (slope in ordinary linear regression) that is smaller than -0.0001 or greater than 0.0001 (this is also set to a very small range as an example). Figure 8 shows an example where we set a greater than 0 threshold for T3 and T4 device event log errors.



CableLabs

PNM Management Server

Config Management

Job Scheduling

System Monitor

testing

CREATE/UPDATE

Edit Job Message

```

16 {
17   "field": "std",
18   "operation": "OR",
19   "threshold": {
20     "greater_than": 0
21   }
22 },
23 {
24   "field": "ordinary_linear_regression_m",
25   "operation": "OR",
26   "threshold": {
27     "out_range": [-0.0001, 0.0001]
28   }
29 }
30 },
31 {
32   "cm_event": {
33     "data_service_api": "/getCmEventsData/",
34     "data_service_address": "http://10.70.35.120:10001",
35     "duration": 86400,
36     "trigger_field_config": [
37       {
38         "field": "T3_T4_count",
39         "threshold": {
40           "greater_than": 0
41         }
42       }
43     ]
44   },
45   "cm_fec_stats": {
46     "data_service_api": "/getCmFecStatsData/",
47     "data_service_address": "http://10.70.35.120:10001",
48     "duration": 86400,
49     "trigger_field_config": [
50       {
51         "field": "uncorrectable_codewords",
52         "weight": 1.0,
53         "threshold": {
54           "greater_than": -1
55         }
56       }
57     ]
58   },
59   "cm_channel_estimation": {
60     "data_service_api": "/cmChannelEstimationStats/",
61     "data_service_address": "http://10.70.35.120:10001",
62     "duration": 86400,
63     "trigger_field_config": [
64       {
65         "field": "tilt_db_per_mhz",
66         "weight": 1.0,
67         "threshold": {
68           "out_range": [-0.0001, 0.0001]
69         }
70       }
71     ]
72   }
73 }

```

rxmer_stats	cm_rxmer_stats	600	2019-07-16 01:16:09	STOP	REMOVE
data_collector_04	data_collector	30		START	REMOVE
data_collector_03	data_collector	30		START	REMOVE
data_collector_02	data_collector	30	2019-07-16 01:18:35	STOP	REMOVE
data_collector_01	data_collector	600	2019-07-16 01:09:38	START	REMOVE
condisp_stats	cm_condisp_stats	300	2019-07-16 01:17:44	STOP	REMOVE
chest_stats	cm_channel_estimation_stats	300	2019-06-26 22:16:33	START	REMOVE

Showing 1 to 8 of 8 entries

Previous 1 Next

© PNM Management Server Version 0.1 - 2019 CableLabs

**Figure 8 – Configuring for T3 and T4 device event log errors.**

CableLabs

PNM Management Server

Config Management

Job Scheduling

System Monitor

testing

Required Fields

Worker Name (Unique Key)

rxmer\_stats

Worker Identifier

cm\_rxmer\_stats

Job Interval (Seconds)

600

CREATE/UPDATE

Edit Job Message

```

1 {
2   "data_service_api": "/getCmRxMerData/",
3   "data_service_address": "http://10.70.35.120:10001",
4   "data_service_store_api": "/cmRxMerStats/",
5   "cm_mac": [],
6   "duration": 86400
7 }

```

Show 10 entries

Search:

Worker Name (Unique Key)	Worker Identifier	Job Interval (Seconds)	Last Run	Status Control	Remove Schedule
trigger	cm_result_threshold_trigger	300	2019-07-16 01:16:09	START	REMOVE
rxmer_stats	cm_rxmer_stats	600	2019-07-16 01:16:09	STOP	REMOVE
data_collector_04	data_collector	30		START	REMOVE
data_collector_03	data_collector	30		START	REMOVE
data_collector_02	data_collector	30	2019-07-16 01:18:35	STOP	REMOVE
data_collector_01	data_collector	600	2019-07-16 01:09:38	START	REMOVE
condisp_stats	cm_condisp_stats	300	2019-07-16 01:17:44	STOP	REMOVE
chest_stats	cm_channel_estimation_stats	300	2019-06-26 22:16:33	START	REMOVE

Showing 1 to 8 of 8 entries

Previous 1 Next

© PNM Management Server Version 0.1 - 2019 CableLabs

**Figure 9 – Configuring the RxMER statistics worker.**

The job scheduling of the RxMER statistics worker (see Figure 9) is another example of how to configure a worker. The RxMER statistics worker only requires brief configurations that include job interval, data service API, data service address, duration, and cm\_mac list. The job interval defines how frequently the

job is scheduled. The data service API and data service address defines where the RxMER statistics worker retrieves RxMER data. Duration defines the data retrieving window from the moment the data is retrieved. The cm\_mac list helps filter interesting cable modem MAC addresses. When the list is empty, all CM data will be retrieved; otherwise, only CM RxMER data from cable modems in the cm\_mac list will be retrieved. The output from the RxMER statistics worker includes "range", "std", "max", "min", "mean", "ordinary\_linear\_regression\_m", and "ordinary\_linear\_regression\_c".

