

**THE COMPLETE
TECHNICAL PAPER PROCEEDINGS**
FROM:



Published by:  **ncta** 

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

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

ISBN Number: 0-940272-57-1
©2018, NCTA – The Internet and Television Association.
All rights reserved.

Jon Baldry - Infinera
5G Is Rapidly Approaching, What Must Cable MSOs Do To Capitalize On This Business Opportunity..... 1

Dave Morley - Shaw Communications Inc./Freedom Mobile
5G Small Cells and Cable – Realizing the Opportunity..... 13

Tobias Peck - Alpha Technologies
A Customer Experience Based Approach to Improving Access Network Power Reliability 33

Larry Wolcott; Michael O'Dell; Peter Kuykendall; Vishnu Gopal; Jason Woodrich; Nick Pinckernell - Comcast
A PNM System Using Artificial Intelligence, HFC Network Impairment, Atmospheric and Weather Data to Predict HFC Network Degradation and Avert Customer Impact 48

R. J. Vale; Martin J. Glapa; Jean-Philippe Joseph – Nokia, Bell Labs Consulting
Achieving Significant Space, Energy, and Cost Reductions with Future Virtualized Distributed Access RPD and RMD Architectures for MSOs 82

Fernando X. Villarruel; Michael Mobley; Curt Dalton - Cisco Systems, Inc.
Lamar West, Ph.D. - LEW Consulting, LLC
Adaptive Power Management for Node Clusters..... 93

John Holobinko; Glenn McGilvray; John Ritchie - Cisco Systems, Inc.
Analysis and Prediction of Peak Data Rates Through DOCSIS Cores..... 117

Michael Klobardans; Shlomo Ovadia - Charter Communications
Assuring Data Delivery of Critical IoT Devices 131

Michael O'Hanlon; Eric Heaton - Intel Corporation, Network Platforms Group
Automation of Virtual CCAP to Reduce OpEx and Enable New Revenue Streams in the Access Network 148

Greg Nicholson - NuTEQ Solutions, provider of GOCare
Automation Opportunities for Subscriber Management in Cable Television..... 168

Erik Gronvall - CommScope
Cable's Role in the 5G Evolution..... 186

Fernando X. Villarruel; Martin Mattingly - Cisco Systems, Inc.
Capacity and Technology Considerations in DAA Backhaul Deployment Strategies 194

Rajat Ghai - Technicolor
CBRS Use Cases with Focus on Localized Indoor Mobile Access (LIMA), Mobility and Service Continuity..... 216

Tong Liu - Cisco Systems, Inc
Characterization of Spectrum Resource Scheduling in FDX DOCSIS..... 252

Harj Ghuman; David Job - Cox Communications
Coherent Access Applications for MSOs 272

Satish Chalapati - Tata Consultancy Services
Comparison of LPWA Technologies and Realizable Use Cases 295

Arun Ravisankar - Comcast Corporation
Computing at the Edge Still has an Edge 307

Amit Singh - Cisco Systems, Inc.
Eric D. Heaton - Intel Systems, Inc.
Converged Multi Access Networks..... 321

Sandeep Katiyar - Nokia Bell Labs Consulting
Converging Edge Caching and Computing Power for Simultaneous Mobile and MSO Networks to Handle Latency Sensitive Services Using Co-operative Caching 336

Jon Baldry - Infinera
Critical Considerations for the Design of a Robust and Scalable DAA Aggregation and Transport Network..... 349

Jing Wang, Ph.D.; ZhenSheng (Steve) Jia, Ph.D.; Luis Alberto Campos, Ph.D.; Curtis Knittle, Ph.D. - CableLabs
Delta-Sigma Modulation for Next Generation Fronthaul Interface..... 359

Eduardo M. Panciera Molanes; Adrian Grimaldi; Norberto Harmath; Gaston Diaz; Marcos Aberastury - Telecom S. A. <i>Deploying IP Video Services, Architectures and Technologies from Head End to the Home Network</i>	384
Yasser Syed; Alex Giladi - Comcast Cable Ali C. Begen - Ozyegin University <i>Designing Video Services for Low-Latency Distributions in IPTV Cable Systems</i>	442
Chris Ruff - Glympse <i>Digitizing the Customer Experience: Win Loyalty and Sell More with Last Mile Service Trackers</i>	454
Patrick Goemaere; Rajat Ghai - Technicolor <i>Edge Compute and Software Life-Cycle Management: Creating Consumer Value and Flexibility</i>	465
Andrew Sundelin - Guavus, Inc. <i>Embracing Service Delivery Changes with Machine Learning</i>	506
Mark Stratton - Hitachi Consulting Corporation <i>Enabling Smart Cities by Leveraging IoT Sensors, Multi-building Modeling and Analysis, and Smart Energy Business Case Analytics</i>	517
Shiby Parayil; Earl Villanueva - Ericsson North America <i>Enhancing Service Agility for the Enterprise Customers using an Integrated Orchestration and Test Automation Solution</i>	528
Bart Vercammen; Jos Delbar - Technicolor <i>Enhancing Wi-Fi QoE with Targeted Approach</i>	539
David Goodwin; Charles Cheevers - ARRIS <i>Evolving The “Box”: The Smart Set-Top Box</i>	560
Karthik Sundaresan - CableLabs <i>FDX & D3.1 Capacity Scenarios</i>	587
John T. Chapman; Hang Jin - Cisco Systems, Inc. <i>FDX DOCSIS Line Extender</i>	620

Jos Delbar; Bart Vercammen - Technicolor <i>Guaranteeing Seamless 4K OTT Content Delivery</i>	642
Narayan Menon - Fontech <i>Harvesting Unlicensed and Shared Spectrum: Opportunities and Challenges</i>	667
Nader Foroughi - Shaw Communications <i>HFC Evolution - The Best Path Forward</i>	682
Narayan Raman; Yadhav Krishnan; Miguel Hernandez; Furquan Ansari - Bell Labs Computing/ Nokia <i>How An MSO Can Leverage SD-WAN To Grow Its Enterprise Revenue</i>	703
Nav Kannan; Charles Cheevers - ARRIS International plc <i>How to Finally Conquer Wi-Fi in the Home – Service Provider Style</i>	720
Zhensheng (Steve) Jia, Ph.D.; L. Alberto Campos, Ph.D.; Mu Xu, Ph.D.; Haipeng Zhang, Ph.D.; Jing Wang, Ph.D.; Curtis Knittle, Ph.D. - CableLabs <i>Impact of Access Environment in Cable’s Digital Coherent System – Coexistence and Full Duplex Coherent Optics</i>	766
Sanjay Dhawan - Ericsson Inc. <i>Implications of 5G Low-latency Requirements on Hybrid Fiber-Coaxial Networks</i>	787
Pravin Mahajan - Infinera <i>Improving the Customer Experience with Network Automation and with AI-Powered Voice</i>	799
Patrick Iannone; Yannick Lefevre; Werner Coomans; Dora van Veen; Junho Cho - Nokia Bell Labs <i>Increasing Cable Bandwidth through Probabilistic Constellation Shaping</i>	807
Tim Johnson - Alpha Technologies Arun Ravisankar - Comcast J. Clarke Stevens - Shaw Communications Chris Bastian - SCTE-ISBE <i>Internet of Things Dynamics: Opportunities and Challenges for Broadband Network Operators</i>	821

Ron Ih - Kyrio
Internet of Things Security: Implement a Strong, Simple & Massively Scalable Solution
..... 839

Michael Fay - Akamai Technologies, LLC
Internet Scale Blockchain Architecture..... 858

Charles Cheevers; Jonathan Wu - ARRIS International plc
Is the Smart Assistant Mutually Inclusive with IoT? 867

Tal Laufer; Jeroen Putzeys - ARRIS Uffe Callesen - Stofa
It's ALIVE! Getting to Successful R-PHY Deployment: Do's and Don'ts 886

Jim Walsh; David Hering - VIAVI Solutions
David Judge - Vodafone New Zealand
Learnings From a DAA/DOCSIS 3.1 Early Adopter: Launching and Maintaining a Next-Gen HFC Plant 921

Bill Beesley; Ladan Pickering - Fujitsu Network Communications
LoRa, the LPWA Choice for Cable Operator's Entry into Smart Home Market..... 935

Joe Walsh - inCode Consulting
Low Power Wide Area Technologies for IoT Use Cases: Technology Assessment for MSOs
..... 942

Gary Gutknecht - Technicolor
Navigating IoT Technologies, Standards and Frameworks for Managed IoT Services..... 956

Jennifer Andreoli-Fang, Ph.D.; Bernard McKibben - CableLabs
Alon Bernstein; Aeneas Dodd-Noble; Elias Chavarria Reyes, Ph.D. - Cisco
Curt Wong - Charter Communications
Network Convergence 984

Ted Boone; Jignesh Patel; Rob Ames; Kyle Cooper; Chaitanya Vasamsetty - Cox Communications, Inc.
Network Planning Automation using Big Data
..... 1009

Fady Masoud - Infinera
Network Programmability – A Reality Check and a Glimpse into the Future 1017

Brian Lavallée - Ciena Corporation
New Packet Network Design for Transporting 5G Fronthaul Traffic..... 1027

Robert Gaydos; Mehul Patel; Joe Solomon - Comcast
Node Provisioning and Management in DAA
..... 1037

Christopher Topazi; Michael Cooper - Cox Communications
Operational Considerations and Optimization of OFDM Deployments..... 1051

Daniel Howard - Enunciant, LLC
Chris Day - Analog Devices
Kevin Gantt - CommScope
John Holobinko - Cisco
Rob Howald; Dan Marut - Comcast
Dick Kirsche - ConsultKirsche
Todd Loeffelholz - Alpha
Kathleen Miles - PG&E
Rene Spee - Copperva
Dean Stoneback - SCTE•ISBE
John Ulm - Arris
Lamar West - LEW Consulting
Dan Whitehouse - Hitachi Consulting
Operational Practices for Energy Conservation and Sustainability Measures in the Cable Outside Plant..... 1072

Shamil Assylbekov, Ph.D.; Devin Levy - Charter Communications
Operational Transformation Via Machine Learning..... 1105

Todd Loeffelholz - Alpha Technologies Inc
Opportunities and Challenges of Implementing Wireless: Small Cells / IoT / WiFi..... 1112

Alon Bernstein - Cisco Systems, Inc.
Orchestration: What is Really Behind this Overloaded and Overused Term?..... 1124

Asaf Matatyaou - Harmonic, Inc.
Richard J. Walker - Shared Services
Practical Deployment Lessons of a Centralized Virtualized CMTS..... 1135

Justin Watson - Irdeto
Roger Brooks; Andrew Colby; Pankaj Kumar; Anant Malhotra; Mudit Jain - Guavus, Inc.
Predicting Service Impairments from Set-top Box Errors in Near Real-Time and What to Do About It..... 1148

Dave Belt - Irdeto
Preventing Unwelcome Guests in Your Home
..... 1159

L. Alberto Campos, Ph.D.; Zhensheng (Steve) Jia, Ph.D. - CableLabs
Larry Wolcott - Comcast
Proactive Network Maintenance Evolution to the Optical Domain in Coherent Optics..... 1168

Zhou Wang, Ph.D., FIEEE - University of Waterloo
Abdul Rehman, Ph.D. - SSIMWAVE Inc.
Quality-of-Experience Monitoring, Optimization and Management: A Unified End-to-End Solution..... 1209

Dr. Claudio Righetti; Emilia Gibellini; Florencia De Arca; Mariela Fiorenzo; Gabriel Carro - Telecom Argentina
Real-time Analytics for IP Video Multicast
..... 1221

Neill A. Kipp - Comcast
Running a Multi-tenant Hybrid Cloud for Large Scale Cable Applications..... 1250

Tom Conklin - Ericsson
Securing a Hyper-Connected Society..... 1266

Francesco Dorigo; Bao Nguyen; Daniel Howell - Comcast
Self-Service Dimensional Data Analysis: Scalable Patterns for Data-Driven Enterprises
..... 1275

Bradley May - Accenture
Shifting Left: Harnessing AI to Deliver a Consistent, Engaging Customer Experience
..... 1292

Prabhu Navali; Raj Nair - Ericsson Media Solutions
Software-Defined Service Orchestration for MABR TV Services 1299

Steve Goeringer; Dr. Jason Rupe - CableLabs
Solving All Our Problems...Sort of...Blockchain Integrity, Security, and Reliability for Cable Use Cases..... 1313

K. Scott Helms - Momentum Telecom
Supporting the Changing Requirements For Online Gaming..... 1328

Arvinder S. Anand - Ericsson
The Benefits of Leveraging Multi-Vendor Orchestration to Achieve True Service Agility
..... 1361

Sandeep Katiyar - Nokia Bell Labs Consulting
The Emerging Impact and Use Cases of Blockchain Technology in the Era of HFC Connected People and Things..... 1371

Jean-Philippe Joseph; Amit Mukhopadhyay; Ashok Rudrapatna; Carlos Urrutia-Valdés; Tom Van Caenegem - Bell Labs Consulting, Nokia
The Future of Fixed Access - A Techno-economic Comparison of Wired and Wireless Options to Help MSO Decision Process..... 1385

J.R. Flesch; Charles Cheevers - ARRIS International plc
The New Home as a Hotspot: Wi-Fi Meet CBRS LTE and Meet Your Long Range Brother LoRA..... 1411

Kieran Mulqueen; Michael O'Hanlon; Marcin Spoczynski; Brendan Ryan; Thijs Metsch; Leonard Feehan; Ruth Quinn - Intel
Toward Automated Intelligent Resource Optimization for vCMTS Using Machine Learning..... 1455

Bruce McLeod - Cox Communications
Trends on the Role of Internet of Things Technology in Senior Care..... 1467

Bernard Burg; Fan Liu; Abel Villca Roque; Sunil Srinivasa; Ryan March; Tianwen Chen - Comcast
Using AI to Improve the Customer Experience
..... 1480

Jason Rupe; Colin Justis - CableLabs
Using Historical Traffic Data to Schedule Service Interruptions for Minimum Customer Impact..... 1495

Venk Mutalik; Dan Rice; Karthik Subramanya; Jon-
en Wang - Comcast

***What Gets Measured Gets Done/What Gets
Analyzed Gets Transformed: Analytics for a
Wider/Deeper Network View..... 1511***

Sandy Wilbourn; Craig Sprosts - Akamai

***When Privacy and Security Collide: New
Approaches are Needed 1532***

5G Is Rapidly Approaching, What Must Cable MSOs Do To Capitalize On This Business Opportunity

A Technical Paper prepared for SCTE•ISBE by

Jon Baldry
Metro Marketing Director
Infinera
125 Finsbury Pavement, London, EC2A 1NQ
+44 7766 146 440
jon.baldry@infinera.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. 2018, The Year That 5G Arrives	3
2. The Cable MSO 5G Opportunity	4
3. 5G Services driving 5G Transport Network Revolution	5
4. Introducing X-Haul.....	7
4.1. X-Haul Performance and Architecture	8
4.2. 4G/5G Coexistence	9
4.3. Multi-Service Edge Computing.....	10
Conclusion.....	11
Abbreviations	11
Bibliography & References.....	12

List of Figures

Title	Page Number
Figure 1 - 5G Standardization timeline	4
Figure 2 - 5G Services with Bandwidth and Delay/Latency Specifications	5
Figure 3 - Data Center Distribution to Support MEC	6
Figure 4 - 5G Services by Focus Area, source ITU-R	7
Figure 5 - Fronthaul and Backhaul Evolution For 5G	8
Figure 6 - Demanding Timing and Synchronization Requirements For 5G	9
Figure 7 - 4G and 5G Cell Coverage	10

Introduction

Over the last year there has been a dramatic shift within the wireless industry in terms of 5G preparation. As the 5G standards have solidified, many wireless operators have shifted their transport network mindset from “4G transport that can evolve to 5G” to “5G-Ready now”. This has had a major impact on how wholesalers, and particularly cable MSOs who plan to sell transport services to wireless operators, need to address the market opportunity.

The most significant advance in wireless transport standardization over this period has been the additional clarity around the eCPRI standard and the performance demands this puts on the transport network. In 5G the underlying transport infrastructure moves from a collection of dumb pipes, albeit high performance pipes, to a highly dynamic network supporting:

- Even higher demands on performance – Low latency, synchronization and higher capacity demands are a given with 5G.
- Multiple transport functions integrated into a single network – The network must support legacy 4G infrastructure in parallel to new 5G Fronthaul I (low-split), Fronthaul II (high-split) and backhaul in a X-haul/any-haul environment.
- Network slicing and support for virtualization of mobile infrastructure – dynamic movement of 5G resources around the network to support Multi-access Edge Computing (MEC) and fog networking requires dynamic transport that is SDN controlled and cloud-optimized. Network slicing at all layers will also play a critical role in supporting this environment.

To address these challenges, MSOs need networks that are flexible and open, and offer high performance now. Some wireless operators are already testing wholesale networks for key 5G performance metrics, such as eCPRI synchronization requirements, to prepare themselves for 5G.

This paper will describe the challenges associated with 5G and show how MSOs must evolve their transport services to adapt and grasp the exciting opportunity that 5G presents to the industry. This paper will address the issues, challenges and opportunities associated with the fiber-based footprint and services. Cable MSOs may well also be able to utilize DOCSIS-based services over the coax plant to augment the capabilities outlined in this paper to extend services further into access networks.

Content

1. 2018, The Year That 5G Arrives

2018 is undoubtedly the start of the eventual 5G onslaught with initial 5G services hitting the street in numerous countries. With the rush to be local leaders in 5G there is considerable variety around the services currently labeled as 5G. Most, but not all, conform to the latest 5G New Radio (NR) standard which is the new air interface for 5G. This means that any 5G NR compliant device should be able to connect to the network and deliver mobility services. This is of course very early days for 5G and initial 5G services over the next couple of years will focus on high speed broadband services, essentially higher speed offerings of the 4G services we enjoy today.

But 2018 also sees the next stage of standardization in 5G. The initial non-standalone (NSA) 5G standard was completed in late 2017 and focused on utilizing the existing 4G/long term evolution (LTE) radio access network (RAN) and core to deliver enhanced mobile broadband (eMBB) services. 2018 will see

the completion of the next stage of 5G standardization, the standalone (SA) variant, which utilizes the new 5G RAN and core network, as shown in Figure 1.

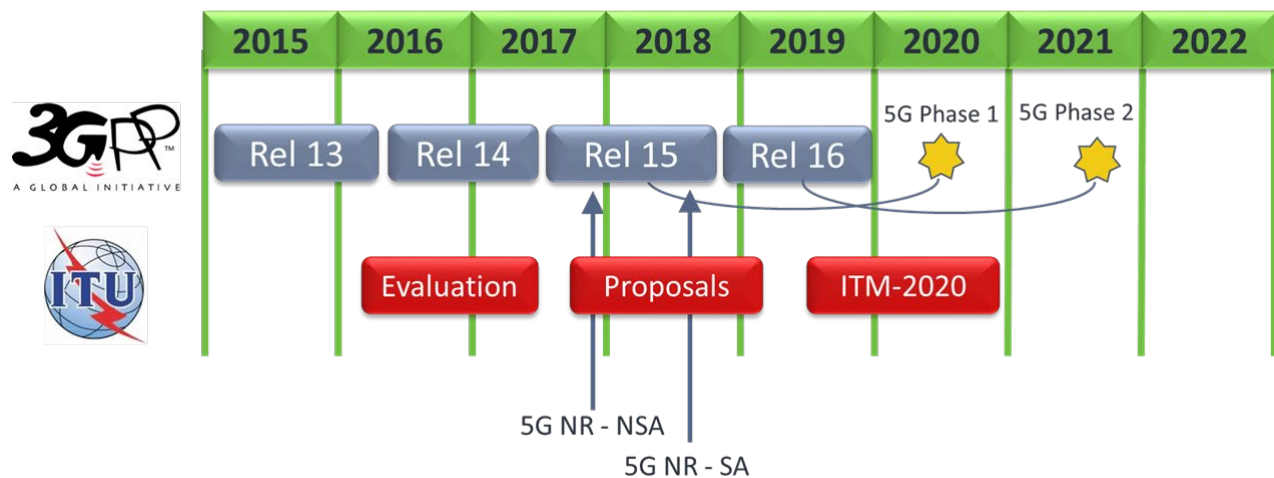


Figure 1 - 5G Standardization timeline²

Ultimately 5G offers the potential for a wide range of services beyond eMBB and these will fully utilize the capabilities of the new 5G network, both RAN and core. Key to these new capabilities will be significant advances in the mobile transport network connecting the RAN to the core. It will however take many years for the full portfolio of envisaged 5G services to fully rollout, but transport networks need to be 5G-ready now to avoid unnecessary and traffic impacting upgrades in the future. This will enable support for all the envisaged services when mobile operators are ready and any new killer apps that we haven't yet considered that can utilize the high performance that 5G has to offer.

2. The Cable MSO 5G Opportunity

Given an ideal world, many mobile operators would like to build their networks with their own fiber or dark fiber from other operators. This would give them total control of their network, but it has many challenges. The cost of building a dedicated fiber network would most likely be overly onerous even if the high cost of new 5G licenses wasn't also taken into consideration. Digging new fiber isn't an option when alternative fiber from wholesale operators exists and is only an option to push fiber deeper into access networks to support new cell sites. Some dark fiber providers will provide dark fiber to mobile operators in some geographies but in others with differing market dynamics, wholesale operators will offer wholesale services rather than dark fiber.

The 5G standards foresee a world where network resources are shared and once the initial rush to grab local headlines with early 5G services is over, network operators will look to use their capital carefully to build unique capabilities when it gives differentiation and to use lower cost shared resources when it doesn't.

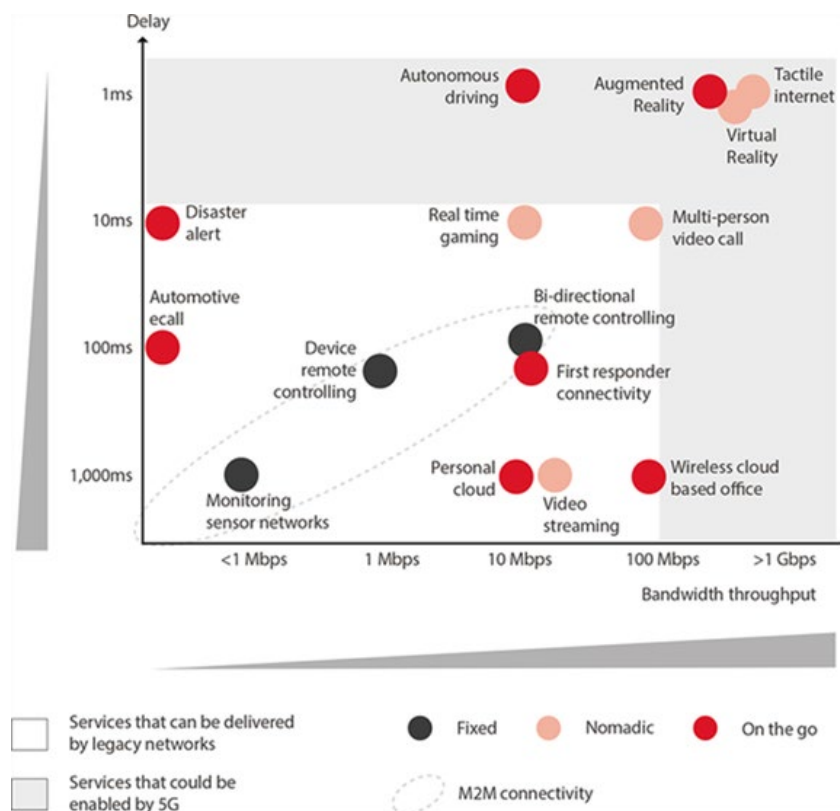
5G creates an opportunity for cable MSOs to utilize their fiber and hybrid fiber coax (HFC) resources to extend any current wholesale mobile services. This is due to the increased pressure on mobile operators to fiber up an ever-growing number of cell sites with high performance transport. The economics of each mobile operator building their own dark fiber-based networks simply don't stack up. As mobile operators look towards wholesale operators, cable MSOs have a unique advantage in their service area due to their extensive fiber footprint that is often in the same places as the mobile operators 5G plans, residential and

business areas. This local access fiber footprint asset will be expanded hugely over the next few years as distributed access architecture (DAA) rolls out pushing fiber deeper from the secondary hub to remote PHY devices. This converts access fibers from analogue optics to digital wavelength division multiplexing (WDM)-based optics, creating the opportunity for a multi-service converged interconnect network (CIN) supporting DAA, mobile and other services.

Key to this 5G wholesale opportunity for cable MSOs is the high performance that transport networks will require in the future. Offering differentiated high-performance services built on a multi-service platform will enable MSOs to provide high quality services with the better economics for both the mobile operator and the cable MSO.

3. 5G Services driving 5G Transport Network Revolution

To understand the requirements on the transport network we first need to consider the range of services that 5G will potentially bring and the requirements these put on the underlying transport network. The plans for 5G services cover a broad range of services from the Internet of Things (IoT), connected and self-driving cars, augmented and virtual reality plus those new killer applications that we haven't envisaged yet. These services require a radical change to the mobile transport network to support the new 5G radio access network (RAN) technology and supporting mobile core functions, enabling them to reach the bandwidth and delay/latency requirements shown in figure 2.



Source: GSMA

Figure 2 - 5G Services with Bandwidth and Delay/Latency Specifications

Existing LTE/4G networks already require high-performance transport to support advanced functionality such as coordinated multi-point (CoMP) and enhanced intercell interference coordination (eICIC). These features will also be used in 5G and require performance characteristics such as high network resilience and low latency to ensure operation. In some geographic regions, high-quality synchronization performance to allow the transport network to deliver frequency, phase and time-of-day synchronization to the cell site. In regions where 4G/LTE synchronization is delivered via other means such as GPS, it is envisaged that 5G cell sites will drive the requirement for alternative means of synchronization as small cells move into locations that aren't suitable to GPS, e.g. inside shopping malls, subway stations and deep within buildings.

From a transport network architecture perspective, 5G brings a radical rearchitecting of the network to support the higher bandwidth, lower latency and tighter synchronization between cell towers required for the envisaged new services. Capacity increases by at least a factor of 10, and the complete network latency is lowered by a factor of 10, from 10 to 1 millisecond. This will be achieved via a two-pronged focus on latency within the transport network and a migration of some of the compute and storage resources previously located in the core of the network out into the network closer to the end user via multi-access edge computing (MEC).

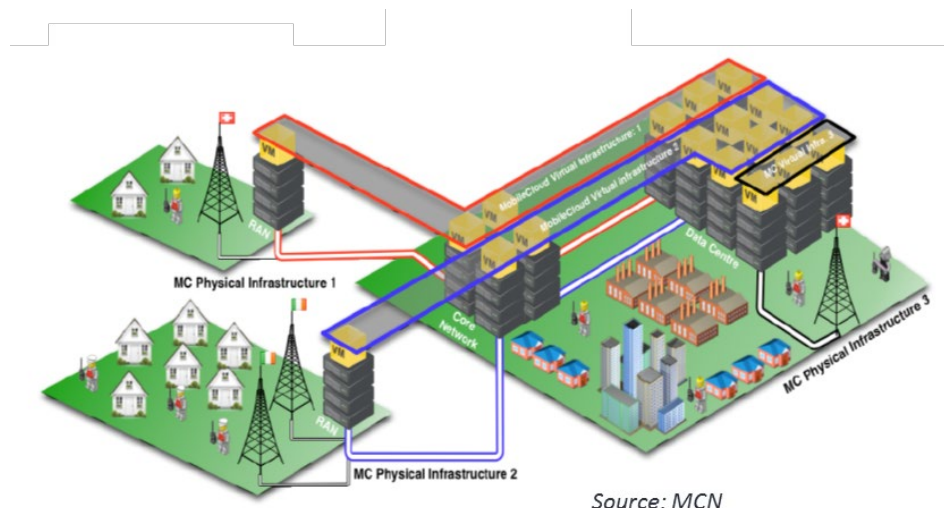


Figure 3 - Data Center Distribution to Support MEC

These MEC capabilities will also be virtualized software packages rather than the previous bespoke hardware platforms, which means the 5G network will become a collection of large core data centers and smaller edge/mini datacenters, as shown in figure 3. These data centers are able to move virtualized capabilities around the network and dial up or down the associated processing power. In order to support this fluid bandwidth and services environment, mobile transport networks will need to migrate from today's collection of effectively dumb pipes, albeit high-performance dumb pipes, to a dynamic network with tightly integrated control mechanisms to the wider 5G network.

One particular challenge for transport networks is that there isn't a single set of performance criteria and capabilities that are optimum for 5G. Figure 4 shows the International Telecommunication Union's radiocommunication sector (ITU-R) view of 5G services and clusters them around three core service types – eMBB, massive machine type communications and ultra-reliable and low latency communications. Each of these 3 focus areas requires differing transport optimization and therefore any

5G transport network will need to blend the technical requirements and economics of bandwidth transport to support these as best as possible.

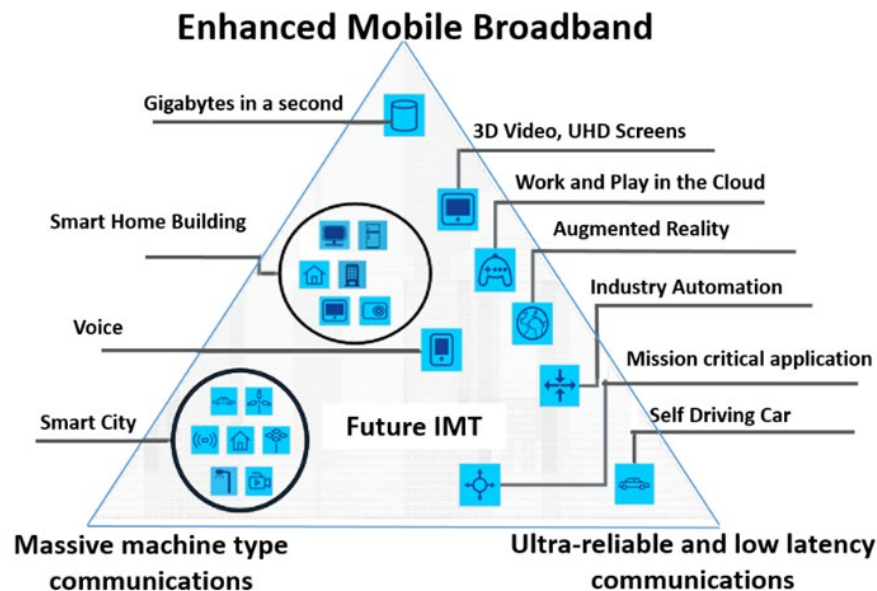


Figure 4 - 5G Services by Focus Area, source ITU-R

4. Introducing X-Haul

Today 4G/LTE networks use a fronthaul and backhaul architecture to support communication between the remote radio head (RRH), the baseband unit (BBU) and the core network. In many cases the RRH and BBU are collocated so only backhaul is required. In other cases networks have started the migration to the cloud-RAN (CRAN) architecture³. Here BBUs are clustered together a maximum of 20 kilometers from the RRH to better support capabilities such as CoMP and eICIC. Essentially from a transport perspective these are two separate networks with the fronthaul network using common public radio interface (CPRI) for RRH to BBU traffic and backhaul using Ethernet for BBU to core. They can however run over a common packet optical physical network to share fiber resources with a packet optical platform that supports both networking types on a wavelength by wavelength basis.

In 5G, this will be required to migrate to a new X-haul architecture, with Ethernet-based backhaul necessary to support the required higher performance for 5G, and fronthaul migrating from CPRI to an Ethernet-like enhanced CPRI (eCPRI) network. This new architecture builds a common network for both traffic types and eCPRI itself comes in high-split and low-split options to give fronthaul-like or backhaul-like performance to match the requirements RRH to distributed unit (DU) or DU to centralized unit (CU) traffic, as shown below.

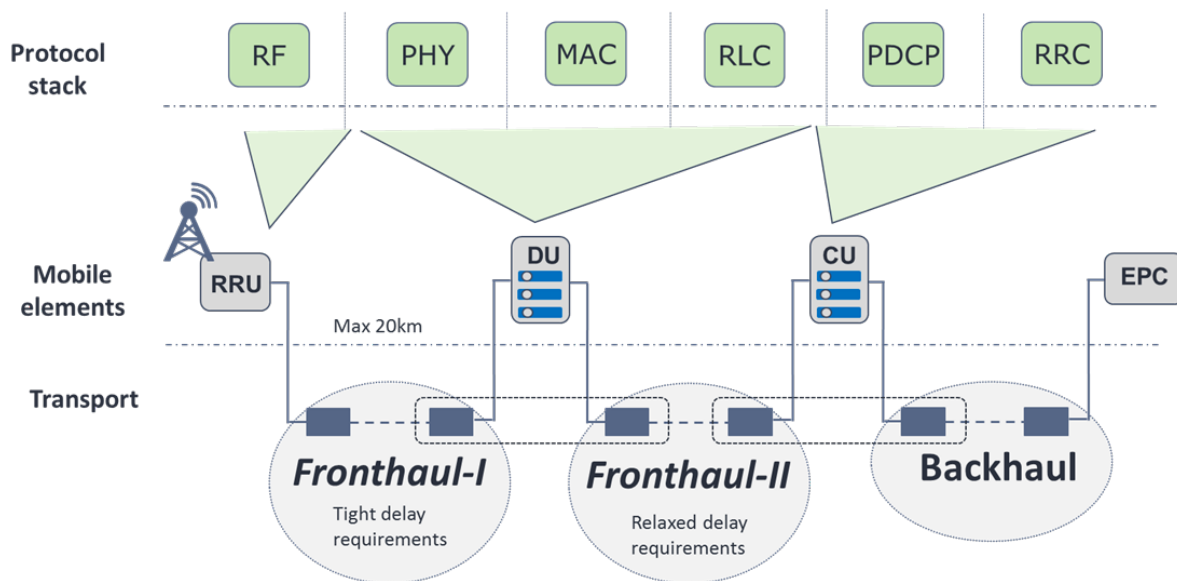


Figure 5 - Fronthaul and Backhaul Evolution For 5G

To support this range of traffic options over a common network, packet-optical-based mobile transport networks will be required to support Ethernet and eCPRI traffic over a combined any-haul/cross-haul/X-haul network, ensuring that each traffic type is supported with the specific performance requirement.

4.1. X-Haul Performance and Architecture

Transport performance is important in many packet-optical networks where Ethernet or IP functionality is required but the service still requires layer 1 like characteristics. Mobile and particularly 5G is a good example of one of these scenarios. These transport performance characteristics include low latency, low jitter (ideally zero) to give layer 1 like constant latency through any particular route through the network and for those networks where high-performance synchronization is delivered to the cell site via the transport network, then high quality synchronization is also required through the network. This includes frequency, phase and time-of-day synchronization.

As an example, LTE-A/4G features such as CoMP and eICIC already require tight synchronization performance of +/- 1.5 microseconds for phase synchronization, so 3 microseconds total budget. The recently completed eCPRI specification¹ reduces this to 65 or 130 nanoseconds for the highest 5G performance services such as beamforming multiple input multiple output (MIMO), roughly a 30-fold improvement in synchronization performance for the network as shown in figure 6. This high level of synchronization performance is only really needed once operators are ready to implement the most demanding 5G RAN features, to support the most demanding 5G services such as the ultra-reliable and low latency communications shown earlier in figure 4, and this is currently viewed to be a few years away yet.

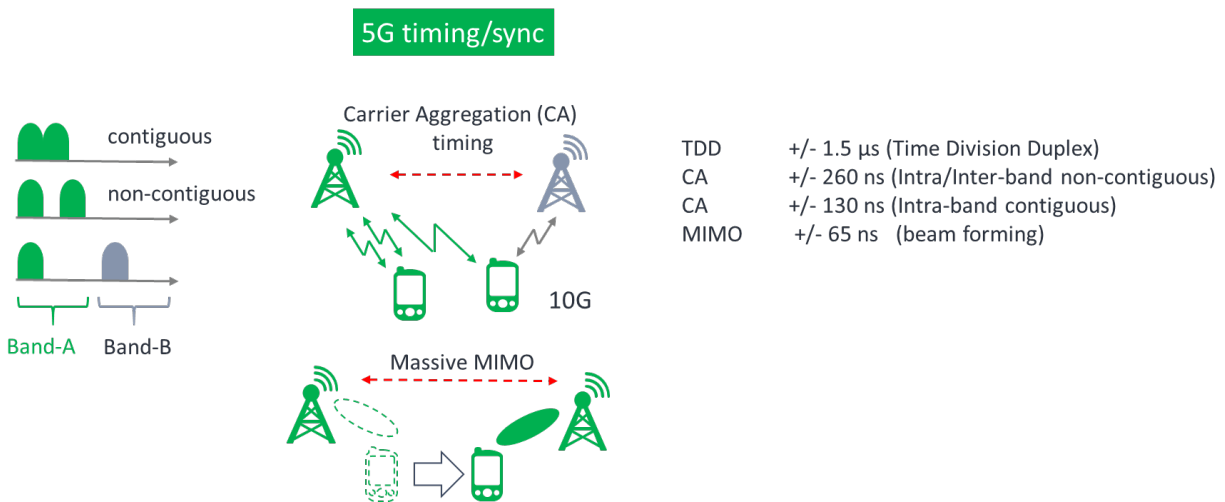


Figure 6 - Demanding Timing and Synchronization Requirements For 5G

But synchronization performance is not something that can be added later via software upgrades as it relies on key design decisions embedded deep in the hardware and software architecture of the platform. 5G-ready mobile transport networks need to be ready for everything 5G RAN will throw at them, without the need for a future network upgrade to fix substandard synchronization performance. So, for cable MSOs to prepare true 5G-ready mobile transport networks, they need to strongly consider the need this high level of performance now. This level of network synchronization is tough to achieve but it is possible in today's 5G-ready networks.

Having 5G-ready levels of synchronization and low-latency performance now can offer cable MSOs competitive advantages today for wholesale mobile services. It has been shown that high quality sync can also help improve existing 4G services. Results from the field have shown that network performance can improve at the handset level with up to 80% better download speeds, 40% better upload speeds and a 40% reduction in latency when backhaul is migrated to packet-optical with high quality sync performance.

One further consideration for any wholesale operator is the ability to support multiple sync domains on a single transport network. Assuming sync as a service is applicable to the end customers then any wholesaler needs to build a network that can support multiple sync domains over the same common infrastructure to enable the economic advantages of a shared network. If the platform used can't support multiple sync domains in parallel, then each new mobile end customer will drive parallel packet-optical platforms per end mobile operator customer with very little economic gain.

4.2. 4G/5G Coexistence

In previous mobile generational upgrades, the move from one generation to the next was a completely new network from edge to core. However, for the migration from 4G to 5G, 4G doesn't go away in the same way that previous generational changes did, in fact 4G infrastructure plays a critical role in 5G. Much of the 4G standardization over recent years was done to enable 4G to support 5G. 4G networks already adequately support today's levels of video streaming, internet browsing, etc., which means new 5G infrastructure should focus on supporting new high-performance services. This is one of the key drivers for CoMP where a 5G device will communicate with multiple cells in parallel to support multiple services or to overcome issues introduced by the new millimeter wave technology needed for 5G. This is one of the primary drivers for high performance synchronization between cell sites.

5G will also use millimeter wave technology to provide the required higher RAN bandwidth, but this comes at the cost of reach and vulnerability to blocking or reflections from objects such as buildings, trees and metallic objects such as cars or production line machinery in in-building networks. Consequently, 5G cells will be smaller, meaning we'll need a lot more of them and it will take a very long time for 5G to reach the same coverage as 4G, if it ever does. If 4G cells are measured in kilometers, then 5G cells will be measured in hundreds of meters, as shown in Figure 7.




Figure 7 - 4G and 5G Cell Coverage

The result of this co-existence is that any MSO wholesale services for 5G need to consider the hybrid 4G/5G scenario and potentially continue to offer 4G fronthaul and backhaul services in addition to 5G X-haul, ideally over the same packet-optical platform.

4.3. Multi-Service Edge Computing

We discussed the introduction of MEC earlier in this paper as a means to help mobile operators reduce network latency for services that require latency performance that is better than current 4G services. To achieve this goal, compute and storage resources are moved from the core datacenters to locations closer to the edge of the network and the end user. These will then run virtualized functions needed for 5G and provide local content or data caching. Essentially the cloud moves closer to the user and is spread around the network, hence this is sometimes referred to as fog networking. From a transport network perspective this drives several requirements:

- SDN control and orchestration – As transport networks will need to dynamically dial up and down bandwidth to match the requirements of the network, open SDN-based application programming interface (API)s will be required. These will allow bandwidth within the network to respond to the overall requirements of the 5G network.
- Highly flexible networking – bandwidth adjustments throughout the network from the cell site to the core need to utilize technology that allows for dynamic use of bandwidth on demand.

- Network slicing – This is a concept introduced in 5G where resources can be sliced to carve out a segment of the network to specifically serve a function, application or customer. Slicing will be done at both the transport and control planes enabling a shared resource, the network, to support either multiple mobile operators, multiple industry segments with specific performance parameters or perhaps multiple applications within a carrier or segment.

MEC and network slicing will be driving forces for SDN control and orchestration of the transport network bringing open APIs to the control plane and the ability to slice both the data plane and the control plane. Mobile operators will need to encompass wholesale services within their plans, requiring cable MSOs offering wholesale services to tightly align with their control and orchestration platforms.

Conclusion

2018 is the ideal time for cable MSOs to reassess their approach to wholesale mobile services. Over the last 12 months 5G standards have matured considerably giving much more clarity on the exact requirements of the mobile operator. In addition, cable MSOs are currently planning their own fiber deep DAA networks that will push fiber and WDM-based packet optical technology deeper into the access creating the opportunity for a multi-service CIN to also support mobile wholesale services.

Cable MSOs who sell connectivity services to mobile operators must consider transport equipment that enables them to roll out 4G fronthaul and backhaul today but that also has the flexibility to migrate to the hybrid X-haul network of the future, supporting 5G X-haul concurrently with remaining 4G fronthaul and backhaul. This requires highly scalable and high-performance transport with low latency, zero jitter and high-quality synchronization to enable the cable MSO to differentiate their services in what will be a very competitive access market.

Lastly 5G will drive considerable changes into the overall transport network architecture with SDN control and orchestration, MEC and fog networking coming to the fore.

Mobile operators have a huge task ahead of them as they race to win in 5G and cable MSOs who can simplify their transport requirements leaving them to focus on the new 5G services that the public is eagerly anticipating have a lot to gain.

Abbreviations

API	Application programming interface
BBU	Baseband unit
CIN	Converged interconnect network
CoMP	Coordinated multi-point
CRAN	Cloud or Centralized RAN
CU	Centralized unit
DAA	Distributed Access Architecture
DU	Distributed unit
eICIC	Enhanced intercell interference coordination
eMBB	Enhanced mobile broadband
HFC	Hybrid fiber coax
ITM	International mobile telecommunication
ITU-R	International Telecommunications Union, Radiocommunications Sector

LTE	Long term evolution
LTE-A	Long term evolution - advanced
MEC	Multi-access edge computing
MIMO	Multiple input multiple output
NR	New radio
NSA	Non-standalone
RAN	Radio access network
RRH	Remote radio head
SA	Standalone
WDM	Wavelength division multiplexing

Bibliography & References

1 - *eCPRI Transport Network V1.2 (2018-06-25); Common Public Radio Interface: Requirements for the eCPRI Transport Network*; CPRI

2 - *IMT for 2020 and beyond*; ITU-R

3 - *C-RAN - The Road Towards Green RAN*; China Mobile Research Institute

5G Small Cells and Cable

Realizing the Opportunity

A Technical Paper prepared for SCTE•ISBE by

Dave Morley

Director, 5G and Regulatory

Shaw Communications Inc. / Freedom Mobile

2728 Hopewell Place NE, Calgary, Alberta T1Y 7J7

+1-403-538-5242

dave.morley@sjrb.ca

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
5G Drivers	4
1. Demand.....	4
2. Technology.....	5
3. Standards	6
4. Spectrum	6
5. Use Cases	7
Network Densification and Small Cells	9
6. Challenges & Opportunities	10
6.1. Site Acquisition.....	10
6.2. RF Coverage.....	12
6.3. Backhaul.....	13
6.4. RAN Architectures.....	14
6.5. Synchronization.....	17
6.6. Converged Access Network.....	18
Conclusion.....	19
Abbreviations	19
Bibliography & References.....	20

List of Figures

Title	Page Number
Figure 1 - Global Mobile Data Traffic (Source: Cisco Visual Networking Index 2017)	4
Figure 2 – Beamforming (Source: Nokia)	5
Figure 3 - Network Slicing (Source: Nokia).....	6
Figure 4 - 3GPP 5G Standards Roadmap (Source: 3GPP).....	6
Figure 5- Global 5G Spectrum Allocations (Source: Qualcomm)	7
Figure 6 - 5G Capabilities	8
Figure 7 - 5G Use Cases (Source: adapted from Nokia).....	8
Figure 8 - Network Densification	9
Figure 9 - 5G Network Deployments by Spectrum Band	10
Figure 10 - 4G Small Cell Deployments.....	10
Figure 11 - Strand-Mount 4G Small Cell Deployments.....	11
Figure 12 - Typical Strand-mount Small Cell Installation.....	12
Figure 13 – Outdoor small cell coverage predictions for LTE @ 2500 MHz (left) and 5G NR @ 3500 MHz (right).....	13
Figure 14 - Possible 5G NR Functional Splits (Souce: adapted from Nokia)	14
Figure 15 - CPRI Line Rates	16
Figure 16 - Evolution of DOCSIS Aggregate Line Rates	16
Figure 17 - LTE Carrier Frequency Error on a production DOCSIS network	18

Figure 18 - Converged DAA/5G Access Network	18
---	----

List of Tables

Title	Page Number
Table 1 - 4G & 5G Small Cell Synchronization Requirements	17

Introduction

As wireless networks evolve from 4G/LTE to 5G, small cells will play a critical role in delivering the high bandwidth, low-latency connections required by the myriad potential 5G use cases. With regulators opening up new low, mid and high-band millimeter wave (mmWave) spectrum for 5G, MSOs are uniquely positioned to create ultra-dense 5G wireless networks by leveraging their existing hybrid fiber coax (HFC) networks and emerging technologies such as full duplex DOCSIS (FDX) and distributed access architectures (DAA). This paper examines the drivers for the adoption of 5G and the potential opportunities and challenges presented by 5G small cells based on insights from Shaw Communications' recent deployments of 4G/LTE small cells, through its subsidiary Freedom Mobile.

5G Drivers

5G promises to dramatically transform the role that mobile technology plays in the world, creating new economic opportunities and promoting social development. By 2020, 5G is expected to add \$1 trillion to the economy of North America alone [1]. It will also impact all aspects of our society from facilitating smart-city energy grids and the Internet of things (IoT) to enabling autonomous vehicles and remote healthcare. The key drivers behind the shift to 5G include growing global demand for wireless services, ongoing advancement of wireless technologies, emergence of new use cases, development of global standards and the availability of new spectrum bands for 5G. These drivers are explained in more detail in the following subsections.

1. Demand

Global demand for mobile data continues to grow unabated as subscribers increasingly consume and share video content. In addition, as mobile devices become even more capable with added processing power, higher-resolution screens, and advanced features such as gesture-based interfaces, consumers are demanding faster, more reliable mobile connectivity.




Figure 1 - Global Mobile Data Traffic (Source: Cisco Visual Networking Index 2017)

As shown in Figure 1, the Cisco Visual Networking Index (VNI) predicts that global mobile data traffic will increase 7-fold between 2016 and 2021, representing a cumulative average growth rate (CAGR) of

47% [2]. Cisco also forecasts that average mobile connection traffic per month will grow by over 5 times during the same period, from 1 GB in 2016 to 5.4 GB in 2021. The same study predicts that the number of mobile-connected devices per capita will reach 1.5 by 2021.

2. Technology

Another important driver for the adoption of 5G is the rapid evolution of wireless technology. Two key developments are massive multi-input multi-output (MIMO) systems and beamforming antennas, which together offer improved network coverage and traffic capacity. As shown in Figure 2, beamforming creates highly-focused beams from the base station to individual users. This counteracts propagation losses, particularly at mmWave frequencies and minimizes interference from other sources, both of which allow higher data rates to be supported than would otherwise be the case.




Figure 2 – Beamforming (Source: Nokia)

Massive MIMO in turn uses multiple antennas to simultaneously serve many users by exploiting multiple radio paths to increase system capacity.

Another key technology driving 5G is advanced channel coding techniques, which provide more efficient and robust data transmission compared with previous coding methods, supporting larger data blocks and more reliable control channels. A scalable OFDMA air interface has also been developed that efficiently addresses different spectrum, deployment and service scenarios by adopting a flexible subcarrier spacing and slot structure.

The 5G core network is also evolving to fully embrace network function virtualization (NFV) and software defined networking (SDN), which together will enable faster, more agile and scalable networks. As shown in Figure 3 below, network slicing will also allow core networks to be dynamically configured on a per vertical or per service basis.



Figure 3 - Network Slicing (Source: Nokia)

3. Standards

5G standards are also proceeding rapidly with the recent finalization of 3GPP Release 15 in June 2018. This release defines the 5G new radio (NR) standards for both non-standalone (NSA) and standalone (SA) operation. NSA operation allows inter-working with the existing 4G evolved packet core (EPC), while standalone (SA) operation is for the new 5G core network (CN). As shown in the figure below, the 5G CN standards are expected to be completed in Release 16 in the 1st half of 2019.




Figure 4 - 3GPP 5G Standards Roadmap (Source: 3GPP)

4. Spectrum

The allocation of low, mid and high band (mmWave) spectrum for 5G is also well underway or completed in several countries around the globe, including the U.S., Canada, Japan, China Korea, and European countries (see Figure 5).



Figure 5- Global 5G Spectrum Allocations (Source: Qualcomm)

In the U.S. and Canada, the initial bands designated for 5G include the 600 MHz, 3500 MHz and 24-28 GHz bands. Additional spectrum bands under consideration include the 3.7-4.2 GHz, 37-40 GHz and 64-71 GHz bands. The use of high-frequency mmWave spectrum bands above 24 GHz is emerging as a key 5G enabler. The use of these bands is very compelling because of the large bandwidths available at mmWave frequencies, enabling extremely high data rates and significant increases in capacity.

5. Use Cases

The key technologies noted above coupled with supporting standards and spectrum, create a range of fundamental new capabilities illustrated in Figure 6. As defined in IMT-2020 [3], these capabilities include 10x higher data rates, 10x lower latency, and 10x higher network efficiency than previous 4G networks.

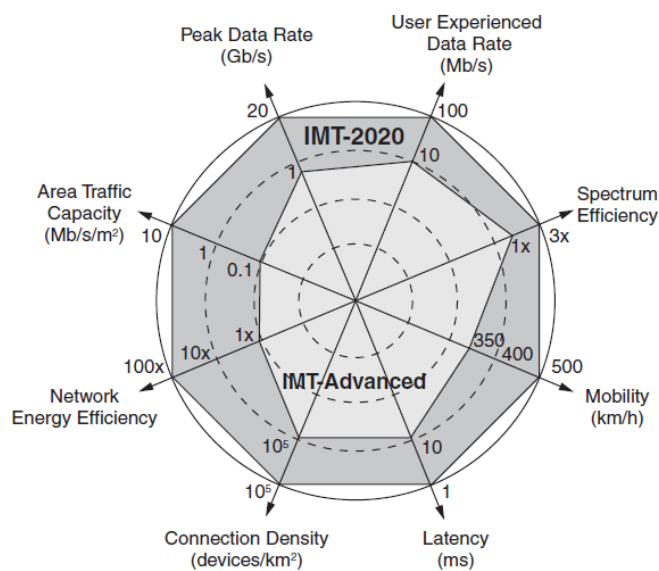


Figure 6 - 5G Capabilities

Given these new capabilities, 5G is expected to support myriad use cases, some of which haven't even been envisaged yet. Those that have can be grouped into 4 main categories: fixed wireless access (FWA), enhanced mobile broadband (eMBB), machine type communications (mTC), and ultra-reliable low latency communications (uRLLC), as shown in Figure 7.

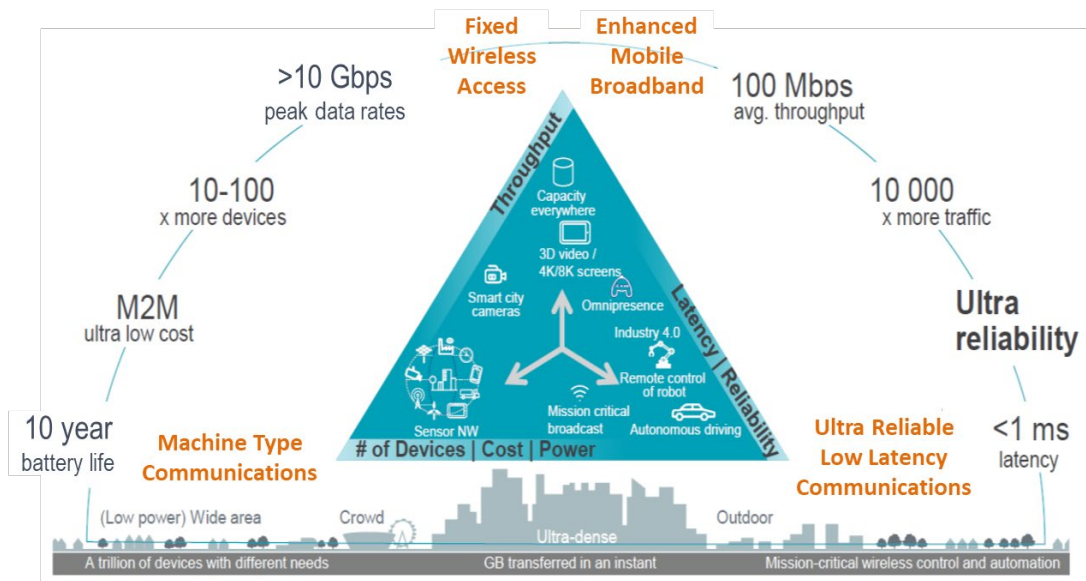


Figure 7 - 5G Use Cases (Source: adapted from Nokia)

FWA promises to deliver multi-gigabit speeds to fixed residential and business premises using mmWave spectrum. As such, FWA is an alternative to wired broadband services and represents a potential threat to traditional wireline providers including MSOs. At the same time, FWA also provides MSOs with the opportunity to deliver broadband services in “brown-field” areas that are not currently served by the

operator such as new residential markets and industrial/business parks. FWA can also be used to provide critical small cell backhaul in areas without existing wireline infrastructure.

eMBB, with its multi-gigabit per second speeds and inherent mobility, will support new immersive experiences such as virtual and augmented Reality (VR/AR), high resolution video, and “flash download/upload” uses cases such as software updates and video downloads. eMBB will also be used to support highly congested environments such as stadiums and airports.

At the other end of the scale, mMTC will require much lower throughputs but will need to support billions of devices, often with very low energy consumption requirements. mMTC has numerous applications in sensor networks, smart cities, agriculture, retail, logistics, and a plethora of other use cases in a world of increasing connected things.

URLLC encompasses mission critical real-time applications such as industrial control, remote robotics, remote surgery/health and autonomous driving.

Network Densification and Small Cells

While 5G opens up a world of new use cases, it also comes with a number of important implications. One that represents both a challenge and an opportunity, particularly for MSOs, is network densification. That is, to realize the higher data rates and area traffic capacities promised by 5G, operators will need to deploy much denser network topologies using small cells. Figure 8 illustrates the transition that needs to be made from today’s networks, which are larger composed of macro cells, to 5G networks of the future with the targeted deployment of small cells.




Figure 8 - Network Densification

For eMBB applications, for example, small cells will likely be deployed in both mid-band and mmWave bands depending on the type of environment. As shown in Figure 9, mid-band (e.g., 3.5 GHz) small cells are best suited for dense urban areas, whereas mmWave (e.g., 28 GHz) small cells will typically be deployed in extremely high traffic areas such as stadiums, airports and pedestrian zones. In addition, mmWave small cells may also be deployed for FWA or flexible use deployments, serving both fixed and mobile subscribers. In contrast, 600 MHz will most likely be deployed at macro sites to provide wide area and in-building coverage.



Figure 9 - 5G Network Deployments by Spectrum Band

6. Challenges & Opportunities

The main challenges associated with deploying 4G/5G small cells include site acquisition, backhaul, synchronization and powering.

6.1. Site Acquisition

Certainly, one opportunity for the cable industry is to leverage their existing Wi-Fi hotspot and/or 4G small cell locations to deploy 5G small cells. With access rights to millions of public Wi-Fi hotspots, often in prime “beach front” properties, the cable industry is uniquely positioned to capitalize on this opportunity, either as a wireless player or a wholesale provider to existing MNOs.

At Freedom Mobile this strategy is currently being used to deploy 4G/LTE-A small cells. Figure 10 shows an indoor small cell that was installed at an existing Wi-Fi hotspot (left) and a new outdoor small cell at a transit station (right). In the future, both locations will be used for the deployment of 5G small cells.




Figure 10 - 4G Small Cell Deployments

In addition, MSOs can leverage their existing cable infrastructure, and in particular aerial plant, to quickly deploy small cells in targeted areas. Deploying small cells on aerial plant addresses three key challenges with small cell deployments: site access, backhaul and power. That is, access is usually already covered by existing pole-line attachment agreements and both backhaul and power are available on the coaxial cable plant.

Figure 11 shows one such deployment in a major metro area by Freedom Mobile. In this case, strand-mount 4G/LTE-A small cells are being installed on aerial strand with inter-site distances ranging from 175 to 225 meters.




Figure 11 - Strand-Mount 4G Small Cell Deployments

Figure 12 shows a typical strand-mount small cell installation, which consists of a small cell gateway, an 4G/LTE-A small cell and directional coupler. The small cell gateway contains a DOCSIS 3.1 cable

modem and power supply, which converts 90 V_{AC} quasi-square wave input power (i.e., Power-over-Cable) to 115 V_{AC} output power for the small cell.




Figure 12 - Typical Strand-mount Small Cell Installation

The 4G/LTE-A small cell can be equipped with up to 3 radio frequency (RF) modules, each capable of supporting 2 x 2 MIMO with up to 2 x 5W transmit power. For dense pedestrian zones, the small cell is typically equipped with a license-assisted access (LAA) RF module, which operates in the 5 GHz unlicensed band. In this configuration, the small cell supports a peak data rate of up to 450 Mbps using two 20 MHz carriers in the unlicensed band and one 15 MHz carrier in the licensed AWS-3 band. The directional coupler connects the small cell gateway to the hybrid fiber-coaxial (HFC) cable plant, which provides both the power and connectivity for the DOCSIS backhaul circuit. At the cable modem termination system (CMTS), the small cell backhaul circuit is routed to the LTE evolved packet core (EPC) network.

As mentioned, a key advantage of using the HFC cable plant for small cell deployments is the availability of power in addition to backhaul connectivity. As a result, small cells can be deployed quickly and easily wherever the MSO has existing coaxial plant. In contrast, passive optical networks (PONs), for example, require a separate power service to be installed, which adds to the cost, time and complexity of the installation.

6.2. RF Coverage

A key consideration in 5G small cell deployments is meeting radio frequency (RF) coverage requirements. Although inter-site distances for the outdoor small cell deployments in Figure 11 were designed for LTE at 2500 MHz band, comparable coverage can be achieved with 5G NR at 3500 MHz. Figure 13 shows RF coverage predictions for LTE at 2500 MHz and 5G NR at 3500 MHz using Infovista's Planet 7 RF network planning tool. This figure shows that the coverage is virtually identical for both bands.

Simulation studies of 5G NR mmWave network coverage by Qualcomm also show that significant coverage is possible when co-locating mmWave equipment with existing 4G macro and small cell sites [4]. Of the 15 cities studied, the outdoor downlink coverage at 28 GHz varied from 50% to 80% of the existing 4G network coverage in 9 of the 15 cities studied. These results suggest that deploying mmWave small cells in urban areas using existing 4G cell sites is entirely feasible, especially in areas with high-

density deployments. While outdoor-to-indoor coverage is not feasible in the mmWave band due to much higher material attenuation (e.g., low-e glass), providing outdoor mmWave coverage would free-up significant capacity in the sub-6 GHz band for outdoor-to-indoor capacity.




Figure 13 – Outdoor small cell coverage predictions for LTE @ 2500 MHz (left) and 5G NR @ 3500 MHz (right)

For indoors locations, we expect that 5G NR coverage at 3500 MHz to be virtually identical to LTE at 2500 MHz as well. Qualcomm also conducted indoor network coverage simulations at mmWave frequencies with similar results. A simulation study was done for the Las Vegas Convention Center using existing 4G antenna locations in the venue. A one-to-one overlay of 5G mmWave radios on the existing 4G/LTE antenna locations resulted in significant coverage (up to 85%) throughout the facility.

In buildings with existing Wi-Fi access points, the above results suggest that coverage parity might be achieved with a one-to-one mmWave small cell overlay since inter-access point distances for Wi-Fi access points are typically much shorter than 4G small cells. Replacing existing Wi-Fi access points and/or 4G small cells with 5G equipment could result in substantial cost and time savings for 5G deployments.

6.3. Backhaul

As noted earlier, another substantial challenge in deploying small cells is securing the necessary backhaul connectivity (or transport) from the small cell to the core network. There are multiple transport solutions available today, including dedicated fiber, wavelength division multiplexing (WDM), passive optical networks (PON), DOCSIS and microwave radio. The preferred solution is dependent on a variety of factors and considerations, including technical performance, immediacy of deployment, capacity, cost, and accessibility.

Historically, Freedom Mobile has relied primarily on dedicated fiber and microwave for macro cell backhaul. With the introduction of 4G small cells, however, DOCSIS 3.1 is now being successfully used to backhaul traffic from both indoor and outdoor small cells in locations that fall within Shaw's cable footprint.

When deployed in areas with a mid-split 1 GHz HFC plant, these cable modems are provisioned for maximum downlink (DL) and uplink (UL) data rates of 500 Mbps and 100 Mbps, respectively. These

maximum data rates coincide with the peak data rates of the 4G small cells currently being deployed. However, since the peak data rates can only be achieved under ideal conditions, the average the DL/UL data rates during the busy hour are much lower, typically 20-25% of the peak rates. Therefore, up to 5 small cells can usually be served by a single cable modem.

The main advantages of DOCSIS compared with the other alternatives are its low cost, scalability, access to power, and ease of deployment. Another key advantage of DOCSIS is availability. Within Shaw's HFC network, for example, there is typically 3 to 5 times more coaxial cable than fiber in major metros markets. Over time this ratio will certainly decrease but in the short-to-medium term there is a huge incentive leverage DOCSIS for 5G small cell densification.

With its higher peak data rates and lower latencies, however, 5G introduces a new set of demands on the backhaul/ transport facilities, which are discussed in further detail in the sections below.

6.4. RAN Architectures

For 5G NR small cells, the transport requirements depend on several factors including: the total channel bandwidth; the number of MIMO Layers; the maximum supported modulation scheme; and the number of transmit and receive antennas. The division of radio functions between the remote radio sites and centralized locations also plays an important role in the transport requirements. These radio functions include RF signal processing, baseband (or Layer 1) signal processing, and other networking layers (i.e., Layer 2 & 3) in the end-to-end protocol stack.

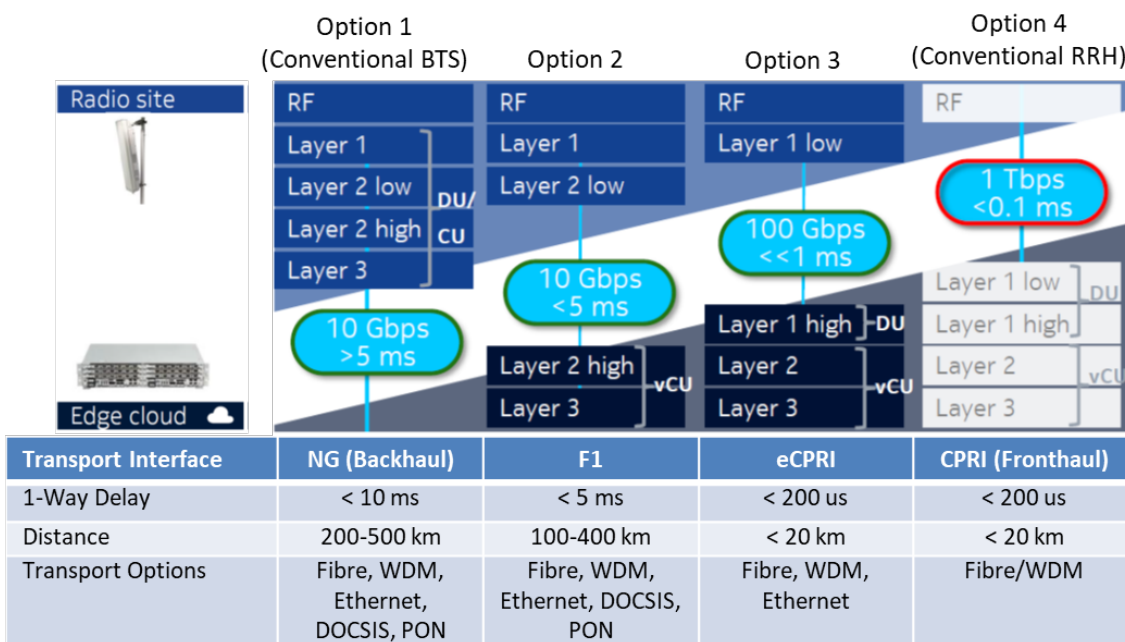


Figure 14 - Possible 5G NR Functional Splits (Source: adapted from Nokia)

As illustrated in Figure 14, there are several possible options for splitting these functions between the radio site and a central location such as a hub site. The upper portion of this figure shows the functions residing at the radio site whereas the lower portion shows the functions at a central site. The throughput and latency requirements are shown over the links between the two.

The first option on the left-hand side (Option 1) represents a conventional base station (BTS) where all functions are implemented at the radio site. The transport facility in the case is referred to as “backhaul”. The advantage of this option is the throughput and latency are minimized because all processing is done at the radio site. In the 5G NR standard, the interface between the BTS and core network is known as the ‘NG’ interface [5]. The backhaul options for the NG interface include dedicated fiber, WDM, DOCSIS and PON.

Option 4 at the other end of the scale represents the conventional remote radio head (RRH) architecture, sometimes referred to a centralized radio access network (C-RAN) or virtualized radio access network (vRAN), if the functions run on virtualized platforms. The RRH performs the RF signal processing only and leaves all other functions to be performed in a baseband unit (BBU) at a central site. With C-RAN, the transport is referred to as “fronthaul” and typically uses the common public radio interface (CPRI) to transport digitized IQ samples of the baseband signal over fiber or WDM from the radio site to the BBU.

The advantages of this approach include simpler radio equipment at the network edge, easier operation, and cheaper maintenance. C-RAN also allows BBU processing capacity to be efficiently reused or shared as demand patterns shift over time. This is particularly relevant for high-density 5G deployments such as stadiums where peak and off-peak demand vary dramatically. The challenge with C-RAN deployments is it requires the RRH and BBU to be connected through a high-speed, low-latency, and accurately synchronized network. Figure 15 shows the required CPRI line rate (without coding) for various channel bandwidths and numbers of transmit/receive antennas. This figure clearly shows that with the higher channel bandwidths and transmit/receive antenna counts associated with 5G, it quickly becomes impractical to use CPRI for 5G fronthaul [6].

With the above in mind, the industry partners responsible for the CPRI specification have developed a new fronthaul specification known as evolved CPRI (eCPRI) [7]. eCPRI offers a ten-fold reduction in the required data rate and allows packet-based transport technologies such as Ethernet to be used. This is Option 3 above.

The final option (Option 2) is based on the F1 interface defined in the 5G NR standard, which assigns the RF, Layer 1 and low-level Layer 2 functions to the radio sites and the other functions in the central site. The advantage of this split is it retains some of the benefits of lower layer splits, while minimizing the throughput and latency requirements.




Figure 15 - CPRI Line Rates

With the continued evolution of DOCSIS standards, we believe that DOCSIS is a viable transport technology for 5G small cells. While the required line rates for eCPRI of up to 100 Gbps are too high for DOCSIS technology, small cells that implement the F1 or NG interfaces could conceivably be carried over DOCSIS. As shown in Figure 16 below, the projected DOCSIS line rates for full duplex DOCSIS (FDX) and extended spectrum DOCSIS, are well within range of those required for the F1 and NG interfaces.




Figure 16 - Evolution of DOCSIS Aggregate Line Rates

Although latency on typical DOCSIS networks currently exceeds the requirements for the F1 interface (i.e., < 5 ms), work by Andreoli-Fang and Chapman [8] demonstrates that latencies of 1-2 ms can be achieved using upstream grant pipelining techniques between LTE/5G and DOCSIS.

6.5. Synchronization

Another important consideration for small cell deployments is synchronization. Unlike Wi-Fi access points, 4G and 5G small cells require precise synchronization to minimize inter-site interference. The minimum frequency and phase requirements for 4G and 5G small cells are listed in Table 1 below.

Table 1 - 4G & 5G Small Cell Synchronization Requirements

Parameter	4G ¹	5G
Frequency	+/- 50 ppb	+/- 50 ppb
Phase	N/R	< 1.5 μ s

1. Assumes FDD mode with no advanced features enabled such as CoMP, eICIC, etc.

In Freedom Mobile's current mobile network, 4G small cells are synchronized via either GPS or precision time protocol (PTP), also known as IEEE 1588v2. GPS/GNSS is currently used for outdoor small cells and PTP is used for indoor small cells. In cases where the outdoor small cell does not have line of sight to a sufficient number of GPS/GNSS satellites, PTP can be used for synchronization as well. Since none of Freedom's 4G small cell operate in TDD mode or use advance features (e.g., coordinated multipoint (CoMP), enhanced inter-cell interference coordination (eICIC)), phase synchronization is not required.

For indoor small cells, PTP is distributed from centrally located grandmaster clocks in each metro area to the small cells over the transport network. This eliminates the need to deploy a local grand master clock on site or install external GPS/GNSS antennas and coaxial cabling for each small cell, both of which would introduce significant deployment costs and time delays.