**Required Fields**

Worker Name (Unique Key)  
data\_collector\_01

Worker Identifier  
data\_collector

Job Interval (Seconds)  
600

[CREATE/UPDATE](#)

**Edit Job Message**

```

1 {
2   "cmts_ip": "10.32.40.68",
3   "dccb_ip": "10.95.254.82:8888",
4   "db_server_address_and_api": "http://10.70.35.120:10001/putData",
5   "namespace": "dccb",
6   "cm_ip_to_filter": "",
7   "cmts_community_string": "CableLabsSNMP",
8   "cm_community_string": "private",
9   "drivers": ["cmPNMDsRxMer"],
10  "drivers_actual_name": ["cmPNMDsRxMer"],
11  "timeouts": [290]
12 }

```

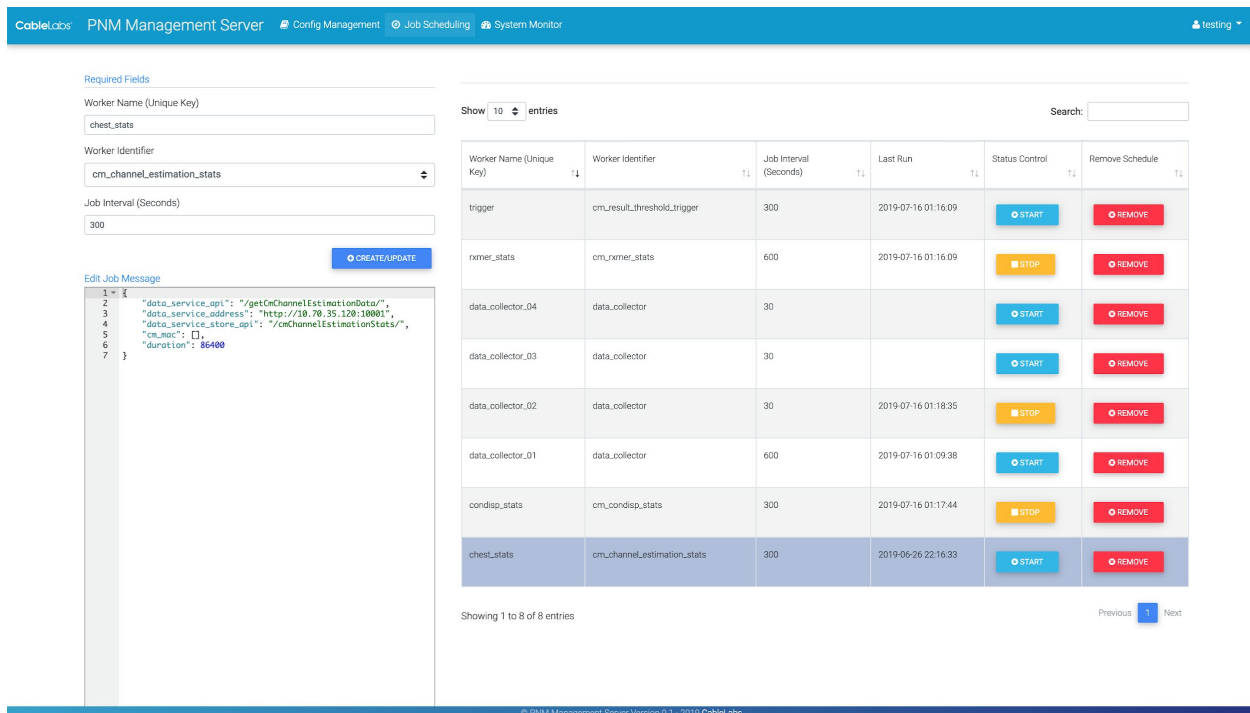
**Table: Worker Configuration**

Worker Name (Unique Key)	Worker Identifier	Job Interval (Seconds)	Last Run	Status Control	Remove Schedule
trigger	cm_result_threshold_trigger	300	2019-07-16 01:16:09	<a href="#">START</a>	<a href="#">REMOVE</a>
rxmer_stats	cm_rxmer_stats	600	2019-07-16 01:16:09	<a href="#">STOP</a>	<a href="#">REMOVE</a>
data_collector_04	data_collector	30		<a href="#">START</a>	<a href="#">REMOVE</a>
data_collector_03	data_collector	30		<a href="#">START</a>	<a href="#">REMOVE</a>
data_collector_02	data_collector	30	2019-07-16 01:18:35	<a href="#">STOP</a>	<a href="#">REMOVE</a>