For indoor sites with fiber backhaul, PTP can be used to reliably distribute both frequency and phase synchronization. For sites with DOCSIS backhaul, however, PTP can only be used to distribute frequency synchronization due to the asymmetric latency between the DOCSIS downlink and uplink and the lack of support for PTP in existing DOCSIS implementations .

During initial small cell deployments there were doubts whether frequency synchronization could be distributed transparently over DOCSIS using PTP. Through extensive testing, however, Shaw has validated that the frequency error meets the 3GPP specification of +/- 50 ppb limit. Figure 17 shows actual frequency error test results from a 4G small cell on a production DOCSIS network. These results show that the frequency is well within specification. To the best of our knowledge, this is the first time that frequency synchronization has been distributed to 4G small cells over a DOCSIS network using PTP.




Figure 17 - LTE Carrier Frequency Error on a production DOCSIS network

For 5G small cells, however, the lack of support for phase synchronization over DOCSIS is a serious shortcoming. Fortunately, the DOCSIS timing protocol (DTP), which is part of the DOCSIS 3.1 specification, and allows phase and time synchronization to be accurately distributed over DOCSIS [9]. Further work is required, however, by DOCSIS chipset, CMTS and cable modem vendors to fully implement DTP before 5G small cell deployments begin, likely in late 2019 or early 2020.

6.6. Converged Access Network

In addition to leveraging DOCSIS for small cell backhaul, MSOs could also take advantage of obvious synergies between 5G small cells and distributed access architectures (DAA) and fiber deeper initiatives. For example, as shown in Figure 18 below, wavelength division multiplexing could be used to aggregate traffic from both DAA nodes and 5G NR macro and small cell sites.

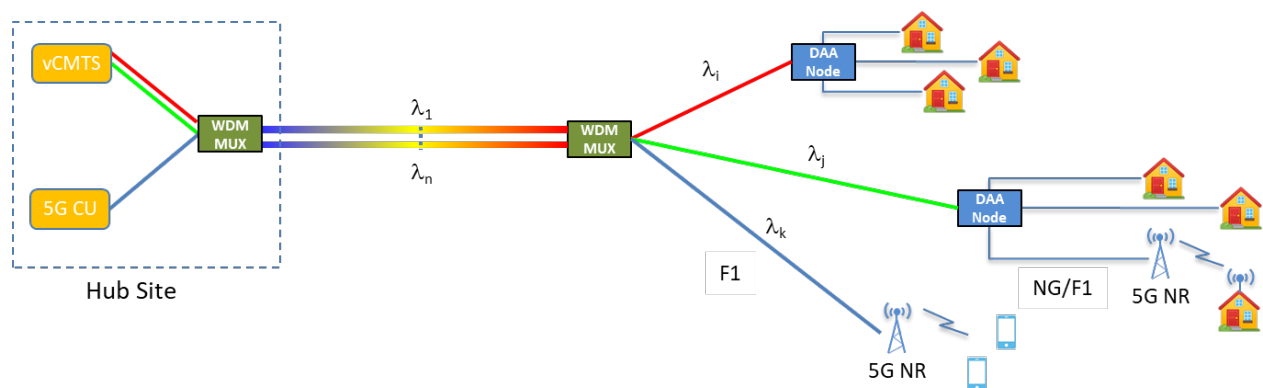


Figure 18 - Converged DAA/5G Access Network

This would allow eCPRI, for example, to be used for high-density 5G NR sites (i.e., both macro and small cells) where peak data rates exceed DOCSIS capacities. Of course, this would not preclude using DOCSIS for backhaul at lower density sites, as required.

Conclusion

5G promises to dramatically transform the role that mobile technology plays in the world. By leveraging their existing and planned network assets, MSOs are well-positioned to capitalize on emerging 5G opportunities, either as mobile network operators or wholesale providers. As wireless networks evolve from 4G/LTE to 5G, small cells will be a key component for delivering the high bandwidth, low-latency connections required by the broad range of use cases 5G makes possible. With hundreds of thousands of Wi-Fi hotspots around the globe and existing support structure agreements, MSOs should have ready access to prime real estate for small cell deployments. This paper also shows that current DOCSIS networks are quite capable of supporting 4G small cells today and evolving to meet the needs 5G small cells in the future. Two key 5G requirements that need to be addressed by DOCSIS chipset and equipment vendors, however, are phase/time synchronization and lower latency. Finally, distributed access architectures (DAA) and fiber deeper initiatives offer the prospect of further 5G/cable network convergence in the future.

Abbreviations

5G	fifth generation
AP	access point
AR	augmented reality
BBU	base band unit
bps	bits per second
CAGR	Cumulative Average Growth Rate
CMTS	cable modem termination system
CN	core network
CoMP	coordinated multi-point
CPRI	common public radio interface
C-RAN	centralized RAN
CU	central unit
DAA	distributed access architecture
DL	downlink
DOCSIS	data over cable service interface specification
DTP	DOCSIS timing protocol
DU	distributed unit
eICIC	enhanced inter-cell interference coordination
eMBB	enhanced mobile broadband
FDX	full duplex DOCSIS
FEC	forward error correction
FWA	fixed wireless access
GB	Giga byte
Gbps	Giga bits per second
HFC	hybrid fiber-coax
HD	high definition

Hz	hertz
ISBE	International Society of Broadband Experts
LAA	license assisted access
MIMO	multi-input, multi-output
mmWave	millimeter wavelength
mTC	mobile type communications
NFV	network function virtualization
NR	new radio
NSA	non-standalone
PON	passive optical network
RAN	radio access network
RF	radio frequency
RRH	remote radio head
SA	standalone
SCTE	Society of Cable Telecommunications Engineers
SDN	software-defined networks
UL	uplink
uRLLC	ultra-reliable low latency communications
VNI	Virtual Networking Index
VR	virtual reality
vRAN	virtual RAN
WDM	wavelength division multiplexing
ZB	Zeta byte

Bibliography & References

- [1] Press Release, GSMA, Mobile Industry to Add \$1 Trillion in Value to North American Economy by 2020, Finds New GSMA Study (Nov. 1, 2016), <http://www.gsma.com/newsroom/press-release/mobile-industry-add-1-trillion-value-north-american-economy-2020-finds-new-gsma-study/>.
- [2] Cisco Visual Networking Index (VNI), 2017.
- [3] ITU-R Recommendation M.2083-0, IMT Vision - Framework and overall objectives of the future development of IMT for 2020 and beyond, September 2015
- [4] Qualcomm, Mobilizing 5G NR Millimeter Wave: Network Coverage Simulation Studies for Global Cities, October 2017.
- [5] 3GPP TR 38.801, “Technical Specification Group Radio Access Network; Study on new radio access technology: Radio access architecture and interfaces”, March 2017.
- [6] ITU-T Technical Report GSTR-TN5G, “Transport network support of IMT-2020/5G”, February 2018
- [7] eCPRI Specification V1.0, "Common Public Radio Interface: eCPRI Interface Specification", August 2017.
- [8] John T. Chapman, Jennifer Andreoli-Fang, “Low Latency Techniques for Mobile Backhaul over DOCSIS,” Proc. of SCTE Fall Technical Forum, October 2017, Denver.
- [9] Jennifer Andreoli-Fang, John T. Chapman, “Mobile Backhaul Synchronization Architecture Proc. of SCTE Fall Technical Forum, October 2017, Denver.

A Customer Experience Based Approach to Improving Access Network Power Reliability

An Operational Practice prepared for SCTE•ISBE by

Tobias Peck

Director of Software and Fiber Product Management
Alpha Technologies
3767 Alpha Way, Bellingham WA
360.392.2247
tpeck@alpha.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. Why Does Power Reliability Matter?	3
2. Evolution of HFC Power Reliability	4
3. Need for a Better Process	4
4. Six Sigma/Customer Focus.....	5
5. Analysis of the Truck-Rolled Node Outage Metric	6
6. Analysis of Stanby Powering Risk in High Outage Nodes	7
6.1. Risk Analysis	7
6.2. Control Comparison	9
6.3. Field Verification	10
7. Measure It to Manage It	11
8. Developing a Reliability Improvement Program	11
8.1. Key Program Challenges	12
8.2. Data Challenges	12
9. Results	13
10. Building on Success	13
Conclusion.....	14
Abbreviations	14
Bibliography & References.....	15

List of Figures

Title	Page Number
Figure 1 – Overlapping Risks in a Truck-Rolled Node Outage (TRNO)	6
Figure 2 – Power Analysis of High Outage Nodes	8
Figure 3 – Stanby Power Risk of High Outage Nodes.....	9
Figure 4 – High Outage Nodes vs Rest of Network	10
Figure 5 – Examples of Risk Factors Seen During Field Verification	10

Introduction

Customers drive the Broadband industry's revenue. Customer experience is a major factor in determining MSO success, however, traditional Outside Plant (OSP) power reliability improvement strategies are based around arbitrary maintenance cycles, and not targeted toward addressing network elements that most effect customer experience. These strategies have an inconsistent effect on overall plant power reliability and fail to take advantage of modern big data analytics to maximize customer impact.

In 2018 a major US MSO set out to revolutionize reliability improvement by flipping this paradigm and beginning the process with a customer-driven metric that truly shows where customers are being impacted. Their Data Driven approach began by collecting and analyzing customer impacting node outage statistics and ended with a goal of significantly reducing customer affecting node outages in 2018. This paper will detail the process of developing a customer focused Reliability Improvement Strategy and show the real-world results of this strategy. This study will demonstrate how the major MSO under study, using the tenants described to focus efforts on the issues that directly impact customers, will maximize spend on preventative and on-demand maintenance in order to greatly improve power reliability and customer experience.

Content

1. Why Does Power Reliability Matter?

MSOs have consistently driven the broadband industry forward. Billions of dollars have been spent to create a network that can deliver data and video at speeds that continue to push the envelope of what was once thought impossible. New technologies that leverage the power of the HFC network are being introduced with increasing frequency. VOD platforms provide customers with a simple, yet engaging user interface that makes the actual experience of searching for new content to binge-watch fun. MSOs have been successful at building an attractive product and while the robust HFC network is designed to consistently deliver this product to customers, key components of the network depend upon reliable power to make that happen.

In the United States this need for reliable power poses a huge risk. Contrary to what many believe, the grid in the US is significantly less reliable than other industrialized nations.

“More than 70% of the grid’s transmission lines and transformers are 25 years old; add nine years to that and you have the average age of an American power plant. According to industry expert Peter Asmus, we rely on twice as many power plants as we actually need because of “the massive inefficiencies built into this system.” As a result, significant power outages are climbing year by year, from 15 in 2001 to 78 in 2007 to 307 in 2011. America has the highest number of outage minutes of any developed nation – coming in at about six hours per year, not including blackouts caused by extreme weather or “acts of God,” of which there were 679 between 2003 and 2012. Compare this with Korea at 16 outages minutes a year, Italy at 51 minutes, Germany at 15, and Japan at 11. Not only do we have more outages than most other industrial countries, but ours are getting longer. The average U.S. power outage is 120 minutes and growing, while in the rest of the industrialized world it’s less than ten minutes and shrinking. According to Massoud Amin, a power systems engineer, on “any given day in the U.S. about half a million people are without power for two or more hours.”

- Gretchen Bakke, PHD., Bloomsbury Publishing USA, Jul 26, 2016

The unreliable US grid is only a part of the story. The utility grid and the HFC network are not aligned, which allows the possibility for the unreliable grid to cause outages for customers who are unaware of any utility outage. If a power supply injecting power to the plant is on a different powering leg of the grid than the customer premise, power to the plant can go down without affecting power that runs customer devices in the home, causing customer frustration and trouble calls. Additionally, an increasing number of US customers, particularly business services customers, are backing up premise power to ensure connectivity.

2. Evolution of HFC Power Reliability

As a result, MSOs have spent significant time building the most reliable backed-up powering networks possible. Since the advent of standby power for OSP, operators have spent immense time and capital dollars ensuring that their networks can withstand the vast majority of nuisance outages caused by the unreliable US utility grid. The HFC network has grown from a simple video delivery network to a secondary utility network capable of powering and backhauling a myriad of devices.

Plant engineering standards have been created for the OSP that ensure battery backup time available during an outage is well above the US average outage time. Many US MSO power supplies are capable of being remotely status monitored and tested. Many MSOs have developed methodologies for using remote battery testing results to prioritize maintenance and plant upgrades. Remote status monitoring has provided immense amounts of data on plant inventory and risk factors to help prevent customer impacting plant outages. When properly maintained and monitored the powered coax network of most US MSOs is significantly more robust than that of the utility grid.

Due in part to the increased power-reliability possible from the powered coax MSO network, a new growing focus on business-to-business services is driving a strong interest in this reliable HFC powering network, as well as the backhaul and real-estate that it also provides. Today, B2B activities like Small Cells, WiFi, IoT, Security & Surveillance (SWISS) are taking advantage of the HFC network in a whole new fashion. Over the past five years, a number of operators have been adding WiFi access points to their networks to create a stickier environment which allows their customers to always utilize their network. There has also been a new demand for IoT networks which allow machine to machine connections to happen across large areas with single access point coverage. Connections of all kinds will be required to create a ubiquitous network across a geography and right now the HFC network is in a great position to take advantage of these new service requirements.

3. Need for a Better Process

This robust design requires significant time and effort to effectively maintain. The Outside Plant is so named because it is, in fact, outside. This fact makes maintaining the powering elements of the network – power supplies, batteries and outdoor hardened transponders and cable modems – an operational nightmare. Weather extremes, dust particulate, and an array of small animals offer constant assault on these powering elements that necessitate constant and consistent maintenance. For example, in the last year alone, the US network experienced temperature extremes from -42 F in the Northeast to 122F in the Southwest, multiple hurricanes in Texas and Florida and a parade of blizzards in the Northeast.

Additionally, it covers an immense geography which makes it expensive and time consuming to maintain. The vast US HFC network consists of enough miles of plant to stretch to the moon and back again twice.

With more than a million miles of HFC plant to maintain and more than 650,000 power supplies the amount of time required to do yearly maintenance on every power supply is quickly becoming prohibitive to operations budgets set aside for maintaining every inch of this network.

As DOCSIS advances have expanded the capability of the HFC network and the options for new revenue generating services off of the network have grown, the need to ensure network power reliability is more crucial than ever. In spite of this fact, the geography and uncontrolled environment of the OSP network continue to present a huge risk that needs to be managed. Without a more focused method to improve and maintain HFC network power reliability, MSOs will struggle to meet customer expectations for reliability.

4. Six Sigma/Customer Focus

The need for a more focused process is exactly what led the MSO studied by this paper to a huge step forward in power reliability assurance. The search for a more focused methodology led to a simple yet brilliant answer which ties in several key concepts of one of the great operational improvement philosophies, Six Sigma.

Volumes have been written on Six Sigma, but in short it is a disciplined, statistically-based, data-driven approach and continuous improvement methodology for eliminating defects in a product, process or service.¹ Developed by Motorola in the 1980's and then adopted and honed by General Electric in the early 1990's, the original focus of the philosophy was on manufacturing. The principles have since been used by hundreds of organizations to improve a variety of operational processes. The term Six Sigma is derived from the goal of achieving six standard deviations (σ or Sigma) between the mean and the nearest specification limit with regard to variation in process outcomes.

The studied MSO used two key concepts from Six Sigma to focus their power reliability improvement efforts. First, every decision was driven by data. For every key decision point all available data was analyzed to provide direction for process improvements. The analysis of that data, decisions made from it, and new processes driven by it will be given significant focus later in this paper.

Second, and perhaps more importantly, data analysis and process improvements were all focused on customers. In the defining book on Six Sigma, former Motorola employees Mikel Harry and Richard Schroeder state:

“The heart of Six Sigma lies in improving products and services that will benefit the customer. Companies need to understand how their customers define [value] and to create products and services that meet their expectations. Six Sigma translates issues critical to customers' satisfaction to what is critical to a product's or service's [value.]”²

Focusing improvement efforts on creating value for customers is a key tenant of Six Sigma, but the reason for this goes beyond just the obvious improvement in customer satisfaction. Many organizations, and specifically many MSOs, spend significant time troubleshooting and fixing customer issues after they've occurred. This reactive response is inefficient and is done to the detriment of any potential efforts to proactively improve services or products. In short, so much time is spent putting out fires that no time is available to prevent them from happening in the first place. By focusing on fixing the root cause of the customer-based issues that are costing the most, the resulting savings will be maximized and will often fund future continual improvement efforts.

By using these guiding principles of being data-driven and customer focused, the studied MSO identified an appropriate metric to help baseline and drive power reliability improvement - truck-rolled node outages (TRNO.) For the studied MSO, several metrics are used to track the reliability of utility power, node health and customer modems in their network, all of which are extremely useful to help ensure network reliability. However, a TRNO incorporates the element of true customer impact, by including only node outages that lead to customer calls and a truck roll to address customer issues, making this the perfect metric to use for this process.

5. Analysis of the Truck-Rolled Node Outage Metric

After identifying a customer focused metric to drive the improvement process, the studied MSO analyzed data and made some key discoveries. For the analysis, a year worth of data on truck-rolled outages was used from August of 2016 to August of 2017. It was discovered that nodes with an average of at least one truck-rolled outage per month anywhere within the node boundary made up approximately 3% of the total nodes, but accounted for 17% of total TRNOs. This key discovery led to the next question – why? What did these nodes have in common that created such a high rate of this type of outage? This paper will delve into deeper analysis later, but the simple answer is risk. These nodes displayed a higher rate of key powering risk factors that can lead to outages.

A power related TRNO is essentially an intersection of multiple risk conditions existing simultaneously. First, a loss of utility power creates the possibility for loss of plant powering. Second, a failure of standby power drops power to the network shutting off service to the customer premise. And, lastly, the loss of service must impact the customer significantly enough to pick up the phone and call for service. The studied MSO also analyzed data pertaining to non-powering related node issues, but found these issues to be much less significant. For this paper, focus will be given to power related node outages. By identifying and mitigating risk around these conditions, TRNOs can be significantly reduced.




Figure 1 – Overlapping Risks in a Truck-Rolled Node Outage (TRNO)

As discussed earlier in this paper, the utility grid in the US is an inherent risk to plant powering. However, understanding which power supply sites have the highest risk or least stable utility power allows MSOs to bolster standby runtime capacity at those specific sites with additional batteries or other more significant back up strategies, if necessary. Having site specific knowledge of utility stability can also allow MSOs to address potential grid stability issues with local utility companies and hopefully see improvement through that process.

While it is the most difficult risk to control, there are some potential ways to mitigate customer risk. A key fact about customer risk is that any customer who has called about a service related issue in the past is significantly more likely to call again in the future than a customer who has never had an issue. Conversely, customers who have never called are more likely to ignore short-lived issues. Therefore, mitigation of other risk factors will inherently help mitigate customer risk. Additionally, it is possible to analyze data on customer services and data usage to determine which customers rely more heavily on the HFC network and thus would be more likely to notice an outage and submit a trouble call. Similarly, it is theoretically possible to analyze the layout of the HFC network versus the grid and better understand which customers' home power and plant power are on different sections of the grid. With both customer reliance risk and grid misalignment risk, MSOs could target upgrades and enhancements to plant powering equipment to help mitigate these risks.

Of the risk elements that go into a truck-rolled node outage, standby powering risk is the easiest to control. Unlike the other risk elements, plant standby powering equipment, once permitted and provisioned, is completely owned and controlled by the individual MSO. Most US MSOs have some manner of remote status monitoring that can provide an immense amount of information on where standby powering risk exists in the HFC network. Once the studied MSO had identified the 3% of their high-outage nodes, they commissioned Alpha Technologies to analyze standby powering risk at these sites and provide data to help determine how to solve the problem.

6. Analysis of Standby Powering Risk in High Outage Nodes

6.1. Risk Analysis

To better understand standby powering risk in the studied MSO's HFC network, historical status monitoring data was analyzed alongside the MSO's TRNO data. The analyzed data showed the following:




Figure 2 – Power Analysis of High Outage Nodes

The significance of the nodes that were not in status monitoring will be discussed in more detail later, but at this point we will dive deeper into the powering risks discussed above. Firstly, utility risk has already been identified as a key factor in power related TRNOs but, it is significant to point out that 35% of these high risk nodes have at least one power supply showing a utility risk of either extreme high outage time, – in the 98th percentile of all monitored power supplies - extreme outage frequency, – at least one outage month, 10X the national average – or both.

Most significantly, 73% of the analyzed nodes had at least one power supply with an identified powering risk factor. So what are these risk factors and why are they significant. Below is a breakdown of the power supplies that show a powering risk factor.




Figure 3 – Standby Power Risk of High Outage Nodes

These risk factors are designed to identify standby devices that will not be able to perform their most important function, providing backup power to the HFC network in the case of a utility outage. Current Failed Self-Tests and Other Current Major Power Supply Alarms are essentially identifying power supplies that will not provide even minimal back-up time. Offline power supplies will be discussed in more detail later, but are at best unknown and at worst already impacting customers. And lastly, due to the historical nature of the TRNO data analyzed, it was necessary to look at historical occurrences of Alarming and offline conditions to identify devices that most likely caused TRNOs, but have since been remedied.

6.2. Control Comparison

While this data seems to be significant, understanding how these high risk nodes compare to the rest of the network provides a meaningful control data comparison. If there is no significant difference in risk factor frequency between high outage nodes and the rest of the nodes in the network, these risk indicators may not be a key factor in causing TRNOs. The chart below shows a comparison between key powering risk factors for the high outage nodes and the rest of the network.

High Outage Nodes Vs Rest of Network

- Poor Utility Power Supplies – 1.8x
- PS with Failed Self-Test – 2.6x
- PS with Other Major Alarms – 2.7x
- Offline Power Supplies – 2.2x
- Aging PS (> 15 years) – 1.6x




Figure 4 – High Outage Nodes vs Rest of Network

Clearly the comparison data shows a significant correlation between Power Supply related issues and TRNOs as high outage nodes have a significantly higher occurrence of all key risk factors identified. Additionally, high outage nodes had a 60% higher likelihood of having at least one aging power supply of greater than 15 years. As these devices age, they present several additional powering risks to the outside plant including: poor output voltage regulation, lowered efficiency, reduced backup runtime, potential failure of inverter electronics and reduced monitoring capability. These are additional risks that can easily lead directly to TRNOs.

6.3. Field Verification

The data analyzed clearly demonstrated that high outage nodes tended to have a greater occurrence of powering risk. In order to verify this data, a small sample of these high outage node sites were identified and field verification was performed by an Alpha field engineer and an experienced employee of the studied MSO. Examples of issues found included significantly aged batteries (2x useful life,) sites where batteries had been stolen or were never installed, improper grounding or line powering, or the absence of a tap or drop to allow remote monitoring of the power supply. The field findings definitively confirmed the data.



Figure 5 – Examples of Risk Factors Seen During Field Verification

7. Measure It to Manage It

The data analysis and field verification show a clear correlation between standby power risk factors and high outage nodes. Now that the data had shown the studied MSO a clear path to improve power reliability and reduce TRNOs, the next step was to create a process to address these standby power risks in high outage nodes. The first step of this process was to tackle the problem of *blind* power supplies. For the purposes of this paper, a blind power supply is any power supply that is not being actively monitored in the customers power Element Monitoring System (EMS) because it has either gone offline and lost connection to the EMS or was never provisioned to be status monitored. Any node with at least one of these power supplies would essentially have a blind spot that could lead to further outages and would also make baselining an improvement process and measuring success almost impossible. The study MSO had at least one blind power supply in more than one third (37%) of the high outage nodes studied.

Why is monitoring power supplies so crucial and how can it prevent future power related TRNOs? Many MSOs have inherent monitoring and management of customer modems that allows for an understanding of plant RF issues and customer outages. Power supply status monitoring gives MSOs several ways to inform tech ops teams of potential issues before they become customer impacting and are seen by systems monitoring customer modems.

First, a power supply EMS has the ability to inform customers of critical issues with a power supply's ability to provide adequate backup during a utility outage, such as bad or missing batteries or a failed inverter. Without the ability to see and address these issues, even momentary outages in utility power could become customer affecting. Next, status monitoring provides visibility of stand-by powering events caused by utility outages which can then be managed by tech ops to ensure that portable backup generators can be deployed successfully to keep the plant up and running. Again, without this visibility, perfectly healthy power supplies with adequate backup time could potentially run for hours without being addressed and lead to TRNOs. Additionally, an EMS can provide visibility to power supplies that have gone offline as a result of an extended outage and provide answers on a proper fix agent to address the lynch pin causing a large customer outage. And finally, status monitoring can provide key inventory data for proactive reliability planning and budgeting. Being able to monitor powering issues is unmistakably a key tenant of any process designed to improve power reliability and was a priority in the program developed and implemented by the study MSO.

8. Developing a Reliability Improvement Program

The customer-focused, data-driven approach had helped to identify key risk factors to be addressed at the study MSO and the decision had been made to focus on ensuring all power supplies were able to be monitored in the EMS. The next step was to develop aggressive but realistic improvement goals around TRNOs and create an action plan to get there. Our sample MSO started with a goal of reducing TRNOs and customer services calls by 15 to 20%.

To maximize the program's impact on plant reliability, high data on high outage nodes was updated starting at the beginning of 2018 and high outage nodes were redefined as any node averaging at least 1 TRNO per month for 3 months. For the study MSO this definition identified 20% of the MSOs nodes which accounted for 55% of TRNOs. This is significant as it shows how this process allowed the MSO to focus on a manageable subset of identified high risk sites and still potentially impact more than half of TRNOs. The focus of this study is on efforts to improve power reliability using the prescribed methodology, however, additional focus was given to addressing a small subset of non-powering risk factors as well.

Now that the focused subset of power supplies had been identified, standards and processes could be created to ensure standby powering risk factors were eliminated at each site. All identified sites would receive a Preventive Maintenance (PM) visit following the SCTE recommended practice (SCTE 205 2014) to ensure that all risk factors were addressed. By using this industry standard recommended practice, the study MSO could be sure that all risk factors, even those that may not be easily seen by status monitoring, were being addressed and sites would be left in the most reliable state possible.

In addition to a standard PM each site would have steps taken to address identified risk factors.

1. All necessary steps will be taken to ensure proper EMS monitoring, including replacement or addition of transponders and installation of a tap and/or drop where necessary.
2. All sites are required to meet a 3-hour estimated back up runtime standard based on plant power loading at that site.
3. Aging power supplies will be replaced in lieu of repair to address risk factors.
4. Accurate collection of field inventory data is a point of emphasis for technicians to enhance future reliability planning efforts.

In addition to performing maintenance and upgrades to identified sites using the guiding principles above, geographical efficiencies will be leveraged where possible to address additional sites that are currently not being actively monitored or sites that display standby power risk factors. These simple principles will guide the process to drive improved power reliability and ultimately reduce TRNOs.

8.1. Key Program Challenges

While the guiding principles of the program were relatively straightforward, several key challenges needed to be addressed in order to ensure success of the program implementation. Many of these challenges are experienced by US MSOs who deal with trying to implement similar power reliability improvement plans and are the result of one key fact – MSOs are Multi-System Operators that are originally made up of multiple independent systems with disparate standards and geographies.

Due to this fact, a key challenge for the studied MSO was the efficient identification of missing nodes and power supplies from original plant maps. Until these nodes and all power supplies within them could be identified, it would be difficult to plan and budget for upgrades to bring these devices into a monitored state. The studied MSO addressed this key issue quickly and decisively by using data from all available tools to create a list of missing sites. From this list EMS and CMTS data could be used to identify a subset of devices that were online but not monitored and bring those devices into a monitored state without a truck roll. The remaining sites were then disseminated to system ops teams to use all tools at their disposal to provide location and inventory data on missing sites, with technician truck-rolls as an absolute last resort. Once this site data was gathered at the system level, data was compiled from all systems and more precise planning and budgeting could take place.

Several other challenges resulting from the compilation of multiple independent systems are the regional disparity of hardware standards, EMS alarming and ticketing methodologies, and SOWs for site upgrades and maintenance. The study MSO created small working groups to document regional differences in each of these and used available data to successfully address the disparities and create corporate standards.

8.2. Data Challenges

Another challenge that often arises when driving progress through an identified metric is the emergence of data outliers. Most metrics have parameters set around them based on assumptions and desired

outcomes. For example, the TRNO metric used by the study MSO does not include widespread storm or utility events for good reason, as these are generally due to large weather events and not operational or plant issues. During these large events customer expectations and behaviors tend to change with the realization that many services will go down, so considering them a part of the TRNO metric does not make sense. However, the numeric definition of what constitutes a storm event can impact the metric drastically.

The logic involved in this storm case and other data outliers were reviewed and adjusted in several cases during the program to ensure that ongoing measurement of TRNOs would be as accurate as possible. While this was the best course of action in the long term, it did create some difficulty in measuring progress against the original baseline. Additionally, the studied MSO is in the process of significant plant architecture changes to improve plant capability which have added some additional complexity to the measurement progress. Neither of these issues is insurmountable, but this type of shifting data should be accounted for as often as possible prior to creating a power reliability improvement program.

Finally, with technicians visiting a large number of sites with no EMS data, the studied MSO had a great opportunity to gather previously unrecorded plant inventory data for future planning. Although this seems relatively straightforward, the key obstacle to be overcome is that the field personnel collecting the data do not generally value it as much as executives who need it to make informed decisions. If the process to collect key information and quickly compile it to be analyzed is not user-friendly for technicians the data collected may not be correct. Additionally, if there is no process set in place to ensure that the data is consistently maintained, it will quickly become useless. The studied MSO mitigated both of these problems by identifying and implementing software tools to make data collection and maintenance by field and NOC technicians as painless as possible.

9. Results

While results of this program are ongoing and will be updated continually, the studied MSO has thus far seen very positive results through the customer-focused, data-driven reliability enhancement program. To date the study MSO has accomplished:

- 9,990 power supplies brought into a monitored state in the EMS
- 9,098 system battery runtime upgrades

The resulting impact on key metrics has been*:

- 19% reduction in trouble calls
- 14% reduction in TRNOs

Reduction of 45,000 maintenance hours in identified nodes**These results are tentative and must be updated before presenting*

10. Building on Success

In order to ensure positive gains achieved by this program and others similar to it are maintained, next steps and a long-term strategy must be considered. Several key next steps should be taken after the initial year of the program. In order to maximize effectiveness in the next program year it will be necessary to do a program review and analyze the first year results. Were initial goals achieved? What were key lessons learned that should be incorporated into improving the program for the next budget year? Data for TRNOs should be reevaluated and target nodes should be redefined for year two. Program spend versus ROI and NPV should be analyzed and adjusted based on new program goals. While it is impossible to completely eliminate TRNOs, the positive results shown in year one should theoretically make year two less daunting.

In addition to analyzing the program and adjusting for a second year deployment, several other parallel initiatives should be considered by the studied MSO. First, nodes with frequent outages caused by plant powering issues should be identified and addressed sooner to alleviate customer frustration from calling multiple times on the same issue. A more effective real-time feedback loop between software tools designed to monitor powering issues and those designed to initiate truck-rolls on customer calls would allow technicians to more easily triage and repair power related node issues at the time of outage and avoid repeat customer calls. Next, continuing to improve processes for collection of site inventory data will allow for a more precise and predictable budgeting process built around useful asset life versus documented age, thereby ensuring adequate inventory levels to deploy for on-demand outage fixes. Lastly, using the improved inventory data in combination with the recently improved monitoring visibility, enhancements could be made to the EMS to better predict system health and backup time. All of these improvements could be used to proactively eliminate additional TRNOs.

Conclusion

The need to continually improve the reliability of the HFC network and meet the growing demands of customers as well as prepare for the opportunity of new revenue generating services, coupled with the financial, geographic and environmental challenges present with the OSP network, demands a more focused method for driving increased power reliability. In order to increase effectiveness toward this goal, the Major MSO studied in this paper was able to achieve significant impacts on customer experience and operational costs by centering their continual reliability improvement process around a customer-focused, data-driven core.

Beginning with the analysis of TRNOs, a key customer focused metric, the studied MSO was able to build a program of hyper-efficient power reliability improvements. Spearheading this program was a push to ensure all powering devices were visible through their powering EMS and ensure reliability improvements could be effectively measured. In order to ensure the best results possible with regard to field execution of this program, SCTE recommended practices were leveraged. As challenges arose around numerical outliers and corporate standardization, data was continually used as a pathway to the optimal solution. And, now that the first year of the study MSO's reliability improvement program is winding down as a resounding success, a blueprint has been created for other MSOs to use a similar customer-focused, data-driven approach to achieve the access network power reliability needs of tomorrow's HFC network.

Abbreviations

B2B	Business to Business
CMTS	Cable Modem Termination System
DOCSIS	Data Over Cable Service Interface Specification
EMS	Element Management System [power focused]
HFC	Hybrid Fiber Coax
INV	Inverter
MSO	Multi-System Operator
NOC	Network Operating Center
NPV	Net Present Value
OSP	Outside Plant
PM	Preventive Maintenance [Visit]

PS	Power Supply
ROI	Return on Investment
SCTE	Society of Cable Telecommunications Engineers
SOW	Scope of Work [for a project]
TRNO	Truck-Rolled Node Outages
VOD	Video On-Demand

Bibliography & References

ANSI C63.5-2006: *American National Standard Electromagnetic Compatibility–Radiated Emission Measurements in Electromagnetic Interference (EMI) Control–Calibration of Antennas (9 kHz to 40 GHz)*; Institute of Electrical and Electronics Engineers

The ARRL Antenna Book, 20th Ed.; American Radio Relay League

Code of Federal Regulations, Title 47, Part 76

Reflections: Transmission Lines and Antennas, M. Walter Maxwell; American Radio Relay League

A PNM System Using Artificial Intelligence, HFC Network Impairment, Atmospheric and Weather Data to Predict HFC Network Degradation and Avert Customer Impact

A Technical Paper Prepared For SCTE/ISBE By

Larry Wolcott

Comcast Fellow, Next Generation Operations Technology
Comcast

1401 Wynkoop Suite 300, Denver, CO 80202
(303) 726-1596

Larry_Wolcott@cable.comcast.com

Michael O'Dell

Director, Network Maintenance
Next Generation Access Networks
Comcast

215 East North St, New Castle, PA
(724) 856-3074

Michael_ODell@cable.comcast.com

Peter Kuykendall

Principal Engineer
Video Infrastructure and Security Engineering (VISE)
Comcast

4100 E. Dry Creek Road, Centennial, CO 80122
(720) 663-7581

Peter_Kuykendall@cable.comcast.com

Vishnu Gopal

Senior Engineer
Software Dev & Engineering, Engineering & Operations
Comcast

183 Inverness Drive West, Englewood, CO 80112
(303) 658-7049

Vishnu_Gopal@cable.comcast.com

Jason Woodrich

Senior Engineer

Software Dev & Engineering, System Engineering
Comcast

183 Inverness Drive West, Englewood, CO 80112
(303) 658-7123

Jason.Woodrich@cable.comcast.com

Nick Pinckernell

Senior Principal Researcher

Technical Research and Development, Product & AI
Comcast

183 Inverness Drive West, Englewood, CO 80112
(303) 658-7305

Nicholas.Pinckernell@cable.comcast.com

Table of Contents

Title	Page Number
Introduction _____	6
Acknowledgements _____	6
History and Background of Environmental Influence on HFC _____	7
Evolution of HFC Network Architecture _____	8
Traditional Network Maintenance Practices _____	10
Advancements in PNM Technology _____	11
Influence of PNM on Network Repair Prioritization _____	12
A Glimpse into the Future _____	12
Weather _____	13
Effects on Compromised Plant _____	13
Effects on Workforce _____	13
Weather Data _____	14
Weather data requirements _____	14
Geographic scope _____	14
Task 1 – correlation establishment (pre-trial) _____	14
Initial assumptions _____	15
Task 2 – Ticket Reprioritization (trial phase) _____	16
Selection of weather data vendor _____	16
Collection and storage of weather data _____	16
Validation of weather forecast data accuracy _____	16
Formatting and integration of weather data _____	18
Actual Weather Data Example: _____	19
Forecasted Weather Data Example: _____	21
Integration of weather data and plant data _____	23
Analysis of Weather History vs. Plant History _____	23
Characterization of Plant Assets _____	25
Artificial Intelligence, Machine Learning, and the Future of PNM _____	25
The History of Machine Learning And Artificial Intelligence in Comcast Operations _____	25
The Commoditization of ML / AI _____	26
The Convergence of PNM and ML / AI _____	26
The Composition of an AI Enhanced PNM Program _____	27
Hardware & Software _____	27
Data and Modeling _____	27
Execution _____	28
Addressing Common Barriers to Adoption _____	28
Technology and Operational Discontinuity _____	29
New Prioritization Around 10 Day Forecasts _____	29
Solving Priority Conflicts _____	29
Existing Priority System _____	29
Proposed Priority System _____	31
Adapting This Into Existing Workforce Scheduling _____	31
Operational Culture and Workforce Training Considerations _____	32

Organizational Alignment and Common Understanding _____	32
Conclusion_____	32
Abbreviations _____	33
Bibliography & References _____	34

List of Figures

Title	Page Number
Figure 1 – HFC Amplifier Cascade	8
Figure 2 – HFC Logical Digaram	10
Figure 3 – Typial RF Spectrum	13
Figure 4 – Impaired RF Spectrum, Temperature Induced	13
Figure 5 – Performance Evaluation	17
Figure 6 – Forcast Accuracy	18
Figure 7 – Data Process Model	18
Figure 8 – Temperature Over Time	23
Figure 9 – ML / AI Flow Diagram	24
Figure 10 – Analysis Flow Diagram	24

Introduction

Proactive network maintenance (PNM) has become a cornerstone technology within the Data Over Cable Service Interface Specification (DOCSIS®), providing tremendous benefit to cable operators and their customers. From adaptive equalization to full band capture, a rich and extensive data model exists to proactively maintain our valuable networks and reduce operational costs. However, due to the complexity, financial and cultural barriers, many operators have been unable to gain traction with implementing such systems. This paper will examine PNM capabilities, weather information, artificial intelligence, machine learning, operational practice and financial implications to provide a meaningful approach for implementation.

Physical networks that are exposed to the environment are subject to environmental influences which can affect their performance, especially, but not limited to, coaxial and HFC networks and also including fiber optic, satellite, and wireless networking solutions. This is due in part to physics and the simple fact that physical things respond to environmental conditions in predictable ways. Some of the most common factors that influence these networks are atmospheric conditions such as heat, cold, wind, rain, snow, humidity, and freezing. The core premise of this work is to correlate the predictable nature of weather and seemingly unpredictability of weather related outages. While it is generally accepted that these things affect the network, very little work has been done to quantify in an empirical way. Thus, having predictable correlation with faulty network elements allows for operators to proactively prioritize and repair the problems before they impact customers.

Further, by understanding the environmental influence throughout the life span of a network, better decisions can be made about design and construction. This allows operators to evaluate the true cost of ownership over the entire life of the network including support and maintenance, which are generally very expensive.

Acknowledgements

This work would not have been possible without a significant contribution of the individuals listed as authors and contributors. This unique collaboration was facilitated by the Comcast CLEAR engineering program, of which we are all grateful. Thank you to Andrew Frederick, Denice Loud, mentors, and the CLEAR committee members for creating this opportunity.

The Society of Cable Telecommunications Engineers (SCTE/ISBE) has also been instrumental in the ongoing support and promotion of PNM. Especially Chris Bastian, Dean Stoneback, the network operations subcommittee for proactive network maintenance and its members.

Finally, but not least important, special thanks to CableLabs, Alberto Campos, Tom Williams, Jason Rupe, Robert Cruickshank and Curtis Knittle for the steadfast industrial support of PNM. Without their support and innovation, PNM would not exist. Also, Justin Menard and Wil Colon (retired) from Comcast for their previous work on the impact of temperature to the customer experience.

History and Background of Environmental Influence on HFC

Weather has always informed CATV plant maintenance activities. From the earliest days of antenna systems, to more modern hybrid fiber coax (HFC) architectures, atmospheric conditions have always had an influence on the performance and maintenance of CATV networks. Heat, cold, wind, rain, snow and ice. Wherever there are cable network components exposed to the open air and elements, operator's behaviors, and network performance will be influenced by some aspect of weather. All telecommunication and utility teams have developed maintenance cultures and "tribal knowledge" that have been influenced by localized weather patterns. These cultures are ingrained in the practices, habits, and annual maintenance timelines as defined by the historical weather cycles.

Looking at the earliest CATV systems like the ones built by Ed Parson in Astoria Oregon, or Bob Tarlton in Lansford Pennsylvania, their antenna structures were necessarily placed in elevated positions to get the best possible line of site to broadcast antenna locations. Those elevated positions can be subject to winds, moisture, and temperature changes requiring maintenance to keep them properly aligned, and the RF connectors tight and free from corrosion. The maintenance of the transmission lines from those elevated positions, connecting the receive antennas to the community served, are also subject to the influences of weather. Coaxial or other cables ran on aerial pole structures can be influenced by wind, temperature changes, tree limbs, snow and ice loading, and other weather influenced circumstances. Fast forward to newer networks, which carry a broader array of services and signal types, and while they are constructed with more modern materials and components, they are still significantly influenced by localized weather.

Satellite dishes, microwave links, long amplifier cascades, fiber and coaxial cables, and distribution passive components: Though the components may have changed, the behaviors remain the same. Northern climates or higher altitudes that are subject to heavy wet snow in the winter have adapted their maintenance cultures to deal with the accumulation on aerial cable spans and satellite dishes. Shaking spans and sweeping dishes in the winter have become common practice born through necessity over time. Nor'easters, in the New England states, lake effect snows from the Great Lakes in the upper Midwest, and wildfires in the dry seasons are all localized weather phenomena that have informed maintenance behaviors. High Plains states can see temperature changes exceeding 30° Fahrenheit from daytime highs, to nighttime lows. Mountainous areas have developed procedures to deal with flash flooding, when heavy rains run off the hills and pool in low lying areas. These areas can experience mudslides/hill slides when there is insufficient foliage cover to hold topsoils in place during the runoff. Other climates, such as Florida, have developed their own cultures based on their own localized weather influences. The peninsula of Florida experiences large volumes of lightning strikes, locally heavy thunderstorms, and tropical storms such as hurricanes on an annual basis. Plains states experience annual tornadoes and high wind events that can be devastating to utility and telecommunications infrastructures. Across the vast and varied landscapes of America, maintenance behaviors have always been informed by weather.

Evolution of HFC Network Architecture

The earliest systems, like those previously referenced in Astoria, Oregon and Lansford, Pennsylvania, were designed and constructed to retransmit or extend the reach of over-the-air television signals to communities that didn't have a direct line of site to the television broadcaster in the community or adjacent community. Whether because of geographic terrestrial land formations, buildings, or simply distance, the available over-the-air broadcast signals were not sufficiently strong to provide for good reception in the community of interest. From the 1940s through the 1960s, these types of (re)transmission systems were built out to serve the growing video entertainment appetite of the country.

The evolution of cable television systems infrastructure naturally followed the public right of ways, and previously established placement and attachment routes developed by the telephone and electric companies. Telephone companies had already been deploying services for nearly 50 years before CATV systems were deployed. The first transatlantic telephone lines were laid in the mid 1950s¹. Aerial cable attachments to telephone or electric service poles are still the most common coaxial cable placements, however underground facilities have become increasingly more common.

The components used in those early CATV delivery systems were relatively few and simple. There were antennas, twin lead or coaxial cables, combiners, RF amplifiers, and distribution splitters or passives. The fledgling bandwidth limited systems often only carried the 12 VHF (Very High Frequency) channels, and were building the foundational engineering practices necessary to intermodulate video signals and transmit them over long distances while maintaining signal quality and integrity. As the popularity of, and demand for "Cable Television" grew, providers expanded the reach of the one-way video delivery system into what would become the Tree-and-Branch architecture.

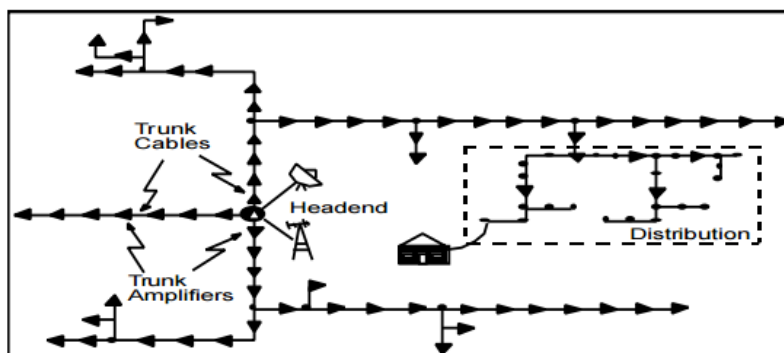


Figure 1 – HFC Amplifier Cascade

The tree-and-branch architecture was comprised of a signal receiving and processing center, called a headend, trunk cables and amplifiers, distribution cables, bridging and line extending amplifiers, distribution passives and "taps" as the drop cable system to deliver the signals to the subscriber television, or terminal. Trunk cable and amplifiers were built to carry the signals long distances from

¹ History of the Telephone - Jason Morris

headend with minimal signal degradation due to amplification. The bridger amplifiers and distribution systems then deliver the signals the “last mile” to subscribing homes via taps and drop cables. Trunk amplifiers can be added one after another in a “cascade” of amplifiers to extend the reach of a cable system, provided that proper engineering and plant management practices are followed to maintain quality signals all the way to the last subscriber terminal. Amplifying intermodulated signals has some potential consequences, however. There are limitations to how many amplifiers can be “cascaded” before signal quality deteriorates to an unsatisfactory level. This tree-and-branch style of architecture was the prevailing build preference for one-way video delivery systems for many years.

According to Roger Hughes, Director of Plant Architectures and Technologies at Armstrong Group of Companies, *“upgrading the amplifiers to push-pull hybrids allowed us to expand the channel plan from 12 channels, up to 22 channels, and extend the reach of the cable plant to longer cascades. Improvements in amplifier technology allowed us to push the bandwidth out to 450 MHz, and cascades longer than 50 amps. That bandwidth expansion resulted in the ability to offer additional video services to our customers, drive penetration, and increase revenue streams. Then in the mid-1980’s fiber optic cables allowed us to cut that cascade in half, and reduce the impact of temperature and weather to that extended trunk cascade.”* In his 40 years with Armstrong, Hughes has seen the architecture grow, and then begin to retract, with the direction of his organization heading toward reduced cascades, smaller data-driven service groups, and ultimately toward passive optical networks (PON).

Cable television programming expanded dramatically in the late 1970s with the introduction of specialized content providers using satellites to make their channels readily available to a great many CATV operators. That technological advancement started a race for expanded downstream bandwidth, and the ability to offer such premium channels as Home Box Office (HBO), the Entertainment and Sports Programming Network (ESPN), and Turner Broadcasting System (TBS). These expanded programming options required more bandwidth to deliver, and upgrades to the cable network amplifier technology. They did not, however, require significant changes to the tree-and-branch architecture, nor the depth of cascade.

The next significant technological advancement to influence cable architecture and design was the introduction of internet services over cable. Cable (internet) modems and cable modem termination systems by companies like Lan City, Com21 and Zenith required that cable systems have operational return paths, and be segmented geographically into smaller footprints or “service groups.”² Fiber optic technology, while still fairly expensive, was already deployed, and becoming more commonplace, and the tree-and-branch architectures were being redesigned using cascade reduction techniques. Fiber optic cables were run toward the further reaches of the cable footprints, along the same routes as the trunk cables, and connected to fiber nodes. These nodes allowed for the long distance transmission of video signals with low losses from attenuation, and reduced distortions attributable to repeated amplification. It also allowed for return path signals to be sent back to the headend, and smaller two-way footprints, or nodes to be established. In this manner, amplifier cascades of 50 or more amplifiers could be reduced to fewer than 10 amplifiers in a single cascade. In 1997, the first iteration of the Data-over-Cable Interface Specification (DOCSIS®) was released, which provided a set of specifications from which vendors could design and build cable modem termination systems (CMTs) and cable modems that would be interoperable for Multiple System Operators (MSOs). DOCSIS® Architecture model: Figure 2.

² How DOCSIS Revolutionized the Cable Industry, 2016 - The Volpe Firm, Brady Volpe

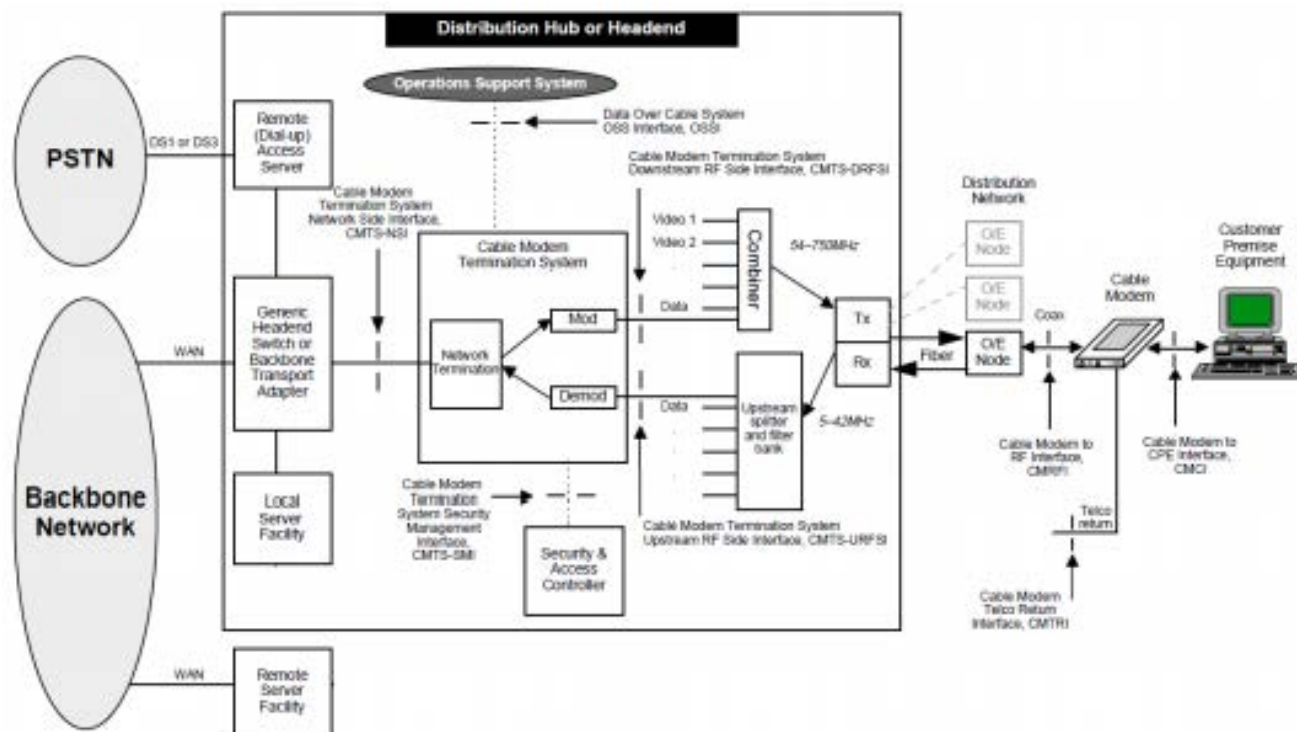


Figure 2 – HFC Logical Digram

The DOCSIS® specified services, such as high speed (broadband) internet and Voice over Internet Protocol (VoIP) became significant income streams for the MSOs, provided that their network or node configurations were built to support the customers' increasing need for broadband services. This data-driven model currently informs both node size and amplifier cascades in the hybrid fiber coax (HFC) systems. From headend-fed, tree-and-branch architectures with 50+ amplifier cascades, to fiber nodes with nary a single amplifier, CATV systems continue to evolve to meet the needs of the subscribers they serve.

Traditional Network Maintenance Practices

As described by Dr. Walter Ciciora in the CableLabs paper *Cable Television in the United States - An Overview*, "Cable television is made possible by the technology of coaxial cable. Rigid coaxial cable has a solid aluminum outer tube and a center conductor of copper-clad aluminum. Flexible coaxial cable's outer conductor is a combination of metal foil and braided wire, with a copper-clad, steel center conductor. The characteristic impedance of the coaxial cable used in cable television is 75 ohms." ³ While there have been improvements in the materials and manufacturing processes, the fundamental physical and electrical properties have remained largely unchanged. As such, the traditional network maintenance practices used to find and fix impairments on coaxial cable networks have also remained largely unchanged. Many of the symptoms of service degradations in the coaxial network can be boiled down to a very few actual causes. Impedance mismatches, shielding integrity issues, and noise are among the chief contributors to service impairments in coaxial networks.

³ Cable Television in the United States - An Overview, 2005 - Walter S Ciciora, Ph. D.

Active sweep and balance, and signal leakage detection remain extremely valuable practices in any architecture type that still relies on coaxial cable as the primary transmission medium. More recently, DOCSIS® protocols include Proactive Network Maintenance (PNM) information derived from Customer Premise Equipment (CPE), which can be included in the symptoms and can provide very detailed information on where network maintenance is required. PNM data can inform where there are reflective cavities from impedance mismatches, frequency suckouts, and standing waves, among other things. These symptoms are indicative of the need to visit the coaxial portion of the plant, and restore the physical and electrical properties of the cable plant.

Beyond the find and fix (Demand Maintenance) functions of a network maintenance team, there are, and have been, preventive maintenance functions that are crucial to the quality service delivery of a cable system. Chief among these, sweep and balance has been integral to keeping the forward and return portions of the plant operating properly. As further explained in Dr. Ciciora's *Cable Television in the United States - An Overview*, "The principal negative of coaxial cable is its relatively high loss. Coaxial cable signal loss is a function of its diameter, dielectric construction, temperature, and operating frequency. A ballpark figure is 1 dB of loss per 100 feet. Half-inch diameter aluminum cable has 1 dB of attenuation per 100 feet at 181 MHz; at one-inch diameter, the attenuation drops to 0.59 dB per 100 feet. The logarithm of the attenuation of cable (in dB) varies with the square root of the frequency. Thus, the attenuation at 216 MHz (within TV channel 13) is twice that of 54 MHz (within TV channel 2) since the frequency is four times as great. If channel 2 is attenuated 10 dB in 1,000 feet, channel 13 will be attenuated 20 dB.... Since attenuation varies with frequency, the spectrum in coaxial cable develops a slope. This is partially compensated with relatively simple equalization networks in the amplifier housings.

"The attenuation of the cable is a function of temperature and aging of components. These amplifiers use a pilot signal to control automatic-gain-control (AGC) circuits. A second pilot signal at a substantially different frequency than the first allows the slope of the attenuation characteristic to be monitored and compensation to be introduced with automatic slope control (ASC) circuits. Thus, long cascades of amplifiers can, once properly set up, maintain their performance over practical ranges of temperature and component aging."

In tree-and-branch topologies with long amplifier cascades, and long cable segments connecting them, losses from attenuation can vary greatly across the course of normal temperature swings in an annual weather cycle. As a result, sweeping these runs twice a year to adjust levels and gain controls in preparation for winter and summer temperatures was critical to ensuring service reliability.

Advancements in PNM Technology

PNM has been around for at least 10 years in mainstream cable operations and many new capabilities continue to be developed. While the existing DOCSIS 3.0 specifications may have reached maturity, new use cases and valuable extensions continue to evolve. Some examples include DOCSIS downstream blind equalization analysis and carrier to interference noise ratio (CNIR) for ingress noise detection. Without a doubt, more capabilities will continue to evolve; however, only adding limited incremental value to any existing PNM stack.

More so, the greatest opportunities ahead exist with artificial intelligence and machine learning to help elevate new use cases and drive operational value in ways that have not been previously conceived. Weather prediction and correlation, such as proposed within this paper, is one example of that.

Within the DOCSIS 3.1 specification, many capabilities remain untapped, largely because they are currently incomplete and unavailable. Even in the case of some functionalities, many operators have not even begun to capitalize on the capabilities that already exist.

Influence of PNM on Network Repair Prioritization

The evolution of intelligence tools, and more recently the proliferation of PNM tools, has had a significant impact on the prioritization of coaxial network repairs. It has been nearly 10 years since Comcast introduced its first PNM software program (the “Scout Flux”) to Beta, and we continue to develop tools that offer greater visibility into the performance of its networks. Add to that the capability of nearly all CPE to report some level of intelligence back to the tool sets, and you have a comprehensive view of all corners of the network. No longer is plant performance informed by a small quantity of DOCSIS channels on a relatively small number of devices. The ability to construct, analyze and match full downstream signatures in multiple portions of each individual premise within a node, in addition to the traditional DOCSIS frequencies, means that an operator can develop an accurate map of all the impairments in the upstream and downstream at virtually every component level within that node. Individual premise issues can be isolated and identified, and with the integration of system design prints, network level impairments can be correlated to active or passive components with a fairly high degree of certainty. This allows for impact scaling, and prioritization of impairment resolution with greater precision, as well as improved task management. Since causality can be attributed to the component level with greater accuracy, dispatch of the proper fix agent is more effective and efficient.

A Glimpse into the Future

Technology consumers today are interacting with our networks and the MSOs themselves in ways that are very different than they were just a few short years ago. They converse digitally with each other, and with service providers, frequently, and across multiple platforms. These new and dynamic communications avenues are going to be critical to the feedback loop and data ingest of our learning machines and decision engines. Intelligent line-of-questions (LOQs) will understand network performance at every level, and inform digital solutions to a dynamic workforce. As a result of the proliferation of Artificial Intelligence and Network Monitoring programs, machines will become more adept at ingesting large quantities of disparate data and connecting symptoms to causality, and dispatching the correct fix agent to affect the repairs. Technology, at its core, should connect people to what’s important to them, and deliver it wherever, whenever, and on whatever device they choose to consume it. Understanding the behaviors of our transmission medium in the context of annualized weather patterns, and building networks that are as impervious to localized weather as possible, will allow us to keep our customers connected to the communications products and services that they rely on us to provide.

Weather

Effects on Compromised Plant

Through systematic observation using PNM capabilities, certain examples of environmentally stimulated plant failures become obvious. One of the most typical signatures recognized by field technicians is referred to as a “suck out.” In actuality, an RF suckout is a half period of a standing wave. That’s to say, if a longer contiguous visible spectrum were available, it would be perceived as a periodic standing wave. Given the physical constraints of velocity of propagation and reflections, suckouts visible within 1GHz of occupied spectrum are limited to very short fault distances. Typically these will be high energy reflections within less than a few inches of reflective cavity. This is important because it makes them very easy to localize with traditional tree-and-branch analysis of the network topology.




Figure 3 – Typical RF Spectrum




Figure 4 – Impaired RF Spectrum, Temperature Induced

Effects on Workforce

The introduction of a weather informed prioritization model to an already dynamic workforce will represent a significant disruption to any work distribution model. The additional layer of weather as prioritization suggests a time element that was not previously a component in the model. Since most weather predictions follow a 10 day model, there becomes a shifting 10 day period that predicts outcomes based on current data sets, and potential vulnerabilities in the node architecture. This level of variability is

quite contrary to many current Demand and Preventive workforce distribution models. The reactive Demand Maintenance activities are already informed by plant failures and weather events, but only after the fact. Preventive Maintenance activities as mentioned previously, typically follow one of two paths: Scheduled or Intelligence tools informed. These models have been optimized by MSOs for many years now, and have provided admirable results thus far. A sliding predictive analysis model that implies prevention of uncertain future events will be difficult to elicit immediate adoption under the best circumstances. Line Maintenance Technicians, like most technicians, are creatures of habit, but tend to be longer tenured and entrenched in traditions and behaviors long earned. Introducing a change with uncertainty to this degree will be a significant challenge to adoption. Additionally, the current points or productivity models used to measure workforce efficiency may also have to undergo some enhancements as tasks are created for repairs prior to the predicted weather behaviors.

Weather Data

Weather data requirements

The weather data is needed for two tasks, each with their own set of requirements. It is assumed that each task will be iterative, with lessons learned being applied to modify the data requirements as needed.

Task 1 is to initially use temporally fine-grained (e.g. hourly) actual data to discover and quantify the strength of correlations between various weather events and plant impairments.

Task 2 is to initially use temporally coarser-grained (e.g. daily) forecast data to predict plant impairments based on predicted weather events, using the correlations established in Task 1. The predictions will in turn feed the plant maintenance ticket Reprioritization Engine (RE).

Geographic scope

The geographic scope of Comcast's nationwide footprint was one of the few weather data requirements known at the beginning. Companies with a smaller footprint may possibly have a wider range of data vendors to choose from.

Task 1 – correlation establishment (pre-trial)

In order to test the viability of the concept we chose six trial markets with diverse climates, and collected weather data for each of them. The trial markets are:

- Albuquerque
- Denver
- Miami
- Minneapolis
- Philadelphia
- Seattle

The correlation establishment phase of the project requires fine geographic and temporal resolution of the actual weather data, in order to establish reliable correlation data between various weather events and various plant impairments. The assumption of the need for frequent updates was based on our assumption that tight temporal and geographical coupling between network impairment data and weather data could increase accuracy in determining the existence and strength of the correlations. Thus, the time scale

would need to be finer than every 12 or 24 hours. The frequent data collection during the trial period will likely be enough to establish the correlations. Once the trial is over we can optionally reduce the collection frequency of actual weather data.

Note that weather forecast data is generally available on a 12 or 24-hour update cycle, not hourly or even finer, as the actual weather data is.

Initial assumptions

The details of the first iteration of requirements were a chicken and egg problem, as we didn't yet have any data to analyze in order to determine what weather data fields would be useful, and how geographically close the weather station needed to be to the network assets to generate reliable correlation data, and what geographic granularity would be desirable in order to determine correlations between specific weather data and specific plant impairment data. Once we collected several months' worth of data we would be able to analyze the correlations and then refine the data requirements based on what was proven to be useful or not.

Our initial guess regarding likely correlations of weather data vs. Plant impairment data included the following hypothesis:

Num	Description	Symptom	Comcast Data source	Notes
1	Cracked coax results in water incursion	High freq rolloff	SPECTRA	Can form a cavity. Water can close it up.
2	Wind results in intermittent RF dropouts	DOCSIS deregisters, video macroblocking	WOPR	Collected every 4 hours in WOPR. Daily snapshots in MELD.
3	Failing amp affects spectrum power	DS MER impairments, many symptoms depending on root cause and line gear type.		
4	Bad shielding leads to noise ingress	DOCSIS US and DS impairments		
5	Fast temp change causes shielding problems due to different metal expansion rates.	Noise incursion, CPD.		Metals plate out, form diodes, becomes a mixer.
6	High temp causes electronics failure (power supply, amp, node)	Low or missing RF		

Based on those initial guesses we required at least these fields in the weather data:
Temperature

Wind speed and direction
Precipitation

Any extra fields, such as humidity, barometric pressure, UV index may prove to be useful as we gain more experience establishing correlations, so will be retained if supplied by the vendor.

Task 2 – Ticket Reprioritization (trial phase)

Once the correlations have been established, it is anticipated there will not be a need for ongoing frequent data collection for that purpose once that phase is completed. Collecting data less frequently after the correlation establishment period will reduce the data cost and computational and storage resource requirements.

Selection of weather data vendor

Vendor selection required some investigation to find out which vendors had actual and forecast data with the desired data characteristics to meet our requirements. These fell into two broad categories: Government-supplied data and commercial data.

The government data lacked the geographic granularity needed, especially for rural areas. Often it is collected at airports, which we assumed would be fine for small or mid-size cities, but not very large cities or smaller towns that lack airports. We felt that the correlation establishment phase would require at least ZIP code level resolution.

Private data vendors offer much finer geographic granularity. We selected Weather Underground, which has some 250,000 weather stations online.

Collection and storage of weather data

Our trial consisted of six cities with diverse climates. For each city we selected one existing weather station from which to collect data, and correlated that to network elements within two miles of that weather station.

Weather Underground has an API that is used for data collection. For our trial we launched a curl command to each weather station in the trial as an hourly cron job.

The data is returned as a JSON file. This is then parsed and inserted into an SQL database by code that Vishnu wrote.

Validation of weather forecast data accuracy

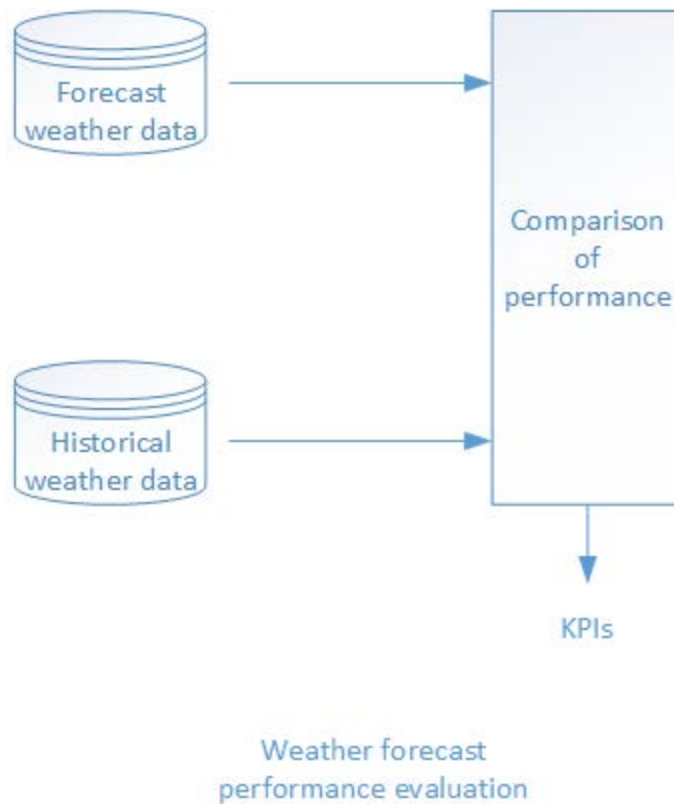


Figure 5 – Performance Evaluation

The point of this project is to have enough confidence in weather forecast data to use it to reprioritize plant repair tickets based on predicted impending weather events. In order to have that confidence, we must analyze the expected errors in the forecast data.

Weather Underground provides a 10 day forecast for various regions, which is updated every day. Note that the forecast region data covers a far broader area than a particular weather station. For example, there is one regional forecast for the Denver metro area, which may contain hundreds of individual weather stations.

In order to analyze the accuracy of those forecasts for each region, we start with the actual data for individual weather stations of interest within that region, for a particular date, and then compare the forecasted values for each field (temp high and low, wind speed, and precipitation) for each of the 10 days. For example, if the actual data is available for the 11th of the month, then we will look at forecast data from the 1st in order to get the 10-day forecast value, forecast data from the 2nd in order to get the 9-day forecast value, etc. The error for the 10-day, 9-day, etc. is then calculated and charted.

An example is shown in Figure 6. This data is from the Denver area trial weather station, with the actual weather temperature data for June 29, 2018. This data was very surprising. We expected the 10-day forecast to have the maximum error, with the error diminishing as the lead time diminished, probably in a highly nonlinear way (much worse at 10 days than at 1 day). Instead, the low temp forecast error

decreased in an approximately linear way, but the high temp forecast error actually got worse as the time date got nearer, with the prior day's forecast missing by nearly 10 degrees!




Figure 6 – Forecast Accuracy

Our beginning assumption is that different markets will have different error trend patterns depending on season. For example, Denver is fairly predictable in the summer and winter, but spring and fall are very unstable, given the city's proximity to the Rocky Mountains. As we collect and analyze more data over the span of months and years we will find out if that's true or not for each market.

Formatting and integration of weather data




Figure 7 – Data Process Model

Understanding forecasted and the actual data was crucial for determining the viability of the concept, and as part of this trial we started collecting the data from Weather underground in January of 2018. The next step was to compare the forecasted data to the actual forecast for the day, and to determine the deviation between the forecasted data and the actual.

We received a 10 day forecast and wanted to determine which day out of the 10 day's forecast we can rely on the most. Finding deviations between the forecasted data and actual gave us that idea, which we could later use in our prediction algorithm using ML & AI.

For the trial we continued collecting using the shell script and a cron job that ran every hour to pull both forecast data as well as the actual data. Forecasted data consisted of forecast for the next 10 days and the actual data was hourly reading of that day's forecast.

Forecasted data was also collected hourly, and the average of high and low temperature taken, to compare against the Max and Min value of different weather points from the actual data to determine the deviation.