data_collector_01	data_collector	600	2019-07-16 01:09:38	<a href="#">START</a>	<a href="#">REMOVE</a>
condisp_stats	cm_condisp_stats	300	2019-07-16 01:17:44	<a href="#">STOP</a>	<a href="#">REMOVE</a>
chest_stats	cm_channel_estimation_stats	300	2019-06-26 22:16:33	<a href="#">START</a>	<a href="#">REMOVE</a>

Showing 1 to 8 of 8 entries

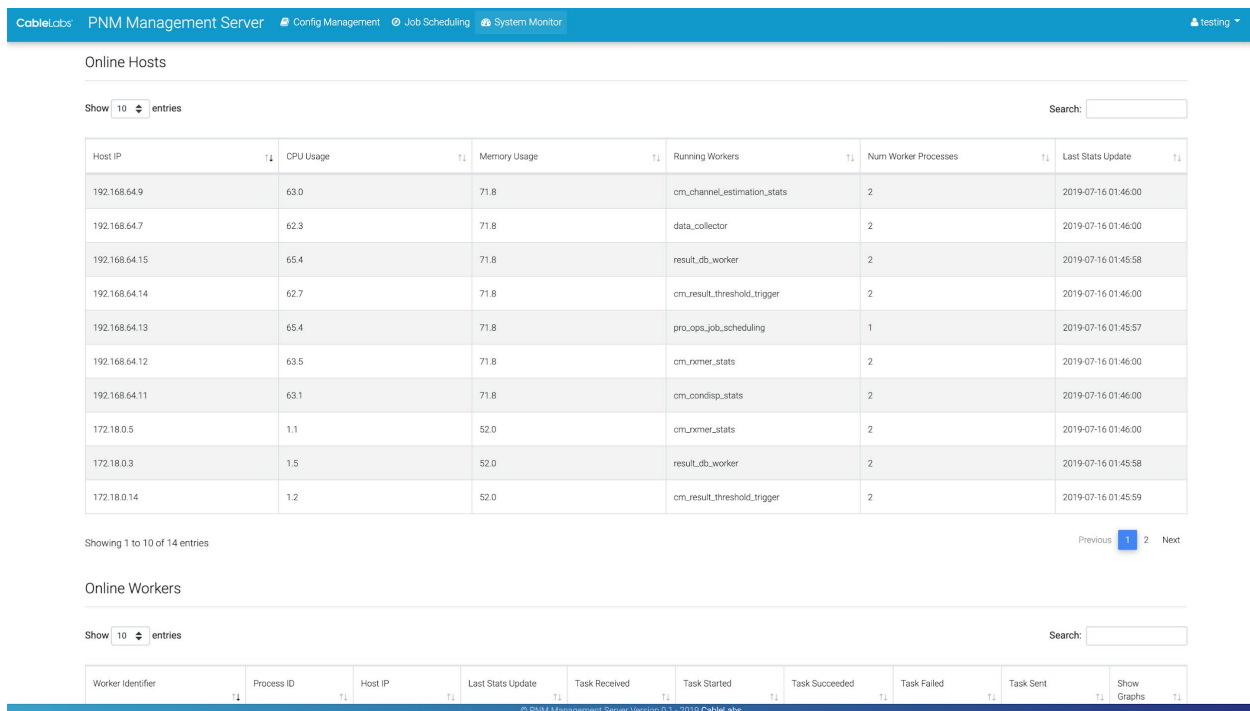
**Figure 10 – Configuring a data collection worker.**

Shown in Figure 10, the data collection worker is a special worker that interacts with the CCF. In this example, based on the configuration, the data collection worker is configured to run every 600 seconds. It points to a CCF instance with an IP address of 10.95.254.82:8888, and it collects cmPNMDsRxMer data from CMTS 10.32.40.68. The collected data are automatically decoded and stored using the API defined in "db\_server\_address\_and\_api". The cmts\_ip field can be a list, and can contain multiple CMTS IP addresses.




**Figure 11 – Channel estimation statistics worker.**

Figure 11 shows the CM downstream channel estimation worker, which has a configuration that is similar to CM RxMER statistics worker's configuration. The output includes calculation results on the tilt of the channel estimation and inferences on potential echo distances.

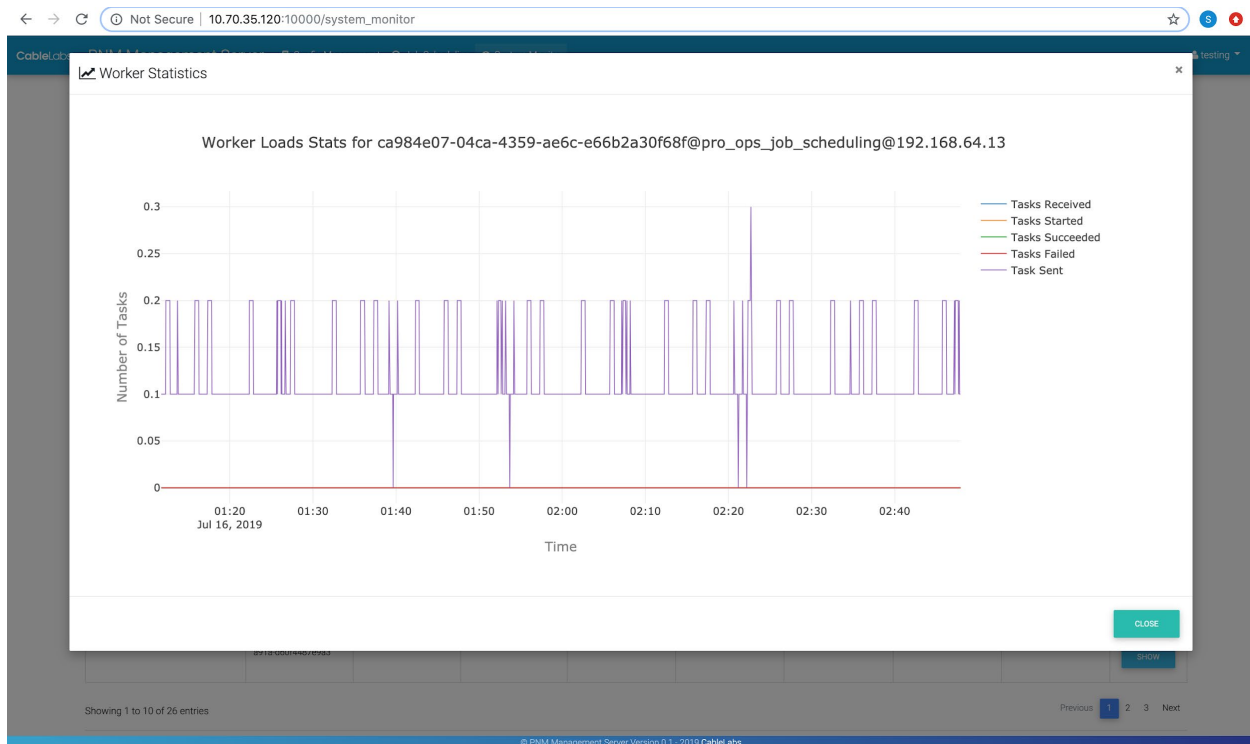


**Figure 12 – Top-level system monitoring.**

Online Workers										
Show 10 entries								Search:		
Worker Identifier	Process ID	Host IP	Last Stats Update	Task Received	Task Started	Task Succeeded	Task Failed	Task Sent	Show Graphs	
result_db_worker	771afbce-9294-4806-9fc4-6931b15c13e6	172.18.0.3	2019-07-16 01:49:25	36	36	36	0	0	 SHOW	
result_db_worker	dd82d590-14a5-42e9-8b63-7f3e08c952dc	172.18.0.3	2019-07-16 01:49:22	36	36	36	0	0	 SHOW	
result_db_worker	e40c42f6-18a7-42bc-a53a-e79f910a14d7	192.168.64.15	2019-07-16 01:49:23	6130	6130	6130	0	0	 SHOW	
result_db_worker	e7cfd36e-1478-462e-8f8e-b85e661c2899	192.168.64.15	2019-07-16 01:49:22	6131	6131	6131	0	0	 SHOW	
pro_ops_job_scheduling	80c524ac-0986-445f-9aa9-e5518cc42b51	172.18.0.11	2019-07-16 01:49:22	0	0	0	0	0	 SHOW	
pro_ops_job_scheduling	ca984e07-04ca-4359-a6c6-e66b2a30f68f	192.168.64.13	2019-07-16 01:49:23	67	67	67	0	69790	 SHOW	
data_collector	12c7ab2f-446e-4c07-b911-e19c6d4ff99b	172.18.0.13	2019-07-16 01:49:23	245	245	0	245	0	 SHOW	
data_collector	3eb309e5-6721-499e-b352-c81f439ce29	172.18.0.13	2019-07-16 01:49:25	246	246	0	246	0	 SHOW	
data_collector	63764412-3b2c-44b7-9775-9c8c05c3f82a	192.168.64.7	2019-07-16 01:49:20	29445	29445	10673	18772	0	 SHOW	