Actual Weather Data Example:

```
{
  "response": {
    "version": "0.1",
    "termsOfService": "http://www.collectionweatherdata.com/weather/api/d/terms.html",
    "features": {
      "conditions": 1
    }
  },
  "current_observation": {
    "display_location": {
      "full": "Littleton, CO",
      "city": "Littleton",
      "state": "CO",
      "state_name": "Colorado",
      "country": "US",
      "country_iso3166": "US",
      "zip": "80122",
      "magic": "1",
      "wmo": "99999",
      "latitude": "39.590687",
      "longitude": "-104.947212",
      "elevation": "1707.2"
    },
    "observation_location": {
      "full": "Denver Centennial, Colorado",
      "city": "Denver Centennial",
      "state": "Colorado",
      "country": "US",
      "country_iso3166": "US",
      "latitude": "39.59",
      "longitude": "-104.95",
      "elevation": "5620 ft"
    },
    "estimated": {
```

```

},
"station_id":"KCOLITTL344",
"observation_time":"Last Updated on July 27, 5:00 AM MDT",
"observation_time_rfc822":"Fri, 27 Jul 2018 05:00:52 -0600",
"observation_epoch":"1532689252",
"local_time_rfc822":"Fri, 27 Jul 2018 05:01:03 -0600",
"local_epoch":"1532689263",
"local_tz_short":"MDT",
"local_tz_long":"America/Denver",
"local_tz_offset":"-0600",
"weather":"Overcast",
"temperature_string":"61.3 F (16.3 C)",
"temp_f":61.3,
"temp_c":16.3,
"relative_humidity":"86%",
"wind_string":"Calm",
"wind_dir":"West",
"wind_degrees":264,
"wind_mph":0,
"wind_gust_mph":0,
"wind_kph":0,
"wind_gust_kph":0,
"pressure_mb":"1019",
"pressure_in":"30.11",
"pressure_trend":"+",
"dewpoint_string":"57 F (14 C)",
"dewpoint_f":57,
"dewpoint_c":14,
"heat_index_string":"NA",
"heat_index_f":"NA",
"heat_index_c":"NA",
"windchill_string":"NA",
"windchill_f":"NA",
"windchill_c":"NA",
"feelslike_string":"61.3 F (16.3 C)",
"feelslike_f":"61.3",
"feelslike_c":"16.3",
"visibility_mi":"10.0",
"visibility_km":"16.1",
"solarradiation":"--",
"UV":"0",
"precip_1hr_string":"0.00 in ( 0 mm)",
"precip_1hr_in":"0.00",
"precip_1hr_metric":" 0",
"precip_today_string":"0.00 in (0 mm)",
"precip_today_in":"0.00",

```



```

    "precip_today_metric":"0"}
}

```

Forecasted Weather Data Example:

```

"simpleforecast": {
  "forecastday": [
    {"date":{
      "epoch":"1532739600",
      "pretty":"7:00 PM MDT on July 27, 2018",
      "day":27,
      "month":7,
      "year":2018,
      "yday":207,
      "hour":19,
      "min":"00",
      "sec":0,
      "isdst":"1",
      "monthname":"July",
      "monthname_short":"Jul",
      "weekday_short":"Fri",
      "weekday":"Friday",
      "ampm":"PM",
      "tz_short":"MDT",
      "tz_long":"America/Denver"
    },
    "period":1,
    "high": {
      "fahrenheit":"86",
      "celsius":"30"
    },
    "low": {
      "fahrenheit":"58",
      "celsius":"14"
    },
    "conditions":"Partly Cloudy",
    "icon":"partlycloudy",
    "icon_url":"http://icons.wxug.com/i/c/k/partlycloudy.gif",
    "skyicon":"",
    "pop":20,
    "qpf_allday": {
      "in": 0.00,
      "mm": 0
    },
    "qpf_day": {

```

```

    "in": 0.00,
    "mm": 0
  },
  "qpf_night": {
    "in": 0.00,
    "mm": 0
  },
  "snow_allday": {
    "in": 0.0,
    "cm": 0.0
  },
  "snow_day": {
    "in": 0.0,
    "cm": 0.0
  },
  "snow_night": {
    "in": 0.0,
    "cm": 0.0
  },
  "maxwind": {
    "mph": 10,
    "kph": 16,
    "dir": "WNW",
    "degrees": 282
  },
  "avewind": {
    "mph": 9,
    "kph": 14,
    "dir": "WNW",
    "degrees": 282
  },
  "avehumidity": 51,
  "maxhumidity": 0,
  "minhumidity": 0
}
]
}

```

Loading of Actual and forecast data was done in Scala and using spark. The flexibility of Scala and spark combined made the coding effort small and this data was loaded into MySQL to analyze and create Tableau dashboards. This gave us the easiness to select date or date ranges to see the deviations in different weather points.

Tableau Dashboard:




Figure 8 – Temperature Over Time

Integration of weather data and plant data

Analysis of Weather History vs. Plant History

As of July 2018, we have collected about six months of actual weather data for one station in each of the six trial markets, along with about two months of regional weather forecast data for the six trial markets.

We also have a long history of plant data, which has been automatically collected for years, as part of our normal network operations.

With data sets for actual weather data and actual plant data for the six trial markets, we are now ready to look for correlations. The general plan is to feed these data sets into three AI analysis tool instances, each using a different algorithm to find correlations between actual weather data and plant impairment data. The algorithms and configuration parameters will be tweaked to find the strongest and most reliable correlations between the data sets.

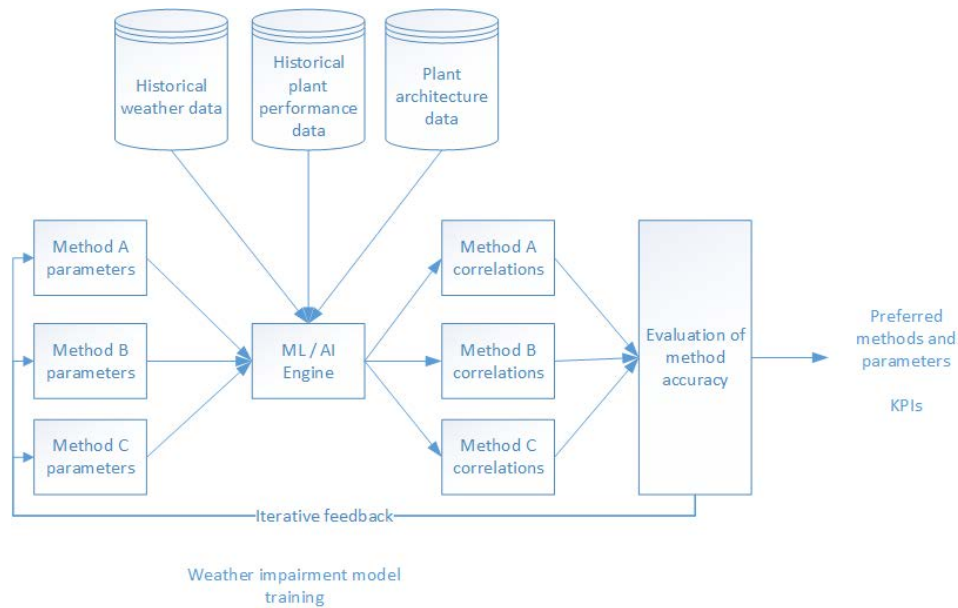


Figure 9 – ML / AI Flow Diagram

Once the machine learning technique is optimized, it can be used to reprioritize the existing plant maintenance ticket backlog, in order to avert customer-impacting problems triggered by impending weather events. The effect of this reprioritization on customer experience can be compared to the BAU process, in order to evaluate the performance of the system. The flow is illustrated in Figure 10, below.

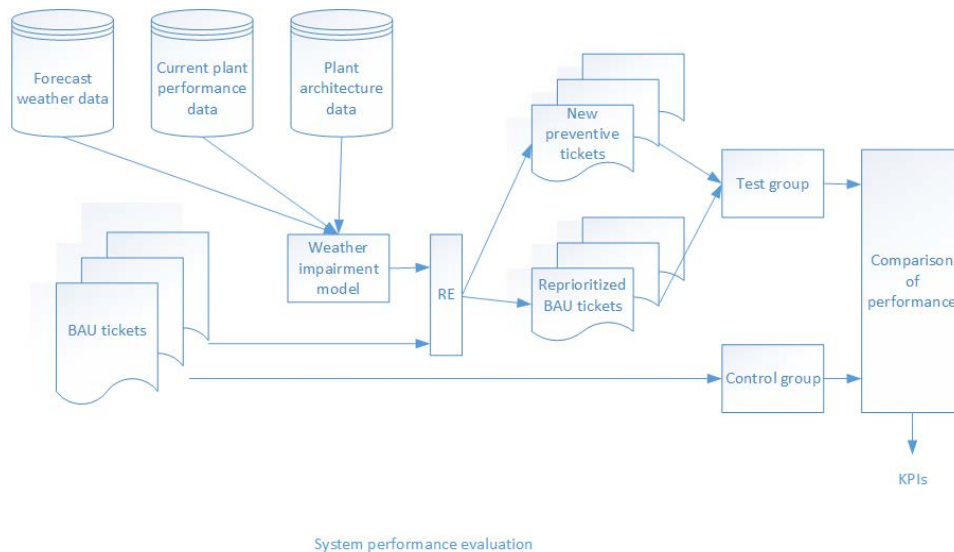


Figure 10 – Analysis Flow Diagram

Characterization of Plant Assets

Previously we discussed the evolution of network architecture, from long tree-and branch-architectures, to the more modern, Node + “x” architectures. As described, the network topology of the cable network is defined by the needs of the local subscriber base. As such, there are many varied and diverse architecture styles still being utilized by MSOs. Within the same service area, operators can have nodes with over 100 amplifiers, and 10 amplifier cascades, as well as nodes with only subscriber taps attached (Node + 0). Each of these architecture styles are matched to their service delivery areas for maximum efficiency, however they could be affected differently by localized weather patterns. For the purposes of determining the impact of weather on coaxial cable networks, it is important that we develop some categories or plant groupings by which we can more clearly understand those impacts. For the purposes of this paper, we have broken the nodes in this study into groups according to several attributes: Aerial cable placements versus underground cable placements, total number of amplifiers in a node, and the maximum number of amplifiers in cascade within a node.

Artificial Intelligence, Machine Learning, and the Future of PNM

With the ever-widening adoption of machine learning into the enterprise, there are many substantial improvements that can be made regarding PNM. Vastly different areas of technology and computing have been affected positively by the research and implementation of these techniques. Take text and video, for example: The relatively new fields of NLP (natural language processing) and CV (computer vision) have made large impacts to those fields. It has also allowed for previously near-impossible tasks to suddenly become possible to give us object recognition in images and more accurate and natural speech to text. As long as the data and telemetry exist, there are many opportunities to enhance current PNM methods as well as enhance predictions beyond the tools currently in use. Additionally, as more telemetry is collected, more often it will become increasingly more difficult to analyze and efficiently gain knowledge from this wealth of new data without the use of machine learning.

The History of Machine Learning And Artificial Intelligence in Comcast Operations

For years, most PNM tools and methods have used conditional logic to implement well known thresholds and alerting from existing device telemetry. These methods have been created and implemented by some of the leading experts in the field and have been instrumental in the shift to proactivity. However, some of these methods can only do so much.

A number of efforts have been researched and worked on. Some of these include predicting when cable modems are active and thus scheduling maintenance activity during periods of predicted low usage, so as to minimize customer impact. Another example is enhanced matching of RF signatures of devices with known issues, so that devices which may have been on the edge of the thresholds are correctly included in an impairment grouping.

Moving forward, many more efforts, aside from this weather correlation study, are underway as well, such as using machine learning and graph theory to better isolate problem RF active and passive elements in the field. Also a number of efforts are analyzing customer impact correlations to certain telemetry.

The Commoditization of ML / AI

With the growth and availability of machine learning software frameworks and tools, it has never been easier to begin leveraging sophisticated techniques that can learn from data and in some cases reduce development time and effort. That said, however, time and investment still needs to be spent on research, development and implementation of these new techniques.

Software

Newer software widely in use today such as scikit-learn python library, or Google's TensorFlow library and framework, lower the barrier of entry and democratize machine learning for a wider audience. In certain circumstances, custom algorithms or implementations still need to be developed, depending on the need of the PNM system or tool -- but mostly these tools should be able to integrate. There are many more fantastic libraries and tools being developed that aren't listed here. Even with all of these, a certain set of skills is required to successfully integrate them into the tool or system.

Skills

There is an oft-cited Data Science Venn Diagram by Drew Conway (reference 3) which portrays the required combination of skills to be successful in implementing machine learning methods. Understanding the difference between classification and regression are good, but to fully grasp why a certain model is producing a specific output it takes some additional knowledge. This isn't to say that every team needs a data scientist, but to dig down into the reasoning behind certain outputs it does help. As the libraries and frameworks move faster forward, the algorithms will continue to improve and become more explainable which will further drive adoption and reduce cost.

Cost / investment

Investment is always an important factor when determining if an existing PNM tool or process should be enhanced, or, in some cases, replaced by a machine learning technique. Usually the highest chunk of the cost is research, followed by development and implementation. The main factor for the increased research cost is simply time. During the research phase there are many items which need to be performed by the researchers and engineers. These are problem definition, data gathering, data cleaning and manipulation, data standardization and normalization, machine learning model development and evaluation and finally implementation.

The Convergence of PNM and ML / AI

Given the fact that machine learning libraries and frameworks are more accessible than ever, greater convergence with PNM will continue. Some of the examples are tool enhancements and new processes, but the time will come when entire systems are replaced with advanced and complementary learning systems. The promise of some of these emerging technologies is that simply based on data, the system can learn optimal paths, unnecessary actions, recommendations and more.

The Composition of an AI Enhanced PNM Program

Hardware & Software

This is very dependent on specific implementations. Our study used weather data from a variety of hardware platforms. From our point of view, the weather station hardware was irrelevant; what mattered was the data was in a standardized format. Computing was done on commodity PC and cloud assets.

Data and Modeling

As previously discussed in the architecture overviews, today's networks are widely varied in their construction and cable placement. For the purposes of this trial, we selected node areas that were comprised of entirely aerial cable placements, or entirely underground cable placements. That singular delineation greatly reduced the volume of nodes considered for evaluation, but those nodes would be more directly informative of the impact of weather on the performance and symptomology being measured. From within those node populations, the next determining factor was cascade length, or depth. There is significant variance in cascade depth among the node populations under review, which ultimately should allow for an analogous population consistent with many architecture styles employed across the country. Understanding that each node could be comprised of both aerial and underground segments, the learnings from this trial, and as the data models evolve, we can take the learnings from the various cable placements, and construct a more hybrid approach to the impact on weather on mixed placement architectures.

Profiles were established for candidate nodes, consisting of of the following features:

- Node size: the number of cable modems being serviced by the node, bucketized into: Node +0 (0), Small (1-19), Medium (20-49), Large (50-99), and Extra-Large (100+)
- Maximum reach: the maximum number of amplifiers between the node and a cable modem, bucketized into: Node +0 (0), Short (1-2), Medium (3-5), Long (6-8), and Extra Long (9+)
- Plant distribution: All Underground, and All Aerial. Nodes with mixed distribution (underground and aerial) were excluded
- Distance from weather station
- Total cable length

The complete feature set used in the machine learning model include:

- Node size: 0 (Node +0), 1 (Small), 2 (Medium), 3 (Large), 4 (Extra Large)
- Maximum reach: 0 (Node +0), 1 (Short), 2 (Medium), 3 (Long), 4 (Extra Long)
- Plant distribution: 0 (underground), 1 (aerial)
- Weather feature
- Impairment feature

Possible weather features include:

- Temperature: degrees fahrenheit
- Relative Humidity: percent
- One hour temperature change

- Average temperature across X hours

Possible impairment features include:

- Noise: numeric measurement or bucketization by severity:
0 (None), 1 (Non-Severe: < 70), 2 (Severe: > 70)
- Suckout: numeric measurement or bucketization by severity:
0 (None), 1 (Non-Severe: < 70), 2 (Severe: > 70)
- Power: 0 (Normal), 1 (Non-Severe), 2 (Severe)
- Forward error correction (FEC)
- Signal to noise ratio (SNR)

Execution

A KPI of the project is the net change in CX brought about by use of the system in response to forecasted weather events. The initial trial will do this via a post-event analysis. This will be done by executing it against the ticket queue of a trial market that subsequently suffered weather damage to the HFC plant which is within the trial radius of the weather station. After the weather event we will have the BAU result, because that's what the market will have executed. By having the Reprioritization Engine (RE) run against the queue prior to the weather event, and analyzing the difference between the actual BAU customer impact scores, versus the calculated customer impact scores had certain tickets been prioritized, we will be able to determine the relative effectiveness of the system.

For example, say a severe rain storm hits a trial market. The BAU result was that a node went down due to a water-induced failure, causing a degraded signal for several hours for 200 customers. This happened because the repair ticket for that node was not high enough in the queue to be repaired before the storm hit. Had the RE been engaged, it would have prioritized that ticket, eliminating the degradation because proactive network maintenance, targeted at the node, suddenly, due to the weather forecast, became a much higher risk asset. Some other enqueued ticket(s), not at increased risk because of rain, got deprioritized. The effect of that will be weighed against the RE's improvement from keeping the node nominal, yielding a net change in customer impact score. That net change is a KPI.

Addressing Common Barriers to Adoption

There are several common barriers to adoption when introducing change into any organization, and particularly where the change is requiring the frontline technical workforce to execute or support the change. The first is the Operational Benefit to the Organization, or the Return on Investment (ROI), and the second is the "What's in it for Me" (WIIFM), as it pertains to the frontline technical workforce (adoption and execution). ROI as it pertains to prevention in the CATV model is generally outlined in a couple of key metric reductions, i.e. outages or trouble calls. As with any preventive program, it is often difficult to measure something that doesn't happen in a limited time period. Year over Year (YoY) performance comparisons in any given metric category can provide early anecdotal data to inform the potential performance gains in a trial or small market rollout. These market trials have historically been the basis upon which programs have been measured with respect to the ROI, and evaluated by the organization as to the worthiness to initiate the change to a larger, if not institutional scale.

Further, results of any level of adoption can be colored by the willingness of the frontline workforce to accept the change. If the initiative is overly complex, and doesn't result in a financial benefit for the frontline users, it may be difficult to institute, to prove out the benefit or ROI for the organization. When dealing with any change to the processes or procedures affecting frontline employees, training hours must be budgeted to outline the change, and fully define the WIIFM associated with the proposed change.

Technology and Operational Discontinuity

It is well established that technology and operations can be successfully integrated, but only through a dynamic integration model that is capable of connecting thought leaders and technologists with execution leaders in operations. Discontinuity, by definition, infers a break in something. Depending on perspective, discontinuity can be perceived as something is broken, like a process or program, or it could be described as a necessary disconnect while the organization course corrects in a more favorable direction. The evolution of network architecture has lent itself to successive changes in construction, design and plant management practices. Any operational or engineering changes suggested by the concepts or ideas within this document, would by definition require a temporary course correction, or discontinuity in the current operations model to provide early estimates on an ROI, and inform decisions on wider scale deployment.

New Prioritization Around 10 Day Forecasts

Solving Priority Conflicts

There are several challenges with any preventative activity or program that have to be overcome, or at the very least, acknowledged in order to evolve a model. First is the understanding that any plant management model is truly preventative. Second is quantifying a return on investment (ROI) calculation, attempting to monetize the impact of an event that was predicted to occur, but did not. Every organization has business imperatives that must be achieved in order to meet obligations to its customers, employees or stakeholders. Those must be balanced against the operational imperative of protecting, improving performance of, and further monetizing their architectural assets. In order for the operational model to evolve to include a weather informed prioritization, it must first be determined to provide tangible benefit to the business imperative. By continuously engaging a weather analytic engine to a local geography, and analyzing key performance indicators against weather forecast data, one could argue that the prioritization model moves closer to achieving prevention in a meaningful way.

Existing Priority System

Early models of network maintenance prioritization relied on call center data to report areas of the tree-and-branch architectures to the line maintenance teams for investigation. In forward-only, broadcast video delivery systems, there were no other telemetry data sets to which line tech teams could respond. All maintenance was scheduled according to an annual visit cycle, or as a result of customer-reported trouble. As network topologies have gotten smaller, the individual node counts have risen significantly, and knowing the lines of demarcation between service areas had gotten more difficult. As technology and intelligence tools have evolved, so too have the potential data sets available to operators to build node performance metrics, and prioritization methodologies. According to Comcast XOC Engineer Larry Scott, Watchtower and Equilibrium are two of the internal toolsets that allow Comcast to performance-rate and prioritize its hundreds of thousands of nodes.

Watchtower is the primary and first level monitoring tool. It ingests information from multiple intelligence programs, and assigns performance values to each node according to the telemetry data from

all of the DOCSIS-based CPE within that particular node. Comcast uses a device poller that calls to each device, retrieves the RF and impairment metrics for that device, then aggregates that data into the Watchtower interface. Each device location is then color-coded relative to its performance data for quick visual analysis, as well as correlated to adjacent devices in close proximity. If multiple devices in the same geographic area are determined to share a similar impairment, then Watchtower declares that those devices are part of an Event. Events are ascribed a score according to multiple logic points built into a scoring algorithm. Each node is then able to be assigned a “node score”. In this model, the higher the node score, the more impaired devices there are within the node. A node score of zero would imply that the devices within that node are operating in an environment that is free from impairments.

Nodes can then be preliminarily prioritized based on their score within a functional team area. Some of the customer premise device attributes that are measured for the purpose of ascribing node score include: Receive and transmit levels, Signal to Noise Ratio (SNR), Modulation Error Ratio (MER), and Codeword Errors (CER). Additional attributes are sent to Watchtower from other toolsets that analyze PNM data points, such as Comcast’s Spectra tool. Spectra analyzes the DOCSIS devices that are capable of providing full band capture information, and matches those signatures, based on impairment logic in the tool, to inform Watchtower when there are groups of devices within a node suffering from suckouts, waves, or other downstream signature impairments.

Another Comcast tool that informs Watchtower device performance is SASQWatch (Service Affecting Stream Quality Watch). SASQWatch looks at node metrics from the CMTS port perspective, as opposed to the aggregate DOCSIS device perspective, to inform data packet loss affecting the upstream CMTS port. Each CMTS port is analyzed at regular intervals, and if a threshold is breached, a message is sent to Watchtower to create an event, increase the score, and by extension, the priority. As Larry Scott explains, not all events are created or measured equally. Certain event types have a base score multiplier consistent with their severity. Like the SASQWatch events just described, a Severe SASQWatch event has a 35 point value multiplier, to increase the score more quickly. For example, if a CMTS port experiences severe packet loss, the amount of customers associated with that event will be multiplied by the time they were impaired, then multiplied by the event multiplier, in this case 35, to derive a total event score. All of the individual events are then summed, to equal the total node score. Node scores are characterized further by the aggregate node scores across incremental time periods (i.e. Current score, One Day score, and Three Day score). This can provide context to the duration of the impaired Customer Experience, as well as which nodes are currently impaired, and are immediately actionable by available line technician resources.

In addition to the device performance metrics, Watchtower also monitors all the customer premise equipment for an offline state, indicative of an outage. These events are “soaked” for a period of time for validity, then issued immediately to the available technicians on duty for assignment and correction via its embedded communications engine. Individual users subscribe to automated messages via e-mail, text message or both. The interface also houses the technicians’ schedules and individual roles. Techs can be assigned to particular maintenance roles, such as Demand Maintenance (Outages, Severely Degraded Service events and Single Customer escalations), or Preventive Maintenance, (Sweep, Optimization and Leakage). Because Watchtower understands the number of techs on duty, and their individual roles, it routes the appropriate jobs to the tech best equipped to handle that particular task.

Node scores are prepared by Watchtower about every four hours, and sent for final prioritization in Comcast’s prioritization engine, called Equilibrium. EQ, as it’s known, is the intermediate funnel between Watchtower’s analysis and scoring engine, and its workforce management interface. EQ also ingests the

local maintenance technicians' schedules and roles, and has settings to manage the number of node repair jobs that get issued into the maintenance team's work queue. Teams can adjust their job volume in different ways: Either schedule-based, according to the number of techs on active duty, or on a static basis. EQ administrators can set the job volume to one job per person or more, based on system needs, or as a static volume of a fixed number of jobs at any given time. In either case, as soon as a technician completes one job, EQ automatically issues the next and highest priority job into the queue. One job closes, the next automatically appears. Equilibrium's primary function, however, is to allow for customized prioritization at the local team level, while still maintaining focus on the primary customer impacting metrics. For instance, two nodes may have the same score, however one node may have a higher customer trouble call count than the other. EQ will issue the node with the highest trouble call count first. It also maintains score history, for the purpose of prioritization. If a node is partially impaired for 30 days, EQ will add score to the node in the background to increase its priority. Equilibrium can also be programmed to execute a function called "Compare and Replace". If a node has been issued to the work queue, but hasn't been issued to a technician yet, EQ may replace the node, with one of a higher priority. It is constantly analyzing every tech team's work queue, and searching the total node population for the most impactful repair to be repaired.

Proposed Priority System

As mentioned previously, Comcast uses a comprehensive prioritization engine, with multiple layers of priority "accelerators". It is likely that all cable operators today have their own versions of prioritization, predicated on weighting symptomology that is closely aligned to the customer experience. They may also layer in trouble call and repair call volumes. It is worthwhile mentioning that none of these engines create additional work. Rather, their intent is to weigh current symptomology with predictive data to manage their human resources, and apply them to the customer population that needs relief most immediately. These models, however, likely do not layer in architectural references in an informative fashion. They may consider node size, but only in so much as it informs the volume of devices attached to the network or node area.

This study is intended to add an additional predictive data point, which the operators can add as a more dynamic "last layer" of prioritization. For instance: In periods of relative stability in the weather forecast (based on forecast accuracy), the prioritization model would continue to operate with the original business rules applied to the program. In this model, the PNM predictive indicators, in conjunction with architecture qualifiers, will lend additional "weight" to the health of the node area for prioritization. As such, an operator would then, by necessity, have to examine its plant assets with a more myopic lens, taking into consideration regional weather patterns, as well as any architecture data available indicating cable placement and cascade depth.

Adapting This Into Existing Workforce Scheduling

Enabling a prioritization model based on the traditional 10-day weather forecasting scheme would require several things. First, a high degree of confidence in the predictive model would have to be established, as well as an agile workforce distribution model. The workforce management model would need to shift a portion of a limited resource pool from Demand activities, to Preventive activities in a dynamic fashion. Lastly the operator would need to add an additional layer of prioritization into its work management system based on regional weather and plant architecture design.

Operational Culture and Workforce Training Considerations

We have already referenced the tribal knowledge, or operational culture of a system or geography, as it has been shaped over the years by operation and business philosophies, as well as weather. What we have yet to really dive into, is the impact of enhanced intelligence tools, or PNM data to the operational training and change communication strategy. What is often overlooked in the proliferation of intelligence that defines symptomology is the repair stratagem of the line maintenance teams themselves, with respect to the actual repairs of the hardline plant. As mentioned earlier, the physical components comprising the Hybrid Fiber Coax (HFC) network have changed remarkably little in the previous six decades of cable television architecture. As such, the coaxial plant management practices themselves have changed remarkably little. Of primary import, are two very simple goals, which, because of scale, are remarkably hard to achieve. The goal of any plant management plan can be boiled down to: Nominal 75 ohm impedance, and shielding integrity. These two properties are the end goal of any plant repair effort. Sweep and Balance programs, leakage rideout programs, and noise mitigation can be argued as the primary daily activities of most line maintenance technicians, and all of those hours of efforts can be attributed to the same end result: 75 ohm impedance, and a closed system. This is not an indictment of PNM data or other intelligence tools. Quite the contrary, in fact. Those intelligence tools have provided unprecedented visibility and localization of impairments, and have allowed operators to more properly scale repair efforts, and ensure plant integrity post repairs. We mention the simplicity of the plant remediation efforts in the context of training, and the impact to the operational culture. With this in mind, and in the context of this trial, training then can be focused on the model used to prioritize plant repairs, rather than introducing a new metric or symptom into the mix. Because the plant remediation methods won't necessarily change (only the order in which they're applied to the node population), training the workforce should not add significant technical complexity to the change management process. Finally, the operational culture should likewise not suffer significant impact, as the remediation processes employed by the technical workforce will also remain largely intact.

Organizational Alignment and Common Understanding

When large scale initiatives and large organizations collide, there is typically a wide range of adoption levels and results, and not all of them positive. It's similar to the old telephone game, where the message delivered by leadership, doesn't always resemble the message received by the end user. Further, different parts of the organization may have different goals, which may not be totally aligned. When introducing any change, it's incumbent on the owners of the initiative to understand the deltas in the ability of their regional teams to be able to ingest, and institute change. Particularly where preventive maintenance initiatives are concerned. Common understanding is an end state that ultimately determines the engines that any changes agency uses. Things to consider in determining the change management engine could include: Current state of the region/system, communications and training structure, or regional/local engineering educational structure.

Conclusion

It is well known experientially and anecdotally that weather behaviors can have significant impact on the performance of the outside plant networks, and the operator's ability to ensure trouble-free service delivery to their customers. The body of this work suggests that weather forecast data with greater than 90% accuracy in key weather metric categories, when leveraged against PNM datasets, can inform the operator of potential failures within their networks. This presents a window of opportunity to address the vulnerabilities, and minimize the potential impacts of that weather event to the customer experience.

Abbreviations

AI	artificial intelligence
AGC	automatic gain control
ASC	automatic slope control
ANN	artificial neural network
BAU	business as usual
CATV	cable television
CER	codeword error ratio
CINR	carrier-to-interference-plus-noise ratio
CLEAR	comcast leadership, engineering and relationships
CMTS	cable modem termination system
CPD	common path distortion
CPE	customer premises equipment
CV	computer vision
CX	customer experience
DevOps	development and operations
DOCSIS ®	data over cable service interface specification
DS	downstream
EQ	equalization, or equilibrium
ESPN	entertainment and sports programming network
FEC	forward error correction
FM	frequency modulation
FTTP	fiber to the premises
HBO	home box office
HFC	hybrid fiber-coax
Hz	hertz
ICFR	in-channel frequency response
ISBE	International Society of Broadband Experts
JSON	javascript object notation
KPI	key performance indicator
LOQ	line of question
MER	modulation error ratio
MSO	multiple-system operator
NLP	natural language processing
PON	passive optical network
PC	personal computer
PM	preventive maintenance

PNM	proactive network maintenance
pre-EQ	adaptive pre-equalization
RE	reprioritization engine
RF	radio frequency
RoI	return on investment
SA	spectrum analyzer
SASQwatch	Service Affecting Stream Quality Watch
SCTE	society of cable telecommunications engineers
SID	spectral impairment detector
SNR	signal-to-noise ratio
SQL	structured query language
TBS	turner broadcasting system
US	upstream
UV	ultraviolet
VHF	very high frequency
VoIP	voice over internet protocol
WG	working group
WIIFM	what's in it for me
YoY	year over year

Bibliography & References

- [1] CableLabs DOCSIS® Best Practices and Guidelines - Proactive Network Maintenance Using Pre-equalization. 2012
- [2] Expo2016, A Comprehensive Case Study of Proactive Network Maintenance – Larry Wolcott
- [3] The Data Science Venn Diagram – Drew Conway, <http://drewconway.com/zia/2013/3/26/the-data-science-venn-diagram>

Achieving Significant Space, Energy, And Cost Reductions With Future Virtualized Distributed Access RPD And RMD Architectures For MSOs

A Technical Paper prepared for SCTE•ISBE by

R. J. Vale

Leader of Strategic Techno-Economics – Future Networks
Nokia, Bell Labs Consulting
+1-972-477-8674
Rj.vale@bell-labs-consulting.com

Martin J. Glapa

Partner and Bell Labs Fellow
Nokia, Bell Labs Consulting
303-517-1273
Martin.Glapa@bell-labs-consulting.com

Jean-Philippe Joseph

Fixed Access Networks Team Lead
Nokia, Bell Labs Consulting
+1-908-582-2903
Jean-Philippe.Joseph@bell-labs-consulting.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
The Shift to DAA	3
DAA Architectures.....	4
Benefits of DAA	5
1. Space utilization analysis results	7
2. Energy analysis results	8
2.1. Edge Facility.....	8
2.2. Outside Plant.....	9
Conclusion.....	10
Abbreviations	10
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1 - CCAP and DAA Architectures	5
Figure 2 - Realization of vCMTS RMD.....	6
Figure 3- Realization of vCMTS RPD	6
Figure 4 - Edge Facility Space Modeling Results	8
Figure 5 - Edge Facility Energy Modeling Results	9
Figure 6 - OSP Energy Growth for Each Architecture Comparing N+6 to N+3	10

List of Tables

Title	Page Number
Table 1 - Space Increase at Edge Facility	7
Table 2 - Edge Facility Energy Increase	8

Introduction

Cable access networks are on the cusp of a major transformation driven by an insatiable consumer appetite for multimedia and Over the Top (OTT) content, social media, learning and communication. This hunger is enabled by the digitization of everything and the automation of the world around us, both of which are enhancing human experiences and forever changing our lives. Satisfying this appetite cost-effectively requires a foundational transformation in cable access network architecture to provide the capacity and performance needed to satisfy these enhanced human experiences and needs.

The use of Distributed Access Architectures (DAA¹) will transform cable networks to a degree that has not been seen since the late 1980s and early 1990s when fiber began to play a prominent role in Hybrid-Fiber Coax (HFC) networks. Today's asymmetrical multi-Mbps capacity HFC networks will shift over time to a hyper-scaled symmetrical multi-Gbps capacity HFC network. DAA will deliver these speeds with space and energy savings when compared to today's network architectures.

However, this shift will impact cable edge facility² space requirements, and edge facility and Outside Plant (OSP) energy requirements. This paper presents an analysis of an exemplary edge facility and its associated OSP. It compares the space utilization and energy consumption of three DAA architectures to the conventional I-CCAP (Integrated-Converged Cable Access Platform) architecture. We demonstrate that when migrating from n+6 to n+3, a virtualized Cable Modem Termination System (vCMTS) Remote MACPHY³ Device (RMD) based architecture provides the greatest savings in edge facility space (66%) and energy consumption (86%) over I-CCAP. Migrating to DAA increases OSP energy needs in all architectures, however, there is a variance of just 5% among the DAA alternatives at n+3.

The Shift to DAA

Bell Labs' Future X⁴ Massive Scale Access vision is driven by new technological capabilities and critical digital network needs that will deliver a cost/capacity/bandwidth value transformation to MSOs and consumers. In summary, these capabilities and needs are:

- Seemingly infinite hyper-capacity at 100x growth over the next decade.
- Unlimited on-demand capacity for any application or service.
- Tera-hyper-scaling of networks to support trillions of connected systems, devices, processes, objects and automata.

MSOs need to meet these digital network needs with a network that delivers increased capacity, greater flexibility and reduced complexity, with increased efficiency and reusability – all at a reduced cost. DAA is a critical part of the Bell Labs vision and will enable several major cable access architecture shifts:

- Capacity expansion - DOCSIS® 3.1 increases capacity through improved spectral efficiency, increased Upstream (US) spectrum (and capacity), increased aggregate US/Downstream (DS) spectrum to 1.2 GHz, and many other means. Moving fiber deeper through physical node splitting and node relocation enables spectrum reuse and higher average bandwidth per consumer. Full Duplex (FDX) DOCSIS, currently under definition, will provide symmetrical multi-Gbps data

¹ Defined later in this paper

² Also, may be known as a hub

³ Media Access Control/Physical Layer, also referred to as MAC/PHY or MAC-PHY

⁴ The Future X Network, A Bell Labs Perspective, 2016, M. Weldon and all, CRC Press – Taylor & Francis Group

speeds on existing coax. Longer-term, fiber can be extended to the last coax drop tap, to within tens of meters of about four to six consumers, exploiting spectrum above 1.2 GHz to reach up to 30 or 40 Gbps. DAA architectures will help enable this transformation to FDX.

- Energy consumption reduction and space reduction - Edge facility energy and total energy (i.e., edge facility and OSP combined) improves as devices become more efficient and as functionality (e.g., MACPHY) is moved from the edge facility to the OSP.
- OSP modernization - DAA will trigger a shift from analog fiber to digital fiber, providing a multitude of advantages including: use of higher-order QAM modulation, reducing maintenance costs, and traversing longer distances.

DAA Architectures

I-CCAP is the baseline centralized architecture being deployed today, with integrated Media Access Control (MAC) and Physical (PHY) layer functions located at the cable edge facility and connected to a remote fiber node via analog fiber. Three primary distributed architectures (DAA) have been defined by vendors, MSOs, and CableLabs®:

- CCAP Core Remote PHY Device (RPD) - A non-virtualized CCAP core containing the DOCSIS® 3.1 MAC function with the PHY function in a remote node connected to the edge facility via digital fiber.
- vCMTS RPD - A vCMTS function running the DOCSIS 3.1⁵ MAC on an off-the-shelf server in an edge facility or centralized location with the PHY function in a remote node connected to the edge facility via digital fiber.
- vCMTS RMD - A vCMTS running DOCSIS 3.1 MAC and PHY functions in a remote node connected to an edge facility via digital fiber.