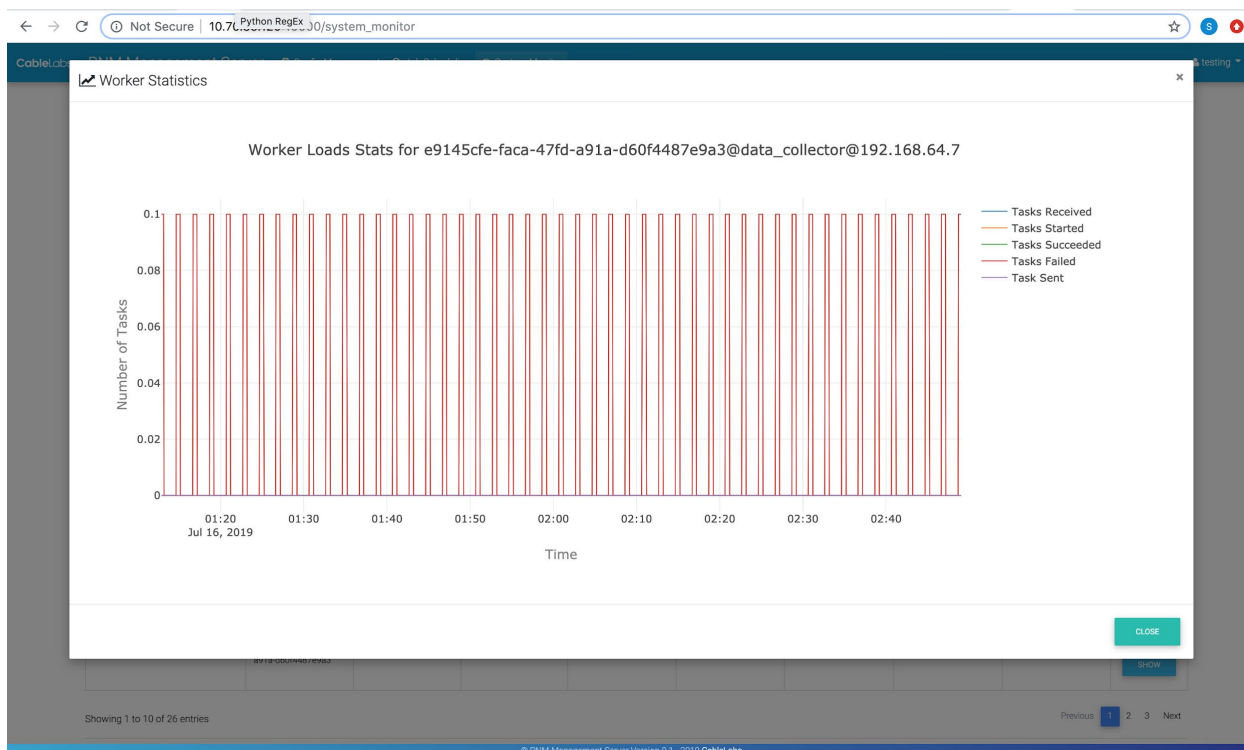
**Figure 13 – Detailed worker system monitoring.**

Figures 12 and 13 show the system monitoring capabilities that come with ProOps. The configuration server has a page that displays worker running statistics and host resources. In Figure 12, workers are running in separate docker containers and these containers show up as different hosts. The system monitoring page shown in Figure 13 also has a table that displays more detailed worker statistics. The statistics are calculated on a per-worker process basis, and include how many tasks are received, started, succeeded, failed, and sent. As an example, some data collection tasks failed here because of lab equipment changes and wiring changes during collection.



**Figure 14 – System monitoring graph.**

Here in Figure 14 is an example visualization of how the job scheduling worker is handling tasks every few seconds. The graph shows average numbers for each counter over time. This output provides a quick and easy way to see whether system resources are sufficient, or ProOps is properly configured to perform well, and as intended for the application configured.



**Figure 15 – System monitoring graph of data collection.**

See Figure 15. Here is an example visualization of how the data collection worker is handling tasks every few seconds. The graph shows average numbers of each counter over time. This is useful for making sure the data collection resources are not over taxed, or the network is not over taxed as well.

## Envisioned Use Cases

We created ProOps to support several use cases for both operators and vendors. Because of this, we had to create it to support several broad advantages.

- **Scalability** – The architecture of ProOps extends CCF in a modular, layered way. As such, we have full options of ways to package the software, and scale it to any use case, from a desktop deployment to a data center to a cloud architecture.
- **Low adoption friction** – Existing deployments of CCF can easily be extended to host ProOps. Once ProOps is installed, applications are simple updates to ProOps as additional, optional software modules that link into it. Each additional module installed within ProOps behaves like an update.
- **Flexibility and configurability** – Because ProOps allows anyone to add workers to the environment and configure rules for their interaction, each deployment of ProOps can be configured independently and in a broad range of ways. That flexibility provides full control to the user, which means it can be configured poorly, but because it is flexible it can be easily corrected and adjusted to fit changing needs too.
- **Actionable output** – The use of layers of workers as explained previously will allow ProOps to be configured so that the output can be acted on with confidence. If conditions change and the repair

work that it suggests is no longer as desired, then adjustments to the configuration can correct that easily.

- Interworkability – Use of CCF allows one data source to support many applications. Likewise, ProOps is a platform that can support many applications in the same platform, allowing them to share information throughout the process of turning data into action. New applications built in ProOps can run with other existing applications in the same platform deployment. Hardware and virtual machine (VM) resources are the main consideration only.
- Modular – ProOps is modular, so you can replace elements like the worker queue, databases, and worker modules with your own creations or to fit our own corporate information technologies (IT) policies.
- Speed – CableLabs has and will continue to improve the speed of execution of ProOps, but what is important to note is that an application can be created in ProOps rapidly, so it brings speed of execution and deployment together. As we develop more basic workers for the environment, much more will be possible with just a simple reconfiguration of the platform to utilize existing workers in different ways.

Several operator use cases can be supported by ProOps.

- Experiment with PNM – You can connect ProOps to a small test network and test its capabilities against your expectations, and tune it for your own needs. Then depending on the outcome of your experiment, you can take the solution and deploy it in your network in many different ways, whichever suits your situation best.
- Build and test a new PNM idea or solution – As ProOps is a flexible platform, you can create your own PNM method and test it in an example or live network as you see fit, and even tune the platform to your new solution before full deployment.
- Development environment – Because of its flexibility and open architecture, you can develop prototype solutions inside of ProOps, essentially making it your team's development environment for operations solutions.
- Network sampling – ProOps was created to schedule work, so you can schedule network sampling for any operations need, including creating requirements for PNM, or to build a business case for a conceived PNM solution.
- Grow your own PNM program – Because you can use ProOps for the entire chain of PNM solutions, you can experiment, build, and deploy what you come up with, and support it yourself if you wish. Each development in ProOps works with the previous, so you can grow the entire solution and keep it updated as network needs change, keeping your solution up to date, and tailored for specific problems, architectures, or needs in an area.
- Gather requirements for vendor supported solutions – Because you can use it as an experiment platform or to develop a business case for a PNM effort, it can be used to support a vendor-supported PNM solution too.

Vendors also can use ProOps to their advantage.

- Rapid prototype new potential solutions – Vendors can use ProOps to develop their new PNM products or services, and
- A framework to support products and services – Vendors can use ProOps as a vehicle to deploy their solutions, to ease adoption and reduce the operations impact of the deployment, and to work in harmony with home-grown or other vendor solutions in an integrated manner.

- Free sample solutions – Vendors can deploy their proprietary solutions as example offerings with limited capabilities so operators can test their solution before purchasing the full enterprise supported version.
- Network sampling for developing operator specific solutions – Vendors can deploy a network sampling version of ProOps in an operator's network to help them determine the benefits of implementing a specific PNM or operations solution, to help get over the business case uncertainties that may hinder purchase decisions.
- Rapid, flexible data collection for consultation – Vendors who work as consultants to operators can use ProOps to collect network data to look for a specific problem, or to support a specific network issue they are trying to resolve.

ProOps was made to turn data into action, so we expect network operations and engineering personnel will discover new ways to make use of it that we have not yet defined.

## Conclusion

ProOps is free to use by CableLabs vendor and operator members. Contact the authors of this paper to obtain a copy of the software, and to get help configuring it for your needs. We encourage everyone who can access ProOps to use it in any way they envision, from reviewing the architecture to taking those advantages in their own developments on up to full use and deployment of the code base.