⁵ Could also be DOCSIS 3.0, however DOCSIS 3.1 is the current technology

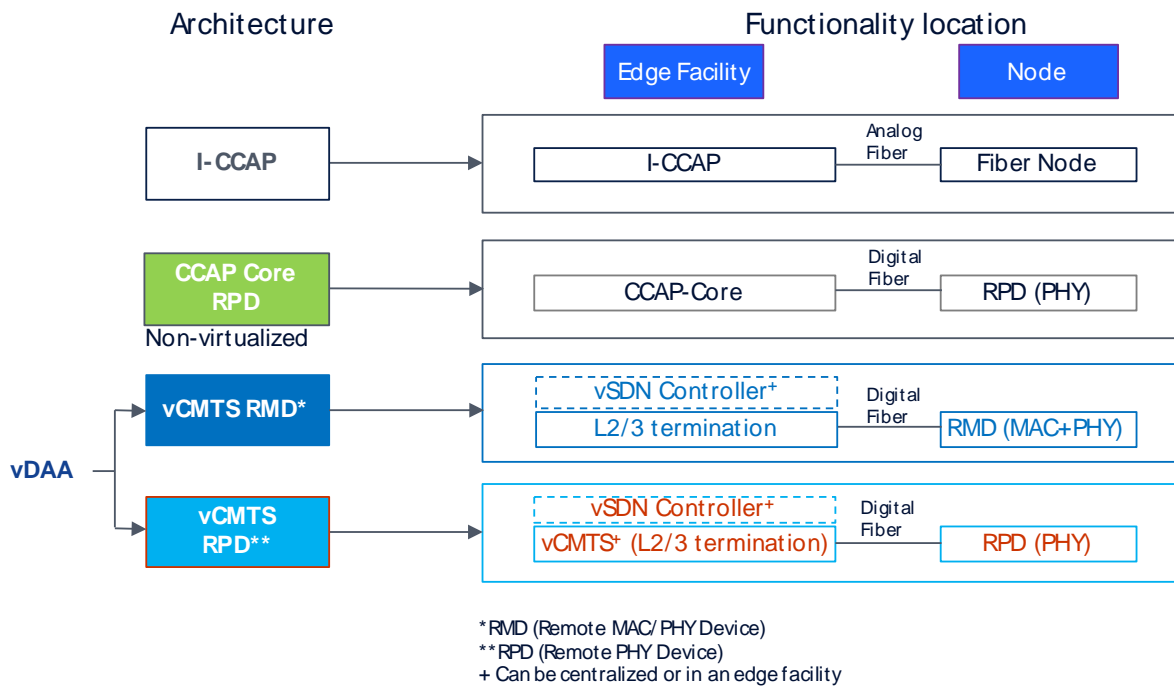


Figure 1 - CCAP and DAA Architectures

Benefits of DAA

Bell Labs Consulting modelled the CCAP Core RPD, vCMTS RPD, and vCMTS RMD based architectures to compare them to the I-CCAP based architecture for space utilization and energy consumption. To provide a holistic view, Total Cost of Ownership (TCO) modelling analysis included equipment required to deploy each of the DAA architectures and the I-CCAP architecture from the edge facility out to the node, including off-the-shelf servers, switches, routers, video equipment and racks. To keep the comparison between these architectures consistent, the modeling was performed on an exemplary edge facility supporting 50,000 households passed and a representative OSP topology. Significant differences exist in how these architectures compare across space utilization and energy consumption. Figure 2 illustrates the equipment in a standard datacenter rack that support a vCMTS RMD solution.

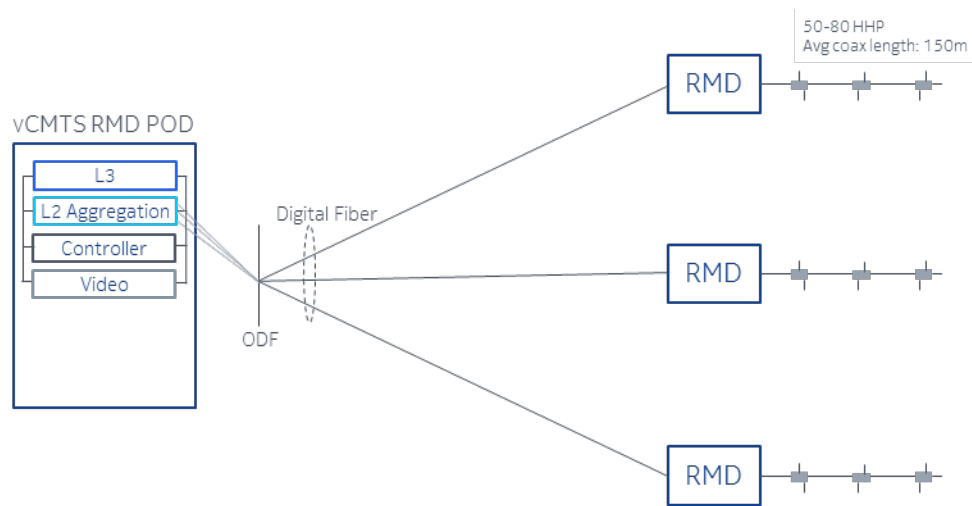


Figure 2 - Realization of vCMTS RMD

Figure 3 illustrates the equipment needed in a standard data center rack for a vCMTS RPD solution realization. A leaf-spine architecture is used to interconnect multiple Commercial Off-The-Shelf (COTS) servers with essential vCMTS terminating components.

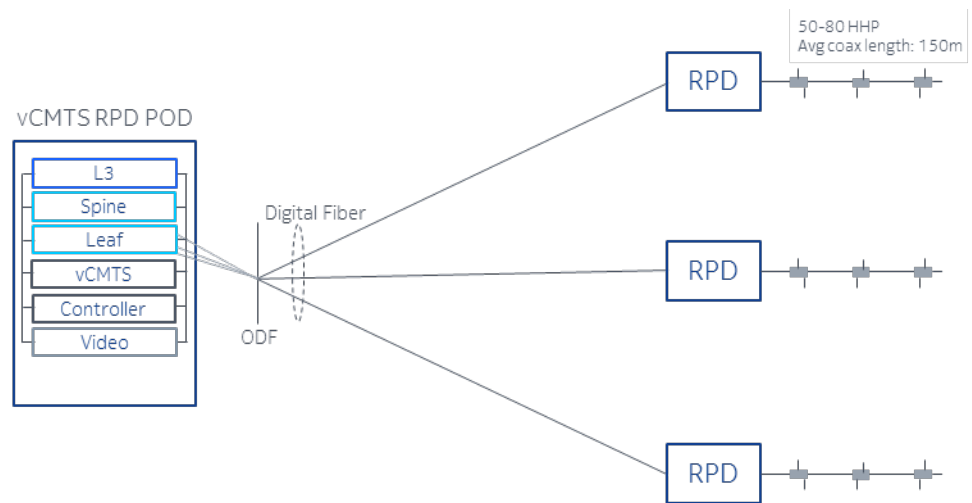


Figure 3- Realization of vCMTS RPD

Key modeling assumptions

Several hundred assumptions are included in the models. Assumptions that have the greatest impact on the analysis results are listed here.

Assumptions common amongst all architectures:

- 50K HHP coverage by edge facility
- 30K consumers served/edge facility
- 42 Rack Unit (RU) rack size
- 10 KW per rack limit

- 2 service groups per node

I-CCAP specific assumptions:

- 16 RU chassis with a density of 96 service groups

CCAP Core RPD specific assumptions:

- 16 RU chassis with a density of 192 service groups

Assumptions common to all vCMTS architectures:

- 600 watts, 1RU COTS compute server trays
- Video servers of similar capability

vCMTS specific RPD assumptions:

- 32 service groups per server tray
- 20 server trays per POD hosting vCMTS function – no redundancy

Controller assumptions:

- 500 service groups per server tray for controller (1+1 configuration⁶)

1. Space utilization analysis results

Edge facility space requirements increase to accommodate additional equipment required as the number of service groups grows and as the OSP architecture correspondingly evolves from N+6 to N+0. The following table illustrates this growth to N+3, using N+6 with 96 SGs (i.e., 48 nodes) as the baseline.

Table 1 - Space Increase at Edge Facility

N+6 - Baseline Architecture	N+3 - Edge Facility Space Increase (over N+6)
I-CCAP	3x
CCAP Core RPD	3x
vCMTS RPD	3x
vCMTS RMD	2x

Remote placement of vCMTS RMD MAC and PHY functions in the OSP provide significant edge facility space utilization (rack unit) savings over I-CCAP and all other DAA approaches.

As seen in Figure 4, the gradual rise of the vCMTS RMD plot reflects its ability to scale more efficiently than other DAA solutions.

⁶ One active, one standby




Figure 4 - Edge Facility Space Modeling Results

2. Energy analysis results

2.1. Edge Facility

Edge facility energy usage increases as the number of service groups grows and as the OSP architecture correspondingly evolves from N+6 to N+0. The following table illustrates this growth to N+3, using N+6 with 96 SGs as the baseline.

Table 2 - Edge Facility Energy Increase

N+6 - Baseline Architecture	N+3 - Edge Facility Energy Increase (over N+6)
I-CCAP	5.2X
CCAP Core RPD	4.5X
vCMTS RPD	2.8X
vCMTS RMD	1.8X

Figure 5 illustrates this growth in edge facility energy consumption showing estimated edge facility equipment energy usage (in watts) for the four architectures. Since heating, ventilation and air conditioning would account for a constant multiple based on Power Usage Effectiveness (PUE), it is not shown in the calculations below.




Figure 5 - Edge Facility Energy Modeling Results

All DAA architectures consume less edge facility energy than I-CCAP due to the distribution of key functions into the OSP. Requiring only a controller function at the edge facility and distributing both the MAC and PHY functions to remote nodes, vCMTS RMD reduces energy consumption in the edge facility more than any of the other architectures.

2.2. Outside Plant

OSP energy usage also increases as the number of service groups grows and as the OSP architecture correspondingly evolves from N+6 to smaller N+x. The following figure illustrates this growth as the OSP evolves from N+6 to N+3. It is normalized to RPD at n+3.




Figure 6 - OSP Energy Growth for Each Architecture Comparing N+6 to N+3

A vCMTS RMD-based OSP consumes about 5% more energy than an RPD-based OSP at n+3, due to the MAC and the PHY being collocated. However, the vCMTS RMD solution offers significant savings in the energy consumed at the edge facility as shown in Figure 5 which makes it attractive from an OpEx perspective.

Conclusion

Both vCMTS RPD and vCMTS RMD solutions reduce edge facility space and energy consumption when compared to I-CCAP and CCAP Core RPD solutions. vCMTS RMD provides the greatest savings in edge facility space (66%) and energy consumption (86%) compared to I-CCAP at n+3. vCMTS RMD consumes slightly more OSP energy (about 5%) than the RPD alternatives. Although, these space and energy savings are an interplay of many factors, the results herein are representative of the typical needs and benefits for larger cable edge facilities.

Abbreviations

CMTS	cable modem termination system
COTS	common off the shelf
DAA	distributed access architecture
DOCSIS	data over cable system interface specification
FEC	forward error correction

HFC	hybrid fiber-coax
HHP	households passed
I-CCAP	Integrated-common cable access platform
ISBE	International Society of Broadband Experts
MAC	media access control (layer)
Mbps	megabits per second
MSO	multiple system operator
OSP	outside plant
OTT	over the top
PHY	physical (layer)
PUE	power usage effectiveness
RMD	remote mac and phy device
RPD	remote phy device
RU	rack unit
SCTE	Society of Cable Telecommunications Engineers
TCO	total cost of ownership
vCMTS	virtualized cable modem termination system

Bibliography & References

The Future X Network, A Bell Labs Perspective, 2016, M. Weldon and all

Adaptive Power Management for Node Clusters

A Technical Paper prepared for SCTE•ISBE by

Fernando X. Villarruel

Architect, Office of the CTO
Cisco Systems, Inc.
5030 Sugarloaf Pkwy, Lawrenceville GA, 30444
villarf@cisco.com

Michael Mobley

Technical Leader
Cisco Systems, Inc.
5030 Sugarloaf Pkwy, Lawrenceville GA, 30444
mobleym@cisco.com

Curt Dalton

Software Development Mngr.
Cisco Systems, Inc.
5030 Sugarloaf Pkwy, Lawrenceville GA, 30444
daltonc@cisco.com

Lamar West, Ph.D.

President
LEW Consulting, LLC.
7840 Holly Springs Road
Maysville, GA 30558
lamar@lamarwest.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Background	4
Node Cluster	5
Sensors and Measurement	6
Power Knobs for HFC RF Nodes	7
1. Example: RF FWD Amplifier	8
2. DAA Devices	9
Load Variations and Cluster Power Savings.....	9
Cloud-Based Adaptive Power Manager	15
1. General Assumptions	15
2. Protocol adapters	16
3. Device Abstraction Layer	16
4. Collector / Receiver Service	16
5. Analytics Engines	16
6. Data Base.....	16
7. Machine Learning.....	16
8. Policy Manager.....	17
9. Applications	17
10. OSS / BSS.....	17
Predictive Analytics	17
1. Heat Map Application	19
1.1. Data Storage	19
1.2. Data Analysis	19
2. Predictions.....	20
Power-Managed HFC Lifecycle	21
Recommendations	22
Abbreviations	22
References.....	23

List of Figures

Title	Page Number
Figure 1 - Node cluster definition. Power supply and its node parts.	6
Figure 2 - Power measurement and reporting of elements within a DAA node.....	7
Figure 3 - Power consumption knobs for an HFC RF node.....	8
Figure 4 - FWD amplification structure for HFC RF Node	9
Figure 5 - Four Optical Nodes Connected to a 90 Volt Line Power Supply.....	10
Figure 6 - Six Optical Nodes Connected to a 90 Volt Line Power Supply	12
Figure 7 - Power supply usage over load variance for cluster example	14

Figure 8 - Cloud-Based Adaptive Power Manager	15
Figure 9 - Hadoop analytics engine structure	18
Figure 10 - Distributed data storage of clusters.....	19
Figure 11 - MapReduce Functions for Heat Map Application.....	20
Figure 12 - Machine learning progression for adding nodes to cluster example.	21
Figure 13 - Function of an adaptive power manager throughout the HFC plant.	22

Introduction

Arguably the biggest challenge for the HFC plant moving forward is implementing next generation technologies in nodes that operate in a limited power consumption environment. This, for example, is the backdrop for the evolution to distributed access architecture (DAA,) full duplex (FDX), and networking nodes which are expected to require more average, and larger ranges of power consumption. There is also a discrepancy between the higher power/thermal dissipation capability of new nodes and the maximal power that is allotted to them per multiple system operator (MSO) product specifications. This creates an opportunity for an analytics-based application of power management that maximizes the performance capabilities of nodes but at the same time keeps in check or reduces the consumption of their system power footprint.

Because of MSO desires to avoid adding new power supplies when deploying additional nodes as part of, for example, fiber-deep deployment, nodes are now part of a power consumption cluster, which is effectively a collection of nodes serviced by a power supply where ultimately the power envelope that matters is that for the collection of nodes in the cluster, and not of any single node in particular. In effect, nodes are reverting back to a more centralized powering schemes from their previous distributed powering architectures, where the power supply placement is often no longer optimal due to new nodes being added downstream of the original node. Significant additional losses in power due to the Joule heating or I^2R losses in the coax used to transport power to the new node locations are now added to the powering requirements of the new nodes themselves. The new power consumption means that many power supplies may become challenged to supply sufficient power as new devices such as wireless strand-mounted devices are added to the HFC plant.

Therefore, power sensory information of node function is now necessary. Sensors that are hardware parts can be added to nodes to make the reading and reporting of power state information possible. This sensory information can be collected and maintained centrally, within a general cloud infrastructure. Making such energy consumption information available from sensors allows for an analytic comparison and optimization of power and/or performance settings of elements in the cluster to optimize performance while improving energy efficiency. And finally, in keeping with the current trend for increasing intelligence in network operations, these sensors and associated data can be the inputs to machine learning algorithms that provide predictions and necessary decisions to optimize or evolve a system and thereby facilitate introduction of new or different elements into the cluster.

This paper is a novel proposition for an analytics-based application that manages the problem of wide deployment of new technology in nodes. In this paper we describe the hardware, sensory capability expected, the cloud-based architecture needed for data collection, storage and analysis, the logic applied to analytical engines, and the process for execution of optimal states in the presence of a broader policy mechanism, and finally the inclusion of a machine learning principles for integration of new technologies as nodes evolve.

Background

The HFC plant finds itself in a precarious position when it comes to the topic of power consumption. On one hand there is the recognized desire to minimize as much as possible the power consumption of the plant. The effort of the SCTE/ISBE Energy Management Subcommittee has documented the tremendous cost burden that power consumption will have for MSO's in the near years to come. With this in mind there has been a rally to gather minds and technologies around energy conservation measures. Simultaneously however, the ever-increasing numbers of ports required for more granular service groups,

in DOCSIS and video, have led to the favoring of distributed architectures, where the PHY layer of the converged cable access platform (CCAP) is migrated to the outside plant. This evolution significantly increases the capacity of the fiber feeding the nodes by making the link digital, uses more efficient gallium nitride (GaN) technology that enables higher RF spectrum use, and can add to the intelligence of the node, but it can also dramatically add to the power consumption of these field-distributed endpoints. These trends are just at the beginning, where energy conservation measures will be increasingly important in the future as added capability of intelligent nodes is expected to be more prevalent, along with their need for more power to accomplish their functions.

Another dimension is the operational challenges with rearranging or upgrading the powering infrastructure of the plant. A wholesale revision would be prohibitively costly, and even impractical due to the nature of upgrading utilities. This dynamic leads operators to prescribe strict specifications for the power consumption of nodes because in the absence of any specific knowledge of the power state in a specific node situation or geography, it is the only guidance that can be given. For example, MSO's typically give a power consumption number for all nodes independent of function or location. As mentioned in the introduction, the deployment of fiber-deep architectures in particular with the constraint of not adding or moving power supplies often leads to new I²R losses in the HFC plant that make power consumption per node a more variable quantity.

In this paper we propose the question what if that were not the case? What if a system was not blind to the power state of its parts? What if guidance about power consumption could be given in broader and more granular terms, and not just for isolated nodes? If this were the case then optimization of power consumption could be done in the context of performance. It could vary from place to place, and it could be updated over time. We develop this proposition in the next sections of this paper.

Node Cluster

We define a node cluster as a group of functional nodes that share one power supply. This distribution of parts is generally referred to as the centralized power supply model. (Note, the principles we present would also work with alternate power supplies and node relationships, but we describe the solution here with the most likely scenario.) Figure 1 is visual representation of a node cluster. We note fundamentally that the possibilities for node parts can include various technologies. We list some of the ones we know now, with the understanding that in the future there could be others. We include legacy analog optic/RF nodes as DOCSIS3.1 or D3.1 in the figure. We include DAA nodes as remote PHY and FDX. We also include packet processing nodes such as field aggregation routers, optical transport nodes (OTN) multiplexing transponders or Muxponders and optical line termination units (OLTs).

These nodes within a cluster can require a range of power consumption profiles. Some could be less than the guidance typically given for power consumption and some could be more. The commonality for all nodes is their shared line power supply. For this reason, the best guidance with regards to power capabilities and limitations is given by the line power supply. This is really the only power envelope that matters.




Figure 1 - Node cluster definition. Power supply and its node parts.

While the statement of the power supply envelope is true, at the moment the relationship of the power supply to its subtending nodes is not known. The reason is that nodes at the moment do not measure (or are not known to measure) or have a mechanism to report their power consumption. Third generation power supplies on the other hand have a way to accurately report power state, but have no way to compare their capabilities to the context of the usage of their power-consuming devices in the node. Thus along with the concept of a cluster we propose the basic principle of nodes that are able to report their power state information.

Sensors and Measurement

The inclusion of power management integrated circuits (ICs) to a node design allow for measurement and reporting of electrical current and power consumption. These ICs are typically small and are cost and power-consumption effective. They have the task of measuring current drawn at a range of voltages and also of reporting current or power. Depending on the model, one can measure several lines by toggling, or have dedicated measurement. The reporting structure is typically facilitated by an integrated micro that allows for two way communication via a rudimentary form such as Intra-Integrated Circuit (I2C). The data set that includes power information is not meant to live in the node, so it would have to be stitched into the data plane for signaling that is already being transported to service packet cores. We expand on this in later sections.

Figure 2 below shows an example for power measurement and reporting of a remote PHY node via a power monitor chip. Note that this example tracks the usage throughout the power tree. The node power supply is monitored at entry and exit, and the granular parts of the node are measured as well. From a data collection perspective the input into the node power supply is the minimum required power measured, but the granularity added to measuring various points of the node allow for a broader application of both power consumption and performance tuning. In the next sections we will show that understanding the usage through out the power tree within the node is quite useful, in particular for balancing node component setting for energy consumption and performance targets. Note that Figure 2 is an example of an HFC Remote-PHY node, but a similar approach can be taken for analog, or packet processing nodes.

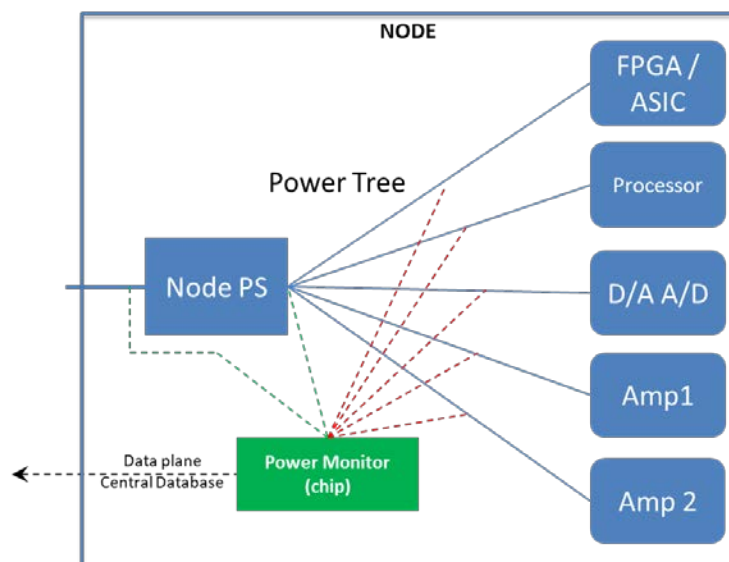


Figure 2 - Power measurement and reporting of elements within a DAA node.

Power Knobs for HFC RF Nodes

The power consumption of an HFC RF node can be influenced by various settings for the parts within the node. The settings are related to throughput and performance expectations of the node. However, it is important to understand that there can be multiple settings that achieve the same customer end-line performance or throughput from a node. Figure 3 shows the various levers or knobs we have available, for which through varying settings can affect the power consumption of the node. Unlike Figure 2 where we call out the node parts that can vary, in Figure 3 we focus on the signal settings and power supply settings of these internal devices that have power consumption effects. In this case it is not just a matter of power consumption but changes that can be done with the perspective of performance. We point out the modulation of the signaling, which with current and next generations of DOCSIS can vary greatly depending on network and end of line conditions. Varying the bandwidth is related to the settings the packet processor of a remote PHY device (RPD) for instance, as we see later on. Varying the RF bandwidth is related to the general settings for the plant of licensed spectrum being used. This of course can impact the necessary settings on both the RF management of the node and the packet processor. RF power refers to the absolute value of RF power being put out by amplifiers in the node design. Depending on the quality of signal needed at the end of line, (through RF drops, or related to the modulation setting, or the bandwidth allocation,) this value can change and can sometimes be maximized, while at other times not so much.




Figure 3 - Power consumption knobs for an HFC RF node

1. Example: RF FWD Amplifier

The forward path amplification structure of HFC nodes allows a good example to understand the dynamic of power consumption and performance. Figure 4 shows the typical forward (FWD) path amplification structure for a four port HFC node. There is a common pre-amplifier and four independent post amplifiers that provide RF output power for the four output ports of the node. The structure itself gives us the first type of power efficiency tactic. In the case where there is an inactive leg of the node, then the mere savings of turning off one post amplifier can be up to 20%. Similarly, if the accompanying RF structure allows for turning off of an amplifier during low traffic conditions (e.g. at night or in summer for university student residences) then those energy savings are also available for selected times of day/year.

The RF amplifiers are effectively the signal (and its noise) interface from the transmission network to customer premises equipment (CPEs), directly or via another set of amplifiers. The RF power value is tied to the needed carrier-to-noise ratio (CNR) in the RF domain and the related signal-to-noise ratio (SNR) after demodulation. Note that in this context, “noise” refers to the sum of thermal noise, interference, and composite intermodulation noise (CIN). Consequently, there is a limited RF output power dynamic range that provides an adequate carrier-to-noise ratio and the related signal-to-noise ratio. And as it turns out, signal-to-noise ratio or equivalently modulation error ratio (MER) for digital (QAM and OFDM) signals, which are used exclusively now, are tied to very particular expectations, where a signal-to-noise ratio that is too low is not usable, and if it is too high, it gives no extra benefit. This allows for the dialing in on a power range that is tied to a particular channel modulation, and by extension the number of channels that are used within a spectrum. As an example, we note that if an amplifier structure that is capable of signaling over the full RF range of DOCSIS 3.1 is operating at half capacity, either by reduction of RF spectrum or reduction of signal order of modulation (and thus lower required MER) it can then accommodate approximately 20% worth of power consumption savings. Note we do not specify any particular numbers as they would be product-specific, but the dynamic should be similar throughout the industry.

Another form of power reduction, a somewhat hot topic now is digital pre-distortion, (DPD), (Chong, 2018). The capability of DPD to help correct noise components from the driven profile of amplifiers is simplified in nodes with remote PHY devices which lend themselves to the digital signal processing (DSP) needed. Overall, accounting for the added DSP power, the saving of DPD can be up to 20% in very straight forward applications. There is potentially more added bonus, but this requires special calibrations and attention to details of loadings, etc.

Overall, it seems that 20% saving from the amplifier usage under reasonable circumstances is a very achievable target.




Figure 4 - FWD amplification structure for HFC RF Node

2. DAA Devices

A critical part to power consumption of new nodes is the silicon that is used for remote PHY or similar devices. These devices effectively take Ethernet / IP signals that were generated from a packet core and convert them to RF signals to be fed into the RF plant. At their core these devices have application specific integrated circuits (ASICs) or field programmable gate array (FPGA) technologies, or both which determine their functionality. Unfortunately, this vast amount of intelligence added to the node also comes with a power consumption penalty. This tradeoff is made with performance in mind as well, where ASICs are a naturally lower power solution but with little flexibility once finalized, and FPGAs have a higher power penalty with greater flexibility for evolution in the field. Because these devices are so power-hungry it is important to be able to understand and leverage any power saving capabilities. Without getting too into product detail, a mix of technologies, older and new versions, along with bandwidths, modulations and other ancillary functions dealing with versions of DOCSIS, the silicon within DAA devices can vary power by up to 50%. For this reason, the numerical understanding of power consumption (measured variations) for DAA devices is quite useful, if not necessary.

Load Variations and Cluster Power Savings

With the availability of varied power states in cluster nodes we turn our attention to how these variations materialize into savings to the overall cluster's power consumption. Connecting the cluster is a sequence of unique coaxial segments. Let's look at how the DC loop resistance of the coaxial cable affects the total power output supplied by the line power supply. This is a highly simplified analysis and is intended to illuminate the issue rather than give a rigorous final result. A rigorous analysis, such as that described in Mitchinson (Mitchson G., 2016) would be extremely complicated and is beyond the scope of what we wish to accomplish here. In order to make the current analysis practical we make the following simplifying assumptions:

1. The switching power supply in a node draws a constant power regardless of the input voltage. Actual switching power supplies vary slightly in efficiency as the input voltage varies. This results in a slight variation in power draw over the input voltage range. However, this variation is relatively small.

2. A group of four optical nodes are being fed from a single line power supply. These optical nodes are located relatively close to each other, yet all are a significant distance from the line power supply. For the purposes of this analysis we will simplify the things by ignoring the powering losses associated with the coax that interconnects the nodes and concentrate on the main coaxial span that connects the group of nodes to the line power supply. The effect of the interconnecting coax might slightly change the final value of the results. However, rigorous analysis of the node interconnections will complicate the mathematics well beyond what is required to demonstrate the nature of the relationships given below. This complexity of this analysis further illustrates the need for multiple individual power sensors located throughout the network in order to obtain empirical data.
3. We will assume that the rms voltage provided by the line power supply is the same as the peak voltage. This assumption is a reasonable one when the line power supply is under no load and the trapezoidal voltage output approaches a square wave. However, as the current from the line power supply increases its voltage waveform, it becomes more rounded, resulting in a decreased rms value. Were this change included in the analysis it would simply tend to reinforce the results that we will provide
4. We will not consider changes in the real power and apparent power as a result of the power factor of the switching power supplies in the node.

Consider the situation where the four optical nodes are connected to a 90-volt line power supply by a coaxial span with a DC loop resistance of 3 ohms, as shown in Figure 5.




Figure 5 - Four Optical Nodes Connected to a 90 Volt Line Power Supply

For the purposes of this analysis we will assume that each optical node draws 75 watts. The four optical nodes will draw a total of 300 watts. We can calculate the current draw, I from the line power supply and the voltage, V_L , across the four optical nodes.

$$300 \text{ watts} = (V_L)(I) \quad (1)$$

Rearranging

$$V_L = \frac{300 \text{ watts}}{I} \quad (2)$$

Considering the voltage drop across the DC loop resistance of the coaxial cable we can calculate the voltage across the optical nodes, V_L , by

$$90 \text{ volts} - (I)(3 \Omega) = V_L \quad (3)$$

Rearranging

$$(I)(3 \Omega) - 90 \text{ volts} + V_L = 0 \quad (4)$$

Substituting V_L using equation (X.2)

$$(I)(3 \Omega) - 90 \text{ volts} + \frac{300 \text{ watts}}{I} = 0 \quad (5)$$

Multiply through by I

$$(I^2)(3 \Omega) - (90 \text{ volts})(I) + 300 \text{ watts} = 0 \quad (6)$$

Solving for I using the quadratic equation

$$I = \frac{90 \pm \sqrt{((-90)^2 - (4)(3)(300))}}{(2)(3)} \text{ amps} \quad (7)$$

Resulting in the two solutions $I = 26.2 \text{ A}$ and $I = 3.8 \text{ A}$. We will ignore the larger solution as this is more current than can be typically supplied by a line power supply and chose the smaller solution. Plugging into equation (X.3) we find the voltage at the optical nodes to be

$$V_L = 90 \text{ volts} - (3.82 \text{ A})(3 \Omega) = 78.54 \text{ volts} \quad (8)$$

We can also calculate the total power being delivered by the line power supply, P_{LPS} ,

$$P_{LPS} = (90 \text{ volts})(3.82 \text{ A}) = 343.8 \text{ watts} \quad (9)$$

Note that 43.8 watts is being dissipated as heat by the 3 ohm DC loop resistance of the coaxial cable.

Now consider a second scenario. In this scenario node splitting has replaced the four optical nodes of the previous scenario with six optical nodes. Additionally, the six optical nodes contain advanced electronics such as DAA electronics and consequently each optical node dissipates 100 watts. This results in a total power utilization by the six optical nodes of 600 watts. This scenario is shown in Figure 6.




Figure 6 - Six Optical Nodes Connected to a 90 Volt Line Power Supply

It is possible to do an analysis of this scenario that is similar to the previous analysis.

$$600 \text{ watts} = (V_L)(I) \quad (10)$$

Rearranging

$$V_L = \frac{600 \text{ watts}}{I} \quad (11)$$

Considering the voltage drop across the DC loop resistance of the coaxial cable we can calculate the voltage across the optical nodes, V_L , by

$$90 \text{ volts} - (I)(3 \Omega) = V_L \quad (12)$$

Rearranging

$$(I)(3 \Omega) - 90 \text{ volts} + V_L = 0 \quad (13)$$

Substituting V_L using equation (X.2)

$$(I)(3 \Omega) - 90 \text{ volts} + \frac{600 \text{ watts}}{I} = 0 \quad (14)$$

Multiply through by I

$$(I^2)(3 \Omega) - (90 \text{ volts})(I) + 600 \text{ watts} = 0 \quad (15)$$

Solving for I using the quadratic equation

$$I = \frac{90 \pm \sqrt{((-90)^2 - (4)(3)(600))}}{(2)(3)} \text{ amps} \quad (16)$$

Resulting in the two solutions $I = 20 \text{ A}$ and $I = 10 \text{ A}$. We will ignore the larger solution as this is more current than can be typically supplied by a line power supply and chose the smaller solution. Plugging into equation (X.12) we find the voltage at the optical nodes to be

$$V_L = 90 \text{ volts} - (10 \text{ A})(3 \Omega) = 60 \text{ volts} \quad (17)$$

As compared to 78.54 volts in the previous scenario. The additional power loading has caused the voltage at the actives to drop by 18.54 volts. We can also calculate the total power being delivered by the line power supply, P_{LPS} ,

$$P_{LPS} = (90 \text{ volts})(10 \text{ A}) = 900 \text{ watts} \quad (18)$$

When we compare the second scenario to the first we see that the power required by the optical nodes increased by a factor of 2 in the second scenario with respect to the first. However, the power required from the line power supply went from 343.8 watts to 900 watts due to the additional coax I²R loss with the higher current. The power required from the line power supply increased by factor of 2.6.

This analysis is not intended to present the results for any given specific architectural case. Rather, it is intended to illustrate the nonlinear relationship between the power required by active devices in the network and the resulting power delivered by line power supplies.




Figure 7 - Power supply usage over load variance for cluster example

This simplified analysis illustrates a more general conclusion. Figure 7 shows the dynamic between the varying loads and power supply usage for the example topology. We note that the power supply output varies increasingly faster than the load variations, particularly as the load reaches higher values. Most operators would prefer to continue to use existing line power supply locations and existing powering architectures wherever possible. But as new system designs and technologies increase the total power utilization of active devices in the outside plant by a factor of two in our example, the power required from the line power supplies may increase in a non-linear manner by a factor that is greater than two due to the DC loop resistance of the coaxial cable supplying the power to those active devices. One slightly mitigating factor is the efficiency of line power supplies increases as the load increases. However, this may not be enough to offset the additional losses in the coaxial cable DC loop resistance.

It is important to note that the DC loop resistance will be unique for different node clusters. For this reason, it is not really practical to write completely effective generic rules for how to deal with the dynamic between nodes and power supplies in a deployment without carving out clusters as the quanta for system solutions. Beyond the identification of clusters however, generally stronger tools are needed to manage the complex nature element relationships. The next sections of the paper guide us through this process.