ProOps is a platform constructed for turning data and information into action. It is built around well defined, general steps that facilitate making decisions automatically, but while keeping control fully in the hand of the users. Rather than relying on assumed experts at hand to review network data to determine what needs to be done, ProOps provides an environment to automate that work. This means limited expert resources can be shared through ProOps, thus extending the effectiveness of expertise. For example, CableLabs expects to work with the PNM community to build into ProOps much of what will be documented in the forthcoming DOCSIS® 3.1 PNM Best Practices document.

ProOps is well suited for PNM, but it can be used for turning any data into action, really. With CCF being fully flexible to gather most any network or system data source through creation of a driver, ProOps likewise can collect that data, analyze it, add context, translate results into potential work, then select the work that is most important to do. Network operations efficiency and improving service reliability are the intended goals of ProOps, but only our imaginations will limit what it can do.

We hope, and fully support, operators and vendors contributing code to the C3 repository for sharing solutions, guiding the industry to solve problems, and sharing ideas. CableLabs intends to build workers with new capabilities and share them in ProOps as workers with which members can build solutions, and for CableLabs to build other proof-of-concept applications. In the months ahead, the power of ProOps will increase due to the contributions from the entire community. We hope capable operators and interested vendors will work with us to develop workers and applications that solve important PNM needs so that both operators and vendors may take advantage from the PNM capabilities that CableLabs provides.



## Abbreviations

API	application programming interface
C3	common code collection
CableLabs	Cable Television Laboratories
CCF	common collection framework
CM	cable modem
CMTS	cable modem termination system
CMVA	cable modem validation application
dB	decibel
DOCSIS	Data-Over-Cable Service Interface Specifications
EWMA	exponentially weighted moving average
FEC	forward error correction
FIFO	first in first out
GIS	geographic information system
GUI	graphical user interface
IP	Internet protocol
IPR	intellectual property rights
ISBE	International Society of Broadband Experts
IT	information technologies (information technology)
MAC	media access control
MIB	management information base
NDA	nondisclosure agreement
OODA	observe, orient, decide, and act
PNM	proactive network maintenance
ProOps	Proactive Operations
RxMER	receive modulation error ratio
SCTE	Society of Cable Telecommunications Engineers
TFTP	trivial file transfer protocol
VM	virtual machine

## Bibliography & References

- [1] <https://danford.net/boyd/>
- [2] *Boyd's OODA Loop and the Infantry Company Commander*, A. Bazin, Infantry Magazine, 2005.
- [3] CableLabs Proactive Network Maintenance Combined Common Collection Framework Architecture Technical Report, CL-TR-XCCF-PNM-V01-180814, August 14, 2018, Cable Television Laboratories, Inc.
- [4] <https://code.cablelabs.com/proactive-operations-platform>
- [5] <https://www.cablelabs.com/cable-network-reliability-proops-platform-for-pnm-and-more>

AUTHOR INDEX  
2019 Technical Paper Proceedings

Aggarwal, Rajeev .....	386	Cloonan, Tom .....	455, 1798
Al-Banna, Ayham .....	1798	Colby, Andrew .....	318
Alcox, Kevin .....	1431	Cole, Jason .....	1
Ansari, Furquan .....	970	Combs, Doug .....	824
Asati, Rajiv .....	1462, 1496	Condra, Steve .....	1416
Bachofen, Ralph .....	439	Cooper, Michael .....	954
Bacon, Bruce .....	187	Cunningham, Ryan Michael .....	1003
Badawiyeh, Basil .....	248	Currian, Bruce .....	89
Baran, Dave .....	30	Cyr, Greg .....	1769
Beesley, Bill .....	1154	Daoud, Mohamed .....	386
Belford, Thomas .....	664	Davis, Drew .....	740
Bender, Andrew .....	1112	Denis, Xavier .....	604
Bendt, Brian .....	1022	Douglas, John .....	945
Bernstein, Alon .....	1462, 1496	Eagles, Michael .....	722
Beshara, Hani .....	1403	Evans, Alan .....	932
Briscoe, Bob .....	1742	Fenby, Asten .....	164
Brooks, Roger .....	318	Ferreira, Jude .....	1311
Brophy, Jay .....	187	Finkelstein, Jeff .....	30
Buhl, Lauren .....	1722	Fish, Roger .....	89
Busch, Chris .....	30	Flesch, J.R. ....	1127
Campos, L. Alberto .....	786	Foroughi, Nader .....	1240
Carothers, Matt .....	1706	Gammons, John .....	1081
Carro, Gabriel .....	288	Ganji, Monsour .....	576
Chapman, John T .....	1567, 1618	Gaydos, Bob .....	824
Cheevers, Charles .....	630, 1127	Ger, Javier .....	970
Cloonan, Ruth .....	1798	Ghatge, Charuhas .....	596

Gibellini, Emilia .....	288	Kolze, Tom.....	89
Gilberton, Phillippe .....	1278	Kreishan, Loay.....	78
Godlewski, Marcin .....	1176	Krishna, Karthik .....	1161
Goemaere, Patrick .....	532	Kumar, Pankaj .....	318
Haefner, Kyle .....	905	Kurkowski, Stuart.....	446
Harb, Maher.....	1311	Lane, Cliff.....	1100
Hassler, Garey .....	1643	Latini, Patricio Sebastian .....	260
Heijnen, Henk.....	1278	Liew, Jay.....	187
Hewavithana, Thushara .....	1567	Lin, James.....	1354
Hillermeier, Rainer .....	1567	Liu, Tong .....	1618, 1654
Hmimy, Hossam .....	386	Lumbatis, Kurt.....	1127
Holloran, Tom .....	203	Maki, Kari.....	1416
Hoole, Elliott .....	427	Marcinko, Tomislav.....	1722
Howlett, Colin .....	51	Maricevic, Zoran.....	1203
Hranac, Ron.....	89	Masoud, Fady .....	922
Hubbard, Matthew .....	386	Matatyaou, Asaf.....	1022
Jain, Mudit.....	318	Medlock, James .....	89
Jain, Nandit.....	318	Meisen, Kai.....	51
Jia, Zhensheng (Steve).....	786	Milley, Andrew Joseph.....	1855
Jin, Hang.....	1567	Mills, Robbie .....	248
Job, David.....	954	Mishra, Mayank.....	1354
Johansson, Kjell.....	1674	Mocerino, Joe .....	1559
Johnson, Douglas.....	51	Morales, Noé .....	164
Jones, Doug .....	887	Mutalik, Venk.....	824
Karam, Edouard.....	806	Neugeboren, Yair.....	1769
Kelkar, Anish.....	740	Ojeda, Miguel Masache .....	970
Keller, Joe.....	519	O'Keeffe, Frank .....	1798
Kelly, Bryan.....	1187	Olfert, Matthew .....	142
Knittle, Curtis .....	786	Pala, Max .....	586

Patel, Vipul .....	248, 604	Stengrim, Chris.....	786
Pendari, Arash .....	722	Stevens, J.Clarke.....	1526
Pettus, Randy .....	187	Strauss, Derek.....	873
Plant, Sam.....	519	Sundaresan, Karthik.....	1354, 1742
Poggio, Esteban .....	970	Tang, Ben .....	970
Priore, Thomas .....	1431	Teflian, Mark.....	187
Purmonen, Arttu .....	1416	Terada, Mike.....	248
Raghavendra, Avinash.....	1394	Tooley, Matthew.....	664
Ramakrishnan, Sangeeta.....	1496	Ulm, John .....	455, 1203, 1798
Ramaswamy, Kishan .....	1100	Urban, David John.....	689
Ravisankar, Arun .....	1065	Vigouroux, Jean-Ronan .....	1278
Rice, Dan .....	824, 1311	Vij, Megha.....	318
Rieckmann, Derek .....	1699	Vogel, Mark.....	751
Righetti, Claudio.....	288	Wall, Bill .....	954
Rodolico, Joe .....	369, 379	Wan, Tao .....	576, 586
Romano, Carlos Germán Carreño.....	288	Wang, Michael Ting .....	856
Rupe, Jason.....	89, 486, 1870	Weerashighe, Srilal M .....	248
Sahin, Yildirim .....	586	Whaley, Damien .....	1706
Sanders, Joshua.....	427	Wheelock, Ian.....	630
Santangelo, Bryan.....	1311	White, Greg .....	1742
Schwechel, Craig .....	1722	Williams, Tom.....	89
Scriber, Brian A .....	992, 1449	Wilson, Giles .....	722
Shakil, Kashif .....	225	Wilson, Marth.....	1722
Shulman, Shaul.....	765	Wolcott, Larry .....	89
Smith, Andy.....	1654	Wolfe, Ron .....	1547
Sowinski, Pawel.....	1654	Xu, Mu.....	786
Spanbauer, Rick .....	1311	Yeo, Elaine .....	330
Spear, Greg .....	806	Zettinger, Chris.....	1769
Stafford, Roger G.....	1038	Zhang, Haipeng .....	786

Zhang, Junwen.....	786
Zhu, Jay .....	1354
Zhu, Jingjie.....	1870