Cloud-Based Adaptive Power Manager

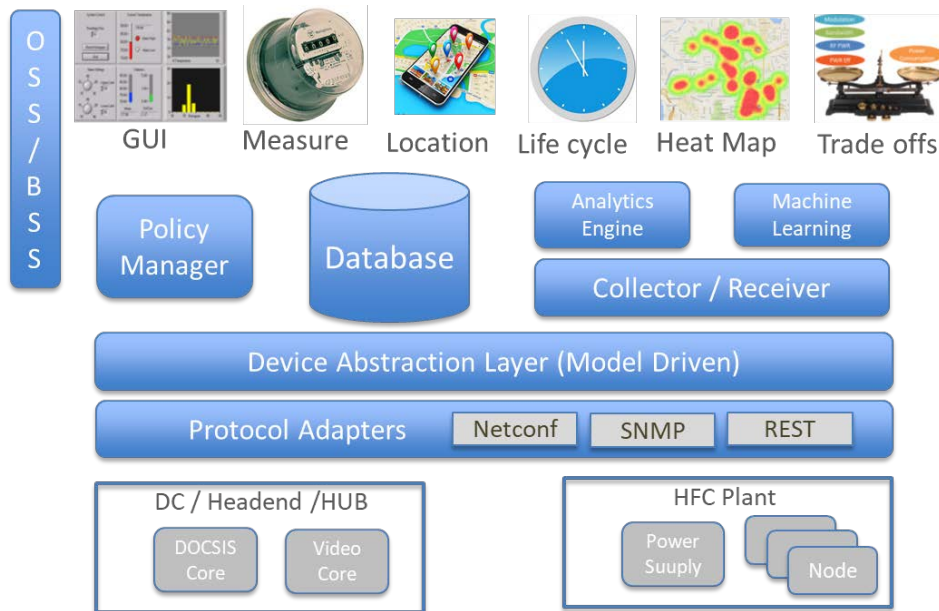


Figure 8 - Cloud-Based Adaptive Power Manager

The previous sections show that the power and performance dynamic of a cluster of nodes is complex enough that simple tools would not be able to extract its most optimal states. It is also the case that what is needed is some form of centralized view of the cluster, with knowledge of all parts of the system and ability to infer how their own sensitivities affect overall group targets, and more importantly, to be able to adjust power consumption of individual devices within the cluster when possible so that the current drawn by the entire cluster can be reduced and thus also the power consumed. Luckily this challenge comes to us at a time where the tools to create such a solution exist, and perhaps even with infrastructure already deployed, in support of other parts of the network. We then build a power manager in the context of the cloud, with otherwise generally understood architecture and approach, see Figure 8. Below we expand on this part of a cloud solution.

1. General Assumptions

Individual network elements must have power monitoring and reporting capabilities. Data may be communicated via a “push” model, whereby the network element sends telemetry data to a receiver in the management system, or a “pull” model, whereby the network element is polled periodically by a collector agent. An intelligent adaptive power management might consist of the following (or similar) components, deployed as micro-services in a cloud-native environment, as shown in Figure 8. Note, there have been discussions in this direction within the energy management community, but in this paper we aim to add more detail around the particular structure and function of supporting applications, (Sandoval, 2016)

2. Protocol adapters

It is expected that different network elements, from different vendors, would use a set of standard protocols to command and communicate power monitoring data and provide a method of control. Therefore, protocol adapters for protocols such as SNMP, REST, and NetConf would be required to communicate with network elements.

3. Device Abstraction Layer

A model-driven device abstraction layer would map device-specific application programmer interfaces (APIs) to a standard set of APIs used within the management system. This is where we assume the Yang model and/or management information bytes (MIBs) for various devices would be translated into generic descriptions. This effort is of course facilitated by the work of the SCTE sponsored APSIS specifications, where robust models have already been defined, or can be further refined. (SCTE, 2018)

4. Collector / Receiver Service

This service or group of services would receive telemetry from network elements sophisticated enough to push power monitoring data up to the management system. For elements that only support a pull model of data retrieval, the collector would periodically poll the desired data from the device. In the illustration, the collector / receiver service collects data from a power supply and distinct types of nodes in an HFC cluster, as well as data from various flavors of associated CMTS and quadrature amplitude modulation (QAM) devices in the cable headend or data center. The collector writes the raw data to the database as-is, with little or no processing of the data itself.

5. Analytics Engines

The analytics engine refers to environment that is particularly structured to handle the scale and workload generated by analyzing big data. If broadly deployed the data that is coming from node clusters could become very large very quickly and the analytics engine would provide a set of tools to organize the nodes cluster information read by the plant and allow us to do relational calculations. We expand on this in the next section.

6. Data Base

The data base is formally a part of the analytics engine as it is the analytics engine that would determine the format for storage distribution and organization of data. Because this is a “big data” exercise there are several qualities to the storage of data that are necessary, like the ability to process in parallel, allow for low commodity hardware and interact with a robust resource manager. We expand on this in the next section.

7. Machine Learning

The machine learning module is really part of the analytics function but we call it out separately. It takes in sensor data, the capabilities of the equipment and statistically tracks both state data from cluster components and makes predictions and recommendations from applications according to programmed rules. In time this training data will allow for the creation of other algorithms such as those in Section 9 below, whose relationships would otherwise be too obscure to conjure up directly, to help maximize the relationship of power consumption and power performance of the cluster. This is exactly where machine

learning can be most useful, and additionally, part of the vision for this module is to overcome the inevitable problem of fully calibrating parts in design and manufacturing before adding to the cluster, both for varied descriptions of power and performance and over time. Machine learning algorithms are often ideal for making predictions and recommendations from incomplete or slightly inaccurate data.

8. Policy Manager

The reduction of power, or the maximizing of performance via the execution of applications just described needs a higher perspective to serve the interests of the wholistic system. These decisions might have to be done in a case by case basis, for example in conjunction with the service level agreement (SLA) of a particular set, or individual end-line customers within the cluster. The policy manager is in charge of executing this higher perspective and must have a view of the other components in the system and their priorities, and this would include other packet cores like the CMTS.

9. Applications

There are a number of possible top line applications. These applications would access relevant information stored in the database to perform their respective functions. Perhaps the most basic function would be to provide a graphical user interface for the rest of the system to operators in the back-office. Another application could take the actual responsibility of measuring the system and providing energy consumption data at a granular and summary level to corporate sustainability offices. Another application could provide the physical location of clusters and another application provide a heat map for power efficiency improvement opportunities. Another application could provide a description of status within the life cycle of products in a cluster, making recommendations on how to optimize upcoming plant design over time, or even do so in modeling before a section of the plant is built. Another application could provide a real time savings calculation in terms of watts or dollars for the cluster—not shown in the picture. We do expect there would be an application that could adaptively modify the element states toward energy savings. There would be an application that makes suggestions on settings of nodes according to their respective power supply capability, towards overall power reduction or performance improvement. Finally, there could be apps that go beyond the settings of the node and extend to plant and network configurations via comparison between clusters. In general, the applications space is an open canvas for creativity and multi-vendor differentiation.

10. OSS / BSS

The operation support system (OSS) and billing support systems (BSS) are the executive coordinators of the whole operation. They determine policy and the ultimate experience of the end line customer. Ultimately, the work of the power management environment must work within the confines set by the OSS and BSS. This is where the policy manager plays a crucial role making sure that any changes in the system do not violate higher directives.

Predictive Analytics

Predictive analytics (PA) has two main targets, the first is to take large sets of data and process it into meaningful relational information and the second is to learn from the experience of that data and anticipate meaningful directions that allow for positive action on the system. In our case the big data set is the power state information for clusters and their components in an HFC plant. The data could be as small as one node cluster, or as large as all the clusters in the national footprint of an MSO. In our case the relational information would be between the node components in a cluster and its line power supply, or

between clusters, or comparisons of particular type of nodes in all clusters. There is no limit to what sort of relations we can study once the data is present. For the anticipation part of PA we could look to avoid failures due to future power inefficiencies and take action to prevent them, or avoid any performance shortcoming when there is power available to make it better. There is also no limit to the actions taken from anticipation learned. Some of the actions could include recommendations for the challenges pointed out in the “Use Cases for Adaptive Power Using APSIS”, (SCTE245, 2018)

While a full treatment of PA would be very extensive, for our particular goal of highlighting the general principles within the context of HFC power management, we focus on three main functions: data storage structure, analysis of data for execution, and prediction of future situations via machine learning. We do this by presenting a practical example. While there are various frameworks to conduct this exercise, we rely here on the infrastructure of an open-source predictive analytics engine called Apache Hadoop. Hadoop stems from a project by Google to organize storage and do computation on big data. It is a now a common tool. (Note that in the case of HFC node clusters, our data analysis does not have to be real time so Hadoop is a good fit.) The framework for Hadoop is shown in Figure 9.

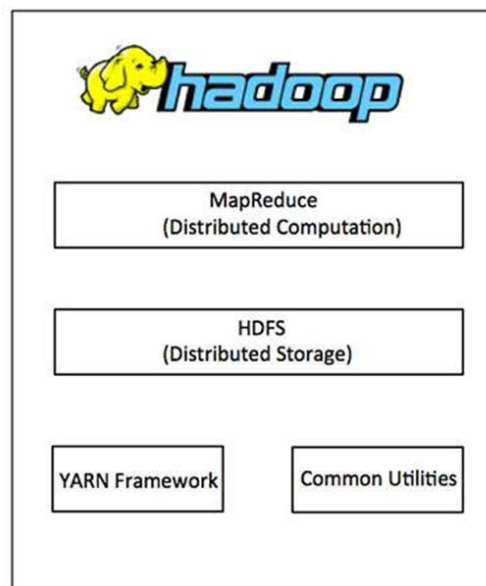


Figure 9 - Hadoop analytics engine structure

The first basic component of this analytics engine is its storage framework, the Hadoop Distributed File Storage (HDFS), (EdurekaHDFS, 2017). As the name implies, when data is stored it is done so in a distributed way, with a management layer or descriptive metadata and separate data clusters that are generic enough that can run on commodity hardware. A more concrete example will be shown below.

The second component is a sequence of algorithms that do analysis on data sets by distributing computation and coalescing results, these are the MapReduce functions (EdurekaMapReduce, 2017). In MapReduce one can execute known or program custom function on data sets. A more concrete example will be shown below.

The third component is Yet Another Resource Negotiator (YARN) which is an operating system of sorts for the architecture of data storage and its usage, (EdurekaYARN, 2017). This resource manager

organizes the data clusters for storage and makes available data sets for computation and parallel processing, along with its scheduling.

There are also other common utilities that are present to facilitate the work of HDFS and MapReduce and while very helpful in practice we do not cover them here.

1. Heat Map Application

As an example, we sketch the heat map application on the PA platform. The heat map application would show the state of interaction between a line power supply and its subtending nodes in and HFC node cluster. This effectively shows whether a group of nodes is within the power envelope of its main supply or not.

1.1. Data Storage

Our first task is to take and organize and store our large data set from the field. In Figure 10 we show this data organization as would be done in HDFS. We note that there are independent raw data clusters and a higher level of metadata, which has description for the raw data clusters. These raw data clusters have redundancy and can be executed on commodity hardware. The management and organization of these raw data cluster sets is done by YARN. In our particular example the raw data clusters have the power state information for a power supply and its subtending nodes. Interestingly, we can organize the data clusters to coincide with HFC node clusters somewhat simplifying the steps that follow in MapReduce. The metadata in this example then just keeps a list of the types of components in the clusters, in our case a list of the power supplies and nodes is available.

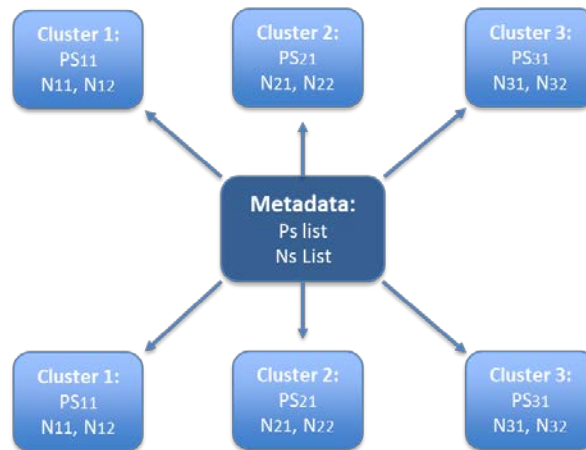


Figure 10 - Distributed data storage of clusters

1.2. Data Analysis

We show our data analysis on a framework that represents MapReduce. Our goal is to compute the comparison of the power available by a line power supply and the composite power usage of its subtended nodes.

In Figure 11 we show the progression for analysis. First is the identification of a data set as input. Note this data set is identified but not retrieved. In our case we need a list that includes power supply and nodes for HFC clusters 1-6. Then this data is organized (split) into usable groupings for distributed computation.

In our case the grouping matches the HFC node cluster for simplicity, so each node cluster's information is a set, line power supply for its two nodes, for six HFC node clusters. Next is the mapping itself, which is the execution of functions in a parallel manner on the data sets identified. In this case it is a logic identifier between the power supply capability and the addition of power consumption for the two nodes it serves. If the line power supply capability is more than its nodes usage then the value "0" is returned. If the line power supply capability is equal to the nodes usage then the value "1" is returned, and if the line power supply capability is less than the nodes usage the value "2" is returned. This is done in parallel for all six groups. The shuffling exercise is the organization of results from computations done in parallel. In our case it coalesces the power supplies that are well within their operation capacity, the power supplies which operate on the very edge and those that are being overworked or beyond their means. The reducing function then gives the actionable data we seek. In our case which power supplies we should monitor and which power supplies need immediate attention. This data is then exposed by the Heat Map application in a graphical representation of colors, yellow clusters if they need monitoring and red if they need attention.



Figure 11 - MapReduce Functions for Heat Map Application

This albeit simple example shows the thought progression behind using analytics for actionable intelligence. With the data at hand it is just a matter of creativity for how to use it.

2. Predictions

The predictive part of predictive analytics comes for the application of machine learning. Formally machine learning is a subset of the broader field of artificial intelligence which allows for some newer technologies like self-driving cars. In our case we prefer it because it allows us to apply statistical methods on data to learn from experience and predict the future behavior of a system, the system in our case being a node cluster.

Below we provide a simple example of using machine learning progression to predict what will happen to a cluster's line power supply output when two new nodes are added, as could be the case for an HFC node cluster when the HFC plant evolves to DAA, for instance.

In Figure 12 three plots show the progression of a very simple machine learning sequence. The very top diagram is borrowed from Figure 6 which shows physical layout of a node cluster with a line power supply, unique DC loop resistance, and six subtending nodes with particular power usage load.

The left graph shows a scatter plot of data learned by the system. In our case this is the power relationship between main supply and composite load. This data of course is available because of the analytics framework we have built already. The relationship mentioned above for instance would be built from a MapReduce module made to return corresponding line supply power and additive load.

The middle graph shows a regression for the data learned; this step is the application of some statistical tool for the data available. In our case we use a straight forward linear regression and extend its outcome beyond the data available. This is our prediction model then because as we look to add two nodes we can estimate the behavior of the system as two more nodes are added, as represented by the red circles on the line. We can also estimate the range of possibilities for these points, per the confidence level of the regression, thus the error bars on the points. Note that there are various statistical tools, beyond linear regression, available depending on the machine learning package being used, but generally with some mathematical treatment many relationships can be reduced to approximate linear forms. Thus simple linear regression can be quite common.

Finally, our last graph plots the actual relationship of two real nodes added to the system. Now it is time to test the model proposed by our machine learning algorithm. We see that in comparison to the regressed line the two real new points are not that far off, but we also note that the regressed line is not perfect, and so we can go through the process again and again until the user is satisfied with the fidelity of the line to the system. Once there is confidence in the model it can be used as actionable intelligence. In our case we can use the model to predict the behavior for other clusters of similar topology that need to add more nodes. We can use it to model a system before deployment. We can even use it to find and build optimal topologies for new builds. The possibilities are many. Note that the fine-tuning capability of these predictive techniques is what makes them ideal tools for node clusters because as we saw earlier the energy consumption signatures will be unique for different node clusters.

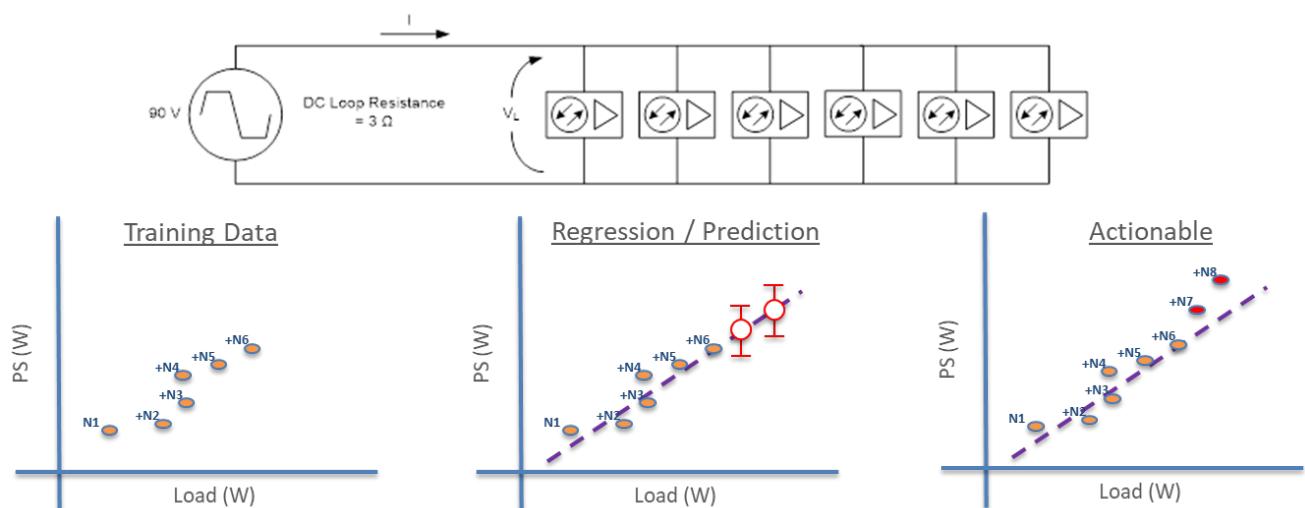


Figure 12 - Machine learning progression for adding nodes to cluster example.

Power-Managed HFC Lifecycle

The interesting takeaway about having a tool that can give insight into the power relationship of the HFC plant is that we can use it throughout the plant's lifecycle, see Figure 13. It can be used for modeling capacities and topologies even before it is built because we can anticipate outcomes with high confidence. It can be used for set up and create birth certificates for plants in deployment. It can be used to monitor and react in the day-to-day operations. It can help to manage service windows for the plant, it can help predict and identify end of life to power sensitive products. It can change the perspective of nodes and components according to expected power consumption, and can use items that are best in a system, not necessarily just efficient at one setting on its own. And it can also anticipate and assist in the sourcing of

replacements parts. Overall it can have a very positive impact on the energy consumption and performance of HFC plant moving forward.




Figure 13 - Function of an adaptive power manager throughout the HFC plant.

Recommendations

The deployment and use of an adaptive power manager can be facilitated and cause several changes in the way the industry would approach the opportunity for HFC plant energy conservation measures and their relationship to network performance. Below we list several recommendations for the MSO community to take in this direction.

- Think of power in manageable chunks: we called them node clusters.
- Create robust and open power data models for all legacy and new nodes technologies.
- Centralize power management via cloud infrastructure.
- Allow predictive analytics to do the heavy lifting.
- Specify node power limits in context of cluster capabilities, not individual targets based on other averages.
- Qualify new node product in context of performance in clusters over time, not solely on individual performance.

Some form of power management is likely inevitable due to the otherwise large energy consumption and performance problems its absence would entail. The above recommendations would facilitate such an implementation and create an interoperable environment ripe for its fast development.

Abbreviations

ASIC	Application-Specific Integrated Circuit
BSS	Billing Systems Support
CIN	Carrier to Composite Intermodulation Noise
CMTS	Cable Modem Termination System

CNR	Carrier to Noise Ratio
CPE	Customer Premise Equipment
DAA	Distributed Access Architectures
DC	Direct Current
DOCSIS	Data Over Cable Service Interface Specification
DPD	Digital Pre-Distortion
FDX	Full Duplex DOCSIS
FPGA	Field Programmable Gate Array
HDFS	Hadoop Data File System
HFC	Hybrid Fiber Coaxial
I2C	Inter-Integrated Circuit
IC	Integrated Circuit
ISBE	International Society of Broadband Experts
MSO	Multiple System Operator
OLT	Optical Line Termination
OSS	Operations Systems Support
OTN	Optical Transport Network
PA	Predictive Analytics
PHY	Physical Layer
QAM	Quadrature Amplitude Modulation
REST	Representational State Transfer
RF	Radio Frequency
rms	Root Mean Square
SCTE	Society Of Cable Television Engineers
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SNR	Signal to Noise Ratio
YARN	Yet Another Resource Negotiator

References

- Chong, K. (2018, July 6). *Director of Product for Quorvo*. Retrieved from Broadband Tech Report: <https://www.broadbandtechreport.com/articles/2018/07/growing-the-hfc-pipe-to-10-gbps.html>
- EdurekaHDFS. (2017, July 17). *HDFS Tutorial: Introduction to HDFS & its Features*. Retrieved from Edureka.com: <https://www.edureka.co/blog/hdfs-tutorial>
- EdurekaMapReduce. (2017, July 17). *MapReduce Tutorial – Fundamentals of MapReduce with MapReduce Example*. Retrieved from Edureka.com: <https://www.edureka.co/blog/mapreduce-tutorial/>

- EdurekaYARN. (2017, July 17). *Hadoop YARN Tutorial – Learn the Fundamentals of YARN Architecture*. Retrieved from Edureka.com: <https://www.edureka.co/blog/hadoop-yarn-tutorial/>
- K. Sundaresan, e. a. (2016). Applications of Machine Learning in Cable Access Networks. *Spring Technical Forum*.
- Mitchson G., H. D. (2016). Measuring and baselining power consumption in outside plant equipment and power supplies. *SCTE Cable-Tec Expo*. www.SCTE.org.
- Sandoval, F. (2016). *APSYS and Open Daylight*. Retrieved from https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=8&ved=0ahUKEwjI89qFrancAhWBc98KHT11BqEQFghSMac&url=https%3A%2F%2Fwiki.opendaylight.org%2Fimages%2F%2Fc4%2FAPSYS_and_OpenDaylight_-_light.pptx&usq=AOvVaw1sHqHxs0z4inNNYefRL8ZK
- SCTE. (2018). *SCTE Energy Management Program*. Retrieved from SCTE: https://www.scte.org/SCTE/Areas_of_Interest/Energy_Management/SCTE/Areas_of_Interest/SCTE_Energy_Management_Program.aspx?hkey=fa20e5a1-38c1-444e-84ee-e5f35541d6bd
- SCTE245. (2018). *SCTE 245 2018, Use Cases for Adaptive Power Using APSIS*. Retrieved from www.SCTE.org: <https://www.scte.org/SCTEDocs/Standards/SCTE%20245%202018.pdf>

Analysis And Prediction Of Peak Data Rates Through DOCSIS Cores

A Technical Paper prepared for SCTE•ISBE by

John Holobinko

Director, Cable Access Business Strategy
Cisco Systems, Inc.
Fort Mill, SC
jholobin@comcast.net

Glenn McGilvray

Manager, Product Marketing
Cisco Systems, Inc.
San Jose, CA
mcgilvg@cisco.com

John Ritchie

Principal Engineer, Engineering
Cisco Systems, Inc.
Lawrenceville, GA
ritchij@cisco.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction	3
1. Traditional DOCSIS Peak Capacity Models.....	3
2. The Evolving Nature of DOCSIS Traffic	3
3. Breakdown of the traditional bandwidth models.....	4
4. CCAP Congestion Points	4
5. Service Group Level Congestion.....	5
5.1. CCAP Platform Line Card Congestion.....	5
5.2. CCAP Core Congestion	6
6. An Alternative Capacity Planning Model.....	6
6.1. Core Peak Bandwidth Data Analysis	6
6.2. Sample Universe	7
7. More Observations On The Nature of CCAP Core Traffic.....	10
7.1. Peak to Valley Traffic Analysis	10
8. Further work.....	13
9. Conclusions	13
Bibliography.....	14

List of Figures

Figure	Page Number
Figure 1 – Potential Congestion Points.....	5
Figure 2 – Peak Data Rate per CCAP Platform.....	7
Figure 3 - Peak Data Only Rate per CCAP Platform by Month	8
Figure 4 - Peak Data Rate per Subscriber by CCAP Platform – Data Only	9
Figure 5 – Predicted Peak Data Demand Based on Subscribers and Growth Rate	9
Figure 6 – Peak to Minimum Traffic Ratio Across CCAP Platforms	10
Figure 7 - Single CCAP Platform Traffic for One Month Taken On 5 Minute Intervals.....	10
Figure 8 – Single CCAP Platform Traffic for One Day Taken On 5 Minute Intervals	11
Figure 9 – Proposed Power Savings By Turning Off Carriers During Low Demand Periods.....	11
Figure 10 – Lowering QAM Modulation Profiles During Off Peak Periods	12
Figure 11 – A Method For Dynamic Power Savings Using Analytics and Smart Nodes	13

Introduction

For nearly as long as DOCSIS has been deployed in cable networks, capacity planning has been a challenge for operators. As the nature of traffic has evolved across the IP network, the accuracy and usefulness of traditional DOCSIS capacity modelling techniques has diminished.

This paper describes an alternative, simpler model for predicting the peak capacity of a DOCSIS network, as measured bidirectionally through the CCAP (converged cable access platform) core, solely based on the number of subscribers attached to the DOCSIS platform and the demographics of the customers. This model is based on analysis of current traffic patterns and traffic types and subscriber behaviors within DOCSIS networks, which is presented herein. Core peak data was sampled over multiple months and across dozens of systems, and is used to show that the accuracy of the technique is superior to other modelling methods.

1. Traditional DOCSIS Peak Capacity Models

Historically, primarily two models have been used to calculate required DOCSIS core network capacity, both dependent on an assumption for the value of statistical multiplexing. Each method was developed based on the results from empirical measurements. Some operators have used a combination of the two models. Both are limited in their efficacy by the rapidly changing nature of DOCSIS traffic, and therefore misunderstood.

The first model employs a multiplier times the maximum data service speed offered to compute the total capacity required for a service group, then multiplies this times the number of service groups, and finally divides this total by an estimated statistical multiplexing conversion factor.

The second model is bandwidth oversubscription. In the oversubscription model, the maximum data speed of each subscriber within the service group is summed together, then divided by a number Y, called the oversubscription ratio. Oversubscription is an expression of statistical multiplexing. The oversubscription ratio Y:1, depends on the number of subscribers per service group, and more importantly on the types of services that are being offered.

Neither of these models is particularly accurate for predicting the peak traffic rate of a specific CCAP platform.

2. The Evolving Nature of DOCSIS Traffic

From the time DOCSIS was implemented and for the following fifteen years, web browsing and file transfers made up the majority of internet traffic. These data flows are characterized by their short, bursty natures. They do not consistently use bandwidth over a long period of time, i.e. minutes to hours. Based on this burstiness, when looking at aggregate data rates across a network, the aggregate data rate was but a fraction of the total of the data rates offered to all subscribers. This enabled a high amount of statistical multiplexing, i.e. “overselling” the same bandwidth multiple times over. For example if there were 500 subscribers, each with a 10 Mbps (megabits per second) service speed, the computed maximum bandwidth would be 5 Gbps (gigabits per second). However, based on the bursty nature of the traffic, operators found that oversubscription ratios of 50:1 (i.e. 100 Mbps vs. 5 Gbps) were not overly aggressive and provided acceptable performance to this number of data subscribers.

With the advent of Netflix, YouTube and other OTT (over the top) video content providers, over the last 5-7 years the nature of IP traffic has changed immensely. According to the September 2017 Cisco Visual Networking Index¹, globally, IP video traffic will be 82 percent of all consumer Internet traffic by 2021, up from 73 percent in 2016. Global IP video traffic will grow threefold from 2016 to 2021, a CAGR (composite annual growth rate) of 26 percent. Internet video traffic will grow fourfold from 2016 to 2021, a CAGR of 31 percent. Video traffic is very different than file transfers and web browsing in that it creates a highly persistent traffic flow, versus an intermittent bursty traffic flow. Even with adaptive bit rates and modern compression that utilizes chunking and large data buffers, video traffic is highly persistent compared to these other traffic types. This means that the “peak” data rate is sustained over long periods of time, i.e. minutes to hours versus seconds. With the ratio of video to overall internet traffic continuing to expand each year, data persistency is only going to increase. This leaves far less benefit to utilizing statistical multiplexing as a means of bandwidth oversubscription.

Statistical oversubscription calculations depend on the bursty nature of internet traffic. As this burstiness continues to be mitigated by persistent data streams, the accuracy of these calculations decline. Whereas 50:1 and even higher oversubscription ratios were used years ago, every operator continues to gradually reduce their oversubscription ratios each year to a fraction of their former values.

3. Breakdown of the traditional bandwidth models

To use maximum data service speed as a means to compute required capacity, a normal rule of thumb is to take a service group of “N” subscribers. The traditional formula is to take X times the maximum downstream data speed offered where X is usually between 1.5 and 2.0, and use this to compute the number of downstream DOCSIS channels required to support this computed bandwidth. The goal is to have sufficient bandwidth to support at least one user in a service group performing a downstream burst speed test while simultaneously supporting other users’ normal data consumption. The required core processing capacity is computed by multiplying this number (the computed bandwidth for each service group) times the number of service groups supported by the CCAP core. The challenge with this formula is that the statistical peak traffic through any service group is dependent upon the number of subscribers in the service group. So, in a large service group with many subscribers, the peak to average will likely be smaller than in a small service group. Secondly, as service groups are summed across a given CCAP core, all peaks will not occur simultaneously. Therefore, computing core capacity using this technique will result in a highly over-engineered core network, while not guarantying adequate bandwidth for each service group.

As maximum data speeds have increased and the nature of IP traffic has changed, the accuracy of previously referenced techniques to compute required network capacity has declined. For example, as data speeds have grown to 1.0Gb/s and even greater, the multiplier “Y” times maximum data speed might remain sufficient for determining the minimum bandwidth required for service group, but it has become a poor method for computing overall required CCAP platform bandwidth, significantly overestimating required bandwidth through the core. Therefore it is instructive to look at the various points of congestion that can occur in a CCAP system.

4. CCAP Congestion Points

In any network, congestion occurs when the total traffic demand exceeds the peak capacity. Any such location is referred to as a bottleneck. In CCAP systems, congestion can occur at the service group level, the line card level, in the switching fabric, in the core, or at the network ingress/egress point. These points are shown in Figure 1.

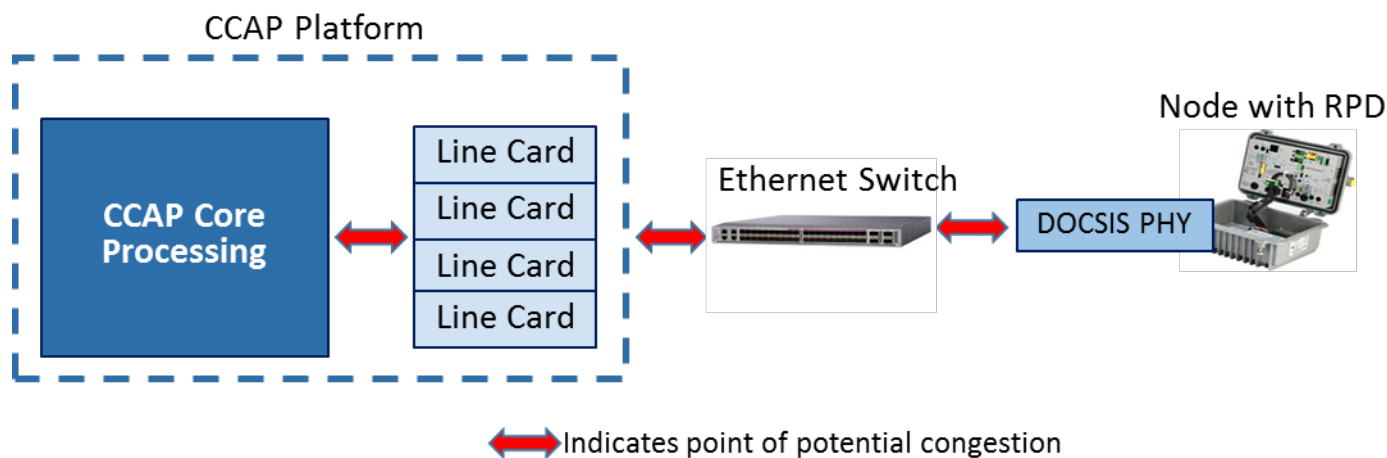


Figure 1 – Potential Congestion Points

5. Service Group Level Congestion

When congestion occurs at the service group level, it can either impact the throughput of all users, or the maximum data speed of the highest speed service tier users. These are two distinctively different forms of congestion. If throughput is impacted across all users, it means that during peak usage periods, the total demand for bandwidth exceeds the bandwidth capacity of the service group. This type of congestion is based on the total number of active users and the type of content they are using. For example, watching high quality videos creates long term persistency of streams, with far less benefit from statistical multiplexing. This congestive impact tends to be over longer periods of time (minutes or hours versus seconds). In contrast, one or more very high speed users may perform speed tests at the exact same time. If insufficient capacity exists, they will not be able to attain their maximum service speed. However, overall network response may not be noticeable to other users depending on the overall QoS settings.

Therefore, when provisioning capacity to a service group, the design goal should be to provide sufficient capacity (i.e. DOCSIS spectrum/bandwidth) to minimize congestion during peak usage periods, while insuring that there is also sufficient bandwidth to support X times (typically $X=1.5$) the maximum speed of the highest tier of service. Depending on the number of subscribers per service group and their relative internet usage, either the first calculation (large number of subscribers) or the second calculation (fewer subscribers, very high top data speed) will dictate the required downstream capacity.

5.1. CCAP Platform Line Card Congestion

Virtually all dedicated CCAP platforms use one or a combination of an ASIC (application specific integrated circuit) or FPGA (field programmable gate array) to perform framing, scheduling and other functions which are shared across a number of service groups. Typically these functions are done at a line card level or equivalent. At this level, the CCAP architecture should provide sufficient capacity (i.e. DOCSIS spectrum/bandwidth) to minimize congestion during peak usage periods. Since line card speeds are typically 50-100 Gbps, one is not concerned about the peak data speed. Rather the limitation becomes the total number of subscribers, i.e. number of subscribers per service group times the total number of service groups on the line card. In contrast to the service group, the design goal is not to provide sufficient bandwidth, it is to ensure that the maximum bandwidth of the card is sufficient to support the combined peak bandwidth demand for the subscribers connected to the card.

5.2. CCAP Core Congestion

The CCAP core is the point where all traffic passes through the platform in both directions and certain functions such as Quality of Service (QoS) are performed. The goal of the core is to be non-blocking, i.e. support the maximum peak traffic defined by the service levels and QoS. Since CCAP platforms are all architected to last for many years, today the only way to create congestion at the CCAP core is to provision an exceptionally large number of subscribers, or to provide insufficient network side optical bandwidth to the network core.

6. An Alternative Capacity Planning Model

A 2014 study by Princeton University¹ showed that if cable subscriber currently subscribed service speeds were adequate for daily use (i.e. >25 Mb/s), and these speeds were increased far beyond the service level they were paying for (e.g. to 250 Mb/s) without the subscribers' knowledge, overall data consumption increased by less than 5%. Some subscriber internet behaviors changed, but overall consumption did not. If one started with an assumption that consumption is directly related to data speed, the expected traffic growth should have been 900% instead of 5%. Similar patterns of behavior were noted in Asia when 1Gb/s data speeds were first introduced, i.e. in Asia, operators noticed an initial spike in data traffic by new users for the first month, then a gradual decline until consumption was almost identical to the prior rate of consumption prior to the introduction 1Gb/s data speed subscription service. Other cable operators in North America have corroborated this phenomenon.

We therefore started with the following premises:

- As an alternative bandwidth model, the peak data rate through a CCAP core can be predicted simply by knowing the total number of subscribers attached to the core and the annual data growth rate for the area in which the platform is located
- As long as congestion at the network edge is not affected, total traffic through the core will not be significantly affected by service group size

These are key premises, because they turn upside down the normal way of thinking about capacity planning: Service groups, data rates, oversubscription, etc.

6.1. Core Peak Bandwidth Data Analysis

In addition to Cisco Systems' annual report on growth of traffic across the internet, multiple large cable MSO's publish or otherwise report their year over year growth projections for downstream and upstream DOCSIS traffic across their footprints.

We set about to measure the peak bandwidth consumption across multiple CCAP cores installed in select cable operators with varying numbers of subscribers and service groups on five minute intervals to determine network behaviors and to see if there is another way to predict aggregate bandwidth through the CCAP core. Our goal was answer the following questions:

- Is core CCAP capacity consumption independent of service groups?
- Given a geographic region, can core CCAP capacity be predicted simply by knowing how many subscribers are connected to the core?
- What is the differential between peak traffic times and off peak times? Can this be used to advantage in other areas of cable network operation?

6.2. Sample Universe

Our total sample size covered tens of CCAP platforms serving multiple hundreds of thousands of DOCSIS subscribers. Taking samples of the core data rate on 5 minute intervals, the total number of samples was in the hundreds of thousands. The total sample period was seven months. Therefore, we feel very confident that the statistical sample size is well beyond the minimum sample size required to achieve confidence in the results.

So as not to reveal any cable operator sensitive data, we will generally summarize the data sets as follows:

- The total sample size consisted of approximately 40 CCAP platforms. Some had converged MPEG video and data/voice services, while others were data/voice only.
- Demographics ranged from dense urban to suburban to rural areas.
- The number of subscribers per platform varied by over 500% with the largest number being approximately 12,000 subscribers. The total number of subscribers across all sampled systems was approximately 300,000.
- The number of service groups per CCAP platform varied by nearly 500%, i.e. the fullest platforms supported four times the number of service groups as the number of service groups the least utilized platforms supported.
- The number of subscribers per service group also varied by over 400% as well with the largest service groups having more than 300 subscribers attached.
- On each CCAP platform core, a sample of the total combined forward and return path combined throughput was taken on five minute intervals. The total sample period for the study was seven months. The total samples for the study approached one half million.

Figure 2 shows the peak data rate through CCAP platform cores over the sample period. If the peak bandwidth were proportional to the number of DOCSIS channels provisioned per service group, the most variation we would expect is 4:1, based on the smallest platforms having one quarter of the service groups as the largest platforms. But instead, we see a ratio of more than 6:1.




Figure 2 – Peak Data Rate per CCAP Platform

This can neither be explained away by converged video on only some of the platforms. Figure 3 shows the same data with video subtracted from calculations; only the data traffic on each platform is presented,

whether or not MPEG video/data convergence was present on a platform. The ratio of the highest peak traffic platforms to the lowest peak data platforms remains the same.




Figure 3 - Peak Data Only Rate per CCAP Platform by Month

However, the data becomes much closer if it is normalized to a peak average per subscriber rather than the individual CCAP system peak throughput itself. This number represents the average bandwidth per subscriber at the point where the platform reaches its absolute maximum throughput (i.e. customer demand) in any given month. This peak average per subscriber is calculated for each platform by taking the monthly peak traffic sample and dividing it by the number of subscribers.

Avg Peak Traffic Rate per Subscriber (APS) = Platform Peak Traffic Rate/Subscribers on Platform

Figure 4 shows the resulting graph comparing the APS across all of the platforms, (no video). The variance across all of the machines is now far closer, 2:1. The highest APS is approximately 2.2 Mbps while the lowest is approximately 1.1 Mbps. (There were a small number of systems that were activated during the study. These are shown with the very high growth rates). Looking more closely at the data reveals that the systems with lower APS tended to be in more rural areas, while more dense metropolitan areas generated higher traffic demand. What the data shows is that normalizing the peak data on a per subscriber basis is a far more accurate predictor of total peak bandwidth through a CCAP platform than using other modelling techniques such as bandwidth per service group along with a multiplexing factor, or total provisioned bandwidth across the platform, number of service groups, etc.




Figure 4 - Peak Data Rate per Subscriber by CCAP Platform – Data Only

The value of this is in the ability to predict CCAP platform total peak bandwidth in the future. The data shows that by knowing the number of subscribers attached to a given platform and whether it is rural or urban, one can make a good estimate of future bandwidth requirements. For example, in an urban environment, we would predict that subscriber APS is conservatively 2.5 Mbps or less as of the date of this paper (recognizing that data grows annually, thus affecting the baseline). Figure 5 below shows graphs based on 30%, 40% and 50% composite data growth for systems serving 5,000, 10,000, and 15,000 subscribers respectively. 2.5 Mbps per subscriber is the starting baseline for these graphs. One could do the same exercise for rural systems, or simply superimpose half the urban APS for rural systems.

It should be noted that these results are for North America. In other regions where viewing habits and internet usage varies dramatically from North America, a different baseline and growth rate may be necessary. However, the fundamental model will be the same.



Figure 5 – Predicted Peak Data Demand Based on Subscribers and Growth Rate

Note that as long as there is not congestion within the bandwidth of the SG's, the number of SGs that are served by the CCAP core do not impact throughput, only the total number of subscribers. This means that modelling the total peak traffic across a disaggregated core with many small SGs becomes very simple. As long as there is sufficient bandwidth in the SG to support the peak bandwidth demand of the SG (i.e. no bandwidth bottleneck), then the results will be the same for a given number of subscribers whether they are in X service groups or divided amongst 4X service groups.

7. More Observations On The Nature of CCAP Core Traffic

If we were to sample the same system for more than a year we would expect the peaks in the second year to be N% greater than the peaks in the prior year. This percentage N would be equal to the annual data growth across the network. However, on a month to month basis traffic is not necessarily N/12, but can vary significantly either up or down from this %, and in some cases be actually negative for a given month. Therefore, all of the graphs in this study have an underlying annual growth rate within the data. We could not adjust monthly traffic apart from the growth rate by subtracting a growth line without having two years of data from which to extract the growth rate. Perhaps in another year, a follow on paper can address this.

7.1. Peak to Valley Traffic Analysis

Figure 6 shows the peak to valley data rate for each platform, on a monthly calculation basis. For some systems, the peak data rate is more than 12 times the valley, while for the least variable systems the peak to valley is approximately 4:1.




Figure 6 – Peak to Minimum Traffic Ratio Across CCAP Platforms

Figure 7 shows the monthly samples for one system in the study. Note on a daily basis the daily peaks to valleys.




Figure 7 - Single CCAP Platform Traffic for One Month Taken On 5 Minute Intervals

Examining one day in detail, one can see the significant difference in traffic as shown in Figure 8.




Figure 8 – Single CCAP Platform Traffic for One Day Taken On 5 Minute Intervals

What can be garnered from this information is that the differences in traffic are significant based on time to day, and further, appear to be well behaved. This may have significant consequences in the future relative to the ability to manage cloud based processing resources and to manage power consumption across the network, gaining appreciable savings by lowering the power required to support low traffic time periods.

There have been previous suggestions to turn off DOCSIS QAM carriers during lower usage times as a means to save power in RF amplifiers. This is illustrated in Figure 9. However, turning off DOCSIS carriers turns out to be impractical. In D3.0 and D3.1 the channels in a bonding group are fixed and are programmed into each cable modem in the SG. If a channel is no longer present, the modems interpret this as a failure and start communicating this back to the CCAP core as a problem. The alternative is to redefine the SG at different points of the day. But to do this requires re-provisioning every modem, taking them offline during the process, resulting in daily service outages.




Figure 9 – Proposed Power Savings By Turning Off Carriers During Low Demand Periods

An alternative potential means to achieve power savings is to reduce the modulation profile and lower the corresponding carrier RF level of each DOCSIS channel accordingly. With regards to modulation throughput, 4096 (2^{12}) QAM (Quadrature Amplitude Modulation) provides double the transmission speed per bit as 64 QAM (2^6). However it takes 18dB more power (translating to approximately 60 times the power) to transmit a 4096 QAM carrier than a 64 QAM carrier. This additional power is a significant operations expense.

Referring again to Figure 7, the peak to valley speed in a single day on this illustration varies by 4X, and as seen from Figure 5, can vary by up to 12X on some systems. However, if we were able to set the highest power level safely above the peak data speed for the month and only reduce the power such that the minimum speed we supported was half of that, this translates to an 18dB power savings at off peak times. We can set the peak power to support a data rate higher than the peak, to enable a safety zone, and a low level significantly higher than the lowest data rate or highest data speed offered.

For example, if our DOCSIS spectrum supports a maximum data throughput of 5.0 Gbps at 4096 QAM, then at 64 QAM the same spectrum supports 2.5 Gbps. We could safely save 10 dB of power across the DOCSIS spectrum, and achieve a minimum data rate of approximately 3.0 Gbps.

Superimposing this on the CCAP platform one day data shows how this might be employed. This is shown in Figure 10. In this example, we use only 10dB for the difference of the signal levels for reasons described below. If the peak rate that can be supported during low traffic time periods is 3.0 Gbps, then we can comfortably have a minimum supported service speed offering of 2 Gbps and still achieve 90% output power savings for that part of the spectrum during low usage time periods.




Figure 10 – Lowering QAM Modulation Profiles During Off Peak Periods

In the CableLabs RPD (remote phy device) spec, there is an adjustable output of 10dB for every carrier. And each one can have its modulation profile changed. So, for all DOCSIS carriers, this means that their levels can be reduced in off-peak periods by up to -10dB which is still highly appreciable (-18dB not possible based on adjustment range). This is depicted in Figure 10.

In an all-IP DOCSIS network, this would result in reducing the required power output of the system hybrids by 10 dB equivalent to approximately 90%. However in a system with MPEG video occupying 50% of the downstream bandwidth, the savings would be only 7-8 dB across the spectrum based on where the DOCSIS carriers are located. In either scenario, since systems never operate at full capacity, we can potentially expect even higher power savings. In today's super high output 1.2 GHz nodes, each RF hybrid consumes over 15 watts at full power. By cutting the composite power even by 6-8 dB will likely result in savings of 7-10 watts per hybrid, meaning saving 28 – 40 watts per node during off peak periods.

Therefore, power savings can become an additional driver to migrate cable systems from MPEG video to full spectrum DOCSIS with all- IP video delivery in the forward path.

But how to accomplish this and not end up in a situation where there is insufficient power and modulation profile to support traffic bursts? The answer is to use analytics. Begin by sampling the downstream traffic on each service group every five minutes. Set the power level and modulation profiles on the DOCSIS carriers such that there is 50% headroom above that level. Compute the appropriate bias voltage/current on the RF hybrids in nodes and amplifiers to support the total composite RF level required. Of course, this requires smart nodes with dynamic remote capability to adjust their power levels. A block diagram of these functions is below in Figure 11.

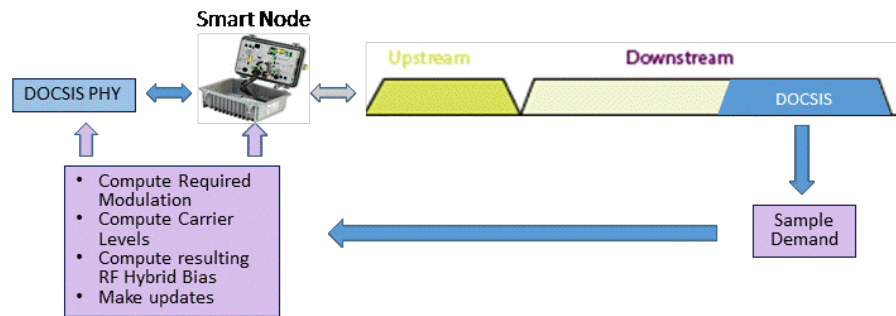


Figure 11 – A Method For Dynamic Power Savings Using Analytics and Smart Nodes

Cisco has applied for a patent that covers this process.

8. Further work

This study is a first step in big data analytics application for traffic prediction. A next step is to take the same sample interval as used in this study and to apply it to measure the peak traffic for every individual service group within a single CCAP platform. One would expect that the larger the average size of the service group the more that individual SG peaks and valleys will reflect the core peak valleys (after all, for a CCAP platform with 50 SGs, one SG represents a sample size of 2%). However as service groups become smaller, either through having more SGs on the same core, or because the CCAP platform is in a rural area serving far fewer total customers, one would expect the variations between an individual service group and the total to be greater. These assumptions have not been validated with data and are an area for future study.

9. Conclusions

We have demonstrated that using the number of subscribers attached to a CCAP core is a better method for predicting traffic through the CCAP core than by any measure employing service group counts, channel counts and an oversubscription factor. We expect the peak data consumption per subscriber to be very consistent across CCAP platforms with similar demographics, such as mostly urban or mostly rural settings. Geographically, we expect a different number per subscriber in each region (e.g. LATAM, Western Europe, and Asia) based on the cultural differences that are reflected in overall internet usage.

One provocative issue that arises from this study is the usefulness of traditional DOCSIS CCAP licensing models which charge for bandwidth on a per service group, or per channel basis. Given that operators revenues are per subscriber and that peak bandwidth is most dependent on the number of subscribers attached to the platform, a per subscriber license plan that enables unlimited bandwidth consumption appears to be the most practical licensing solution. This will become even more important with DAA (distributed access architecture) networks and cloud based CCAP core functions, wherein service groups are no longer an architectural limitation as imposed by dedicated CCAP hardware.

Bibliography

1. Cisco Visual Networking Index: Forecast and Methodology, 2016–2021, Cisco White paper, June 6, 2017
2. Grover S., Ensafi R., Feamster N. (2016) A Case Study of Traffic Demand Response to Broadband Service-Plan Upgrades. In: Karagiannis T., Dimitropoulos X. (eds) Passive and Active Measurement. PAM 2016. Lecture Notes in Computer Science, vol 9631. Springer, Cham

Assuring Data Delivery from Critical IoT Devices

A Method to Create New Services and Mitigate Liability

A Technical Paper Prepared for SCTE•ISBE by

Michael Klobardans

Principal Engineer

Charter Communications

14810 Grasslands Drive. Englewood, CO. 80112

(720) 518-2539

Michael.Klobardans@Charter.com

Shlomo Ovadia

Director

Charter Communications

14810 Grasslands Drive. Englewood, CO. 80112

(720) 536-1686

Shlomo.Ovadia@Charter.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Data Delivery Network Architecture & Operation.....	5
1. Data Delivery Network Architecture	5
1.1. IoT Packet/Header Flow in the Data Delivery Network Architecture	5
1.2. Key Architecture Points for Determining Data Loss.....	7
1.3. Identifying Data Loss at Major Network Segments and Triggered Actions	8
2. Operational Details.....	10
2.1. Registering Devices that Send Critical and/or Important Data	10
2.2. Identifying an IoT Device and Gateway Pair.....	10
2.3. Tracking an IoT Device/Gateway Pair ID.....	11
2.4. Operational Details.....	11
Comparison with Alternative Methods.....	14
Conclusion.....	15
Abbreviations	16
Bibliography & References.....	17

List of Figures

Title	Page Number
FIGURE 1 - PROJECTED NUMBER OF HEALTHCARE DEVICES WORLDWIDE	3
FIGURE 2 - SMART SPEAKER PENETRATION IN US HOMES.....	3
FIGURE 3 - IOT DATA DELIVERY NETWORK ARCHITECTURE: MAJOR SYSTEM COMPONENTS.	5
FIGURE 4 - MESSAGE FLOW DIAGRAM FOR IOT DATA DELIVERY NETWORK.	6
FIGURE 5 - IOT DATA DELIVERY NETWORK ARCHITECTURE: ARCHITECTURE DATA COLLECTION POINTS.....	7
FIGURE 6 - IOT DATA DELIVERY NETWORK ARCHITECTURE: DATA LOSS SEGMENTS.....	9

List of Tables

Title	Page Number
TABLE 1 - IOT DATA PACKET FLOW THROUGH THE DATA DELIVERY NETWORK ARCHITECTURE.....	6
TABLE 2 - MESSAGE FLOW DIAGRAM THROUGH THE DATA DELIVERY NETWORK ARCHITECTURE.	7
TABLE 3 - NETWORK COLLECTION AND ASSURANCE POINTS.	7
TABLE 4 - NETWORK DATA LOSS ACTION MATRIX.	9
TABLE 5 - DEVICES THAT USE THE IPV6 FLOW LABEL HEADER.	11
TABLE 6 – EXAMPLE OF IOT DATA REPOSITORY DATABASE ENTRIES.	13
TABLE 7 - COMPARISON BETWEEN THE LI AND IOT DATA DELIVERY ASSURANCE METHODS.....	15

Introduction

Internet of Things (IoT) devices enable Machine-to-Machine (M2M) data transmissions where a sensor or appliance (machine) at a subscriber's residence collects and sends information to a different machine at another location, such as a mobile phone application, a gateway or even a data center. From that second machine, information is processed, made meaningful and available for human consumption. IoT devices mostly offer casual conveniences, such as changing the lighting color in the dining room or using voice commands to play a genre of music. Lately, a new class of IoT devices are emerging in homes and businesses that send critical and/or important messages such as personal healthcare data and industrial application data (e.g., factory temperature and pressure levels, etc.) to external processing or monitoring service providers.

IoT products have many categories: Gaming (Sifteo, console sensors), Security (Cameras, door/window sensors), Convenience (smart speakers, light bulbs, window shades, smart kitchen appliances), Monitoring (Plant soil moisture, pet food levels, lawn watering, automotive), Healthcare and others. Figure 1 shows the explosive growth of the Healthcare market, which is projected to reach \$137 Billion in just two years [1].

Figure 2 shows, for example, the explosive growth of Smart Speaker devices in the US [2]. This category of IoT devices is interesting because they can control IoT devices as well as provide other services. Statistics from comScore reported that existing products such as Smart Speakers (Google Home, Amazon Echo, Apple HomePod, Sonos One, etc.) increased 50% to 18.7 Million US households in just three months between November 2017 and February 2018. This penetration growth reflects a broad acceptance of automation in US homes.




Figure 1 - Projected Number of Healthcare Devices Worldwide




Figure 2 - Smart Speaker Penetration in US Homes

If vital data packets from these IoT devices are dropped anywhere in the network path between the IoT device and a processing and/or monitoring provider across the Internet, the necessary IoT function can be compromised, which in turn may as a consequence lead to property damage, health or safety problems. While critical IoT data devices use higher-layer data assurance protocols such as CoAP (Constrained Application Protocol) to acknowledge receipt of an IoT transmission, no receipt is possible if the source data fails to reach the processing/monitoring destination [1].

This paper provides one such solution with two main benefits:

1. A new service that alerts subscribers when expected, periodic data is missing.
2. Capability to identify where on the timeline packets are dropped in the Internet or other external network, which can prove helpful in ascertaining root-cause of outages.

The first of the two main benefits is alerting a subscriber when messages from critical and/or important IoT devices drop. This is a new proactive and potentially profitable customer service feature. For example, a medical blood oxygen sensor at the subscriber home sends a regularly-scheduled blood oxygen level information to a 3rd party monitoring company. If these regularly-scheduled data packets are no longer being received and transported through the Cable MSO network, the subscriber and the monitoring company are notified for corrective action to avoid a serious health consequence.

The second of the two main benefits is the ability to track packets in the Cable MSO networks to establish where and when on the transmission route packets were dropped, including drop events occurring after delivery to a non-MSO network such as the Internet. In addition to root cause identification, this information can mitigate liability issues by providing tracking visibility within the Cable MSO's systems of networks.

This paper is organized as follows. Section 1 presents the IoT data delivery network architecture and diagrams of IoT message flow through the data delivery network. Integration with existing Cable MSO transport protocols such as IPv4, IPv6, and MAP-T are explained. This section also discusses how IoT data loss is identified, and the triggers for required actions. Section 2 provides operational details of how the registered IoT device sends its critical data through various points of the data delivery network. Comparison with alternative methods to track critical IoT data through the Cable MSO network is discussed in the next section. The paper concludes with a summary of the key features of the assurance of critical data delivery through the Cable MSO network, and its benefits.

There are three clarifications listed here to give context to this paper: registration, generic naming, and SMB applications.

- **Registration:** Critical and/or important IoT device data must first be identified and registered for tracking and notification services to begin. Section 2 has details on this process, but for initial understanding, assume that the IoT device is pre-registered with the Cable MSO and its data is agreed as critical, important or 'of interest' as defined in section 2.4.4. In this paper, these IoT devices are referred to as 'Registered IoT Devices'. The alerting services and data tracking concepts are confined to these Registered IoT Devices only.
- **Generic Names:** The Cable MSO network device names used in this paper are intentionally generic because each Cable MSO has different network topologies and device function names. For example, there is no single egress point for a Cable MSO of any significant size. In general, the goal is to present concepts that can be applied to the network topologies of any Cable MSO.

- **SMB Applications:** While the focus in this paper is residential subscriber services, the same benefits and procedures apply equally to SMB markets although more formal implementation may be necessary to address the complexities of business demands and networks compared to residential needs.

Data Delivery Network Architecture & Operation

1. Data Delivery Network Architecture

1.1. IoT Packet/Header Flow in the Data Delivery Network Architecture

Figure 3 depicts an architecture of physical devices and software entities in an IoT data delivery system for important and/or critical data from an IoT device. The key architecture components include an IoT device that sends data to a Home Gateway (HGW) where a unique ID is added. The Cable MSO premises contains an Access Network router that records HGW IoT packets and routes them through multiple network elements in the core network including a final egress router that both routes IoT data packets to a non-MSO network and sends IoT IPv6 headers to the IoT Data Repository database. Finally, these IoT data packets are received by a 3rd party monitoring provider. The blue text in the diagram below denotes actions that happen with an associated device and not a flow of packets or header copies. Generally, the device's packets use the UDP transport protocol which has no acknowledgement features. Also, while some IoT devices may use encryption for security purposes, an IPv6 header is always available which contains and tracks the unique ID.

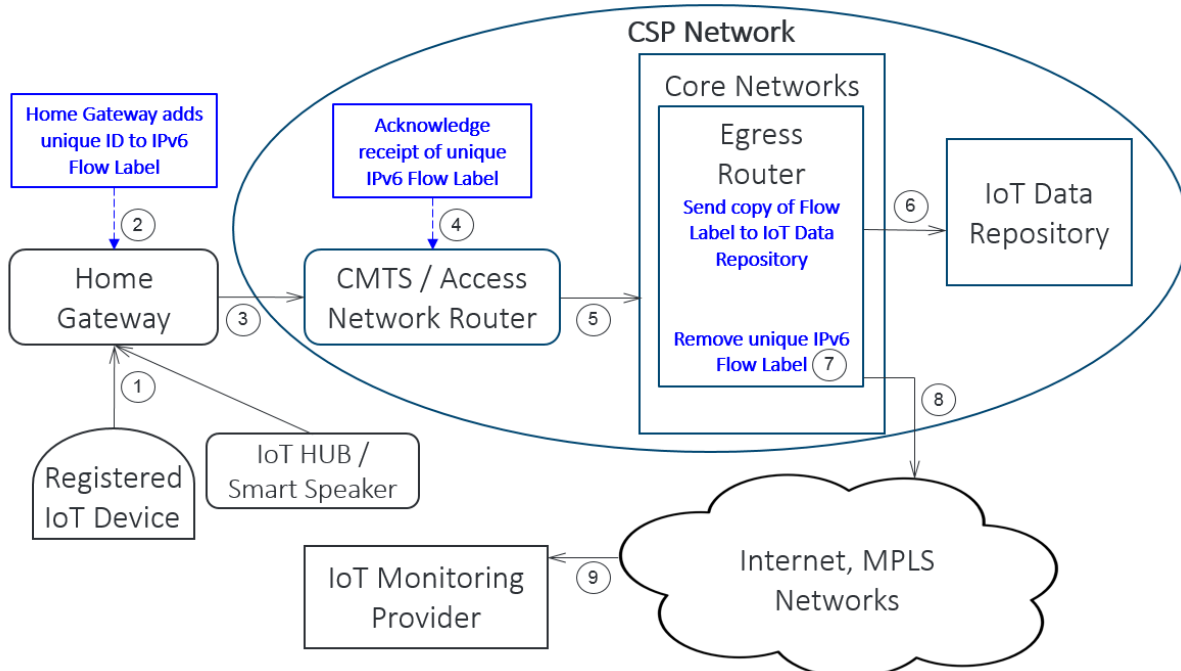


Figure 3 - IoT Data Delivery Network Architecture: Major System Components.

Table 1 - IoT Data Packet Flow through the Data Delivery Network Architecture.

Step	Description of action at each architecture point
1	Registered IoT Device, Hub or Smart Speaker sends IPv4/IPv6 packet(s) to a Home Gateway (HGW).
2	The HGW adds a unique ID to the IPv6 Flow Label field that identifies the IoT device and the HGW.
3	The HGW sends the packets through a Cable Modem (CM) to an Access Network (AN) Router.
4	The AN Router sends a copy of the IoT data header and timestamp to an IoT data repository database.
5	The Access Network Router sends the data frame to the Cable MSO's core network infrastructure.
6	The Core Network (CN) Egress Router sends a copy of the IoT data header and timestamp to the same IoT data repository database as in step #4.
7	The CN Egress Router removes the unique ID from the IP packet.
8	The CN Egress Router sends the data frame to a non-MSO network such as the Internet.
9	The data packet is delivered to the final destination, such as an IoT Monitoring Provider.

Figure 4 shows the message flow diagram through the IoT data delivery network. The arrows indicate an action toward a destination, not necessarily a physical device. The message numbers shown in Figure 4 correspond to the numbers in the IoT data delivery network architecture shown in Figure 3.

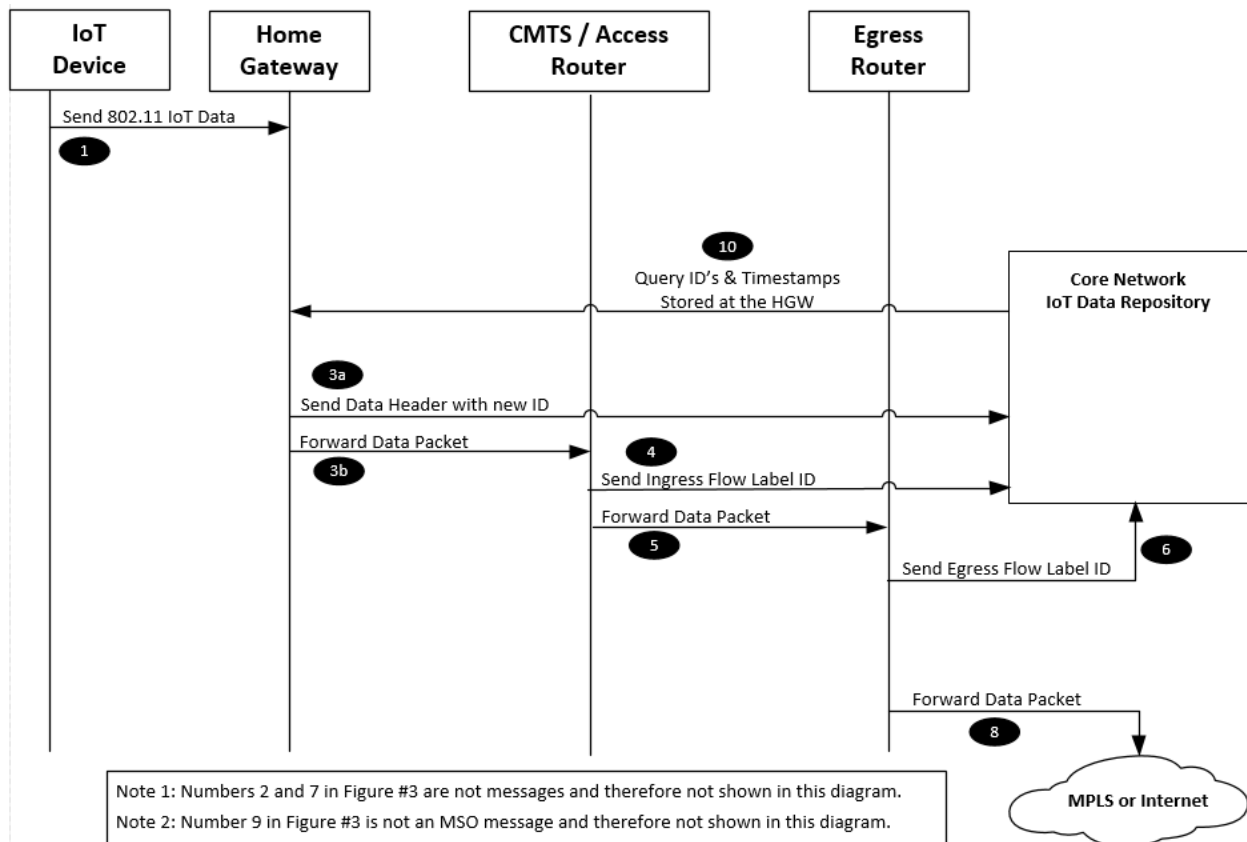


Figure 4 - Message Flow Diagram for IoT Data Delivery Network.

Table 2 - Message Flow Diagram through the Data Delivery Network Architecture.

Step	Message Flow Diagram Description
1	The Registered IoT Device, Hub or Smart Speaker sends one or more packets/frames to the HGW.
3a	The HGW creates a copy of the IoT IPv6 header, w/ID and sends it to the IoT Data Repository database
3b	The HGW forwards the complete IoT packet to the AN Router.
4	The AN Router sends a copy of the IoT IPv6 header to the IoT Data Repository database.
5	The AN Router forwards the complete IoT packet to the Core Network Egress Router.
6	The CN Egress Router sends a copy of the IoT IPv6 header to the IoT Data Repository database.
8	The CN Egress Router forwards the IoT packet using Non-MSO controlled networks onto the final destination.
10	Periodically, the IoT Data Repository database queries the HGW for all IoT messages sent. The period is defined for each device when creating the SLA and depends on the criticality of the device's data.

1.2. Key Architecture Points for Determining Data Loss

The basic IoT data delivery network architecture consists of three data collection points and a data assurance mechanism as shown in Figure 5. The data collection points are identified as points A, B and C representing message ingress at the home (A), Access Network ingress (B) and Core Network egress (C). Point D is the Data Assurance point.

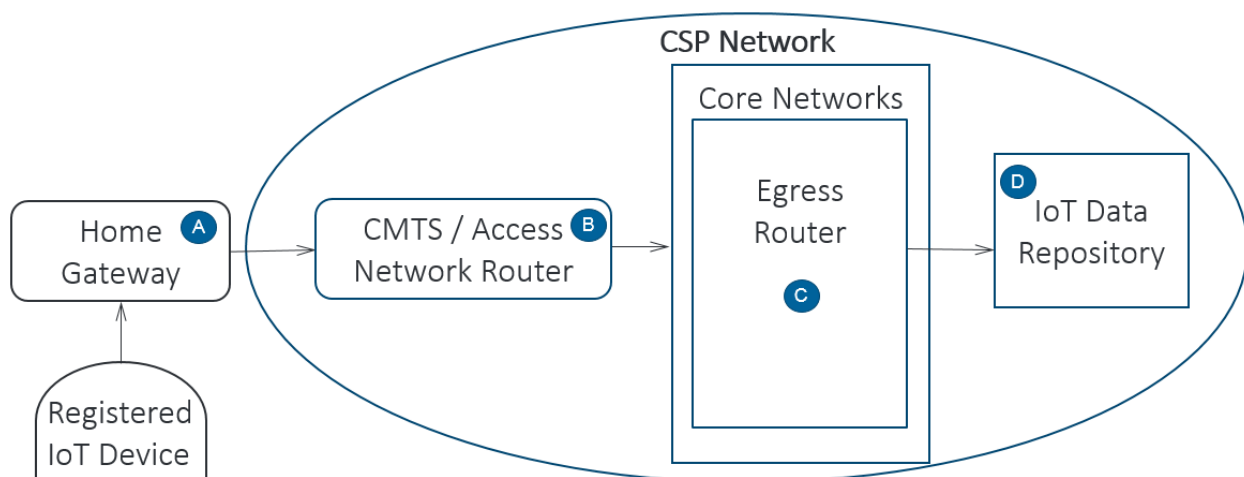


Figure 5 - IoT Data Delivery Network Architecture: Architecture Data Collection Points.

Table 3 - Network Collection and Assurance Points.

Network Point	Description	Function
A	Home Gateway (Home Ingress Point)	Data Collection Point
B	Access Network Router (AN Ingress Point)	Data Collection Point
C	Egress Router (Core Network Egress Point)	Data Collection Point
D	IoT Data Repository database	Data Assurance Point

Network Point A - Unless an IoT device uses cellular or other telecom protocols to communicate, the Cable MSO network is used. There are three common methods used to send IoT data to the HGW:

- a. The HGW Access Point directly receives IoT wireless signals and forwards them to the Cable MSO Access Network.
- b. The HGW Access Point converts wireless signals from an IoT device to a wired protocol and forwards that data to a wired IoT hub. IoT hub messages may then be forwarded back to the HGW for transport to the Cable MSO Access Network.
- c. The IoT device directly communicates to an IoT hub or Smart Speaker (not shown in Figure 5). The hub or speaker then sends messages to the HGW to forward onto the Cable MSO Access Network.

The HGW copies the IoT frame header, adds a unique ID and a timestamp and then temporarily stores this information. The HGW information store is critical because the HGW is the only Cable MSO owned device that resides on the subscriber's LAN and is therefore the only point along the transmission route proximate to the subscriber at which originating IoT device transmissions can be recorded. The HGW then forwards the complete data packet (with the ID in the Flow Label field in the IPv6 header) onto the Access Network using normal routing mechanisms.

Network Point B - The Access Network Router is the entry point into the Cable MSO access network infrastructure which becomes another critical point for data collection because this point is the first time the data is fully 'inside' a totally Cable MSO controlled premises. Data from Registered IoT Devices are recognized at this point and like the HGW, a copy of the packet header is made and sent to the IoT Data Repository database. The IoT header copy is used to mark entrance into the Access Network.

Network Point C - The Egress Router is the final data collection point. Like the HGW and the Access Network Router, the Egress Router recognizes a packet from a Registered IoT Device by finding a non-zero value in the Flow Label header field, makes a copy of the packet header, and sends that header copy to the IoT Data Repository database. This header copy is used to mark the exit of the data packet from the Cable MSO-controlled network. The full data packet is delivered to the adjoining external (non-MSO) network associated at that point.

Network Point D - The IoT Data Repository database software agent periodically queries the HGW for a copy of the home premise IoT information store. The IoT Data Repository database compares the HGW entries with its own copy to verify that no Registered IoT Device data was lost between the HGW and the Access Network Router or the Egress router. All the Flow Label header copies with a matching ID are stored with their timestamps. This information is easily parsed to find any missing entries, and the identification of any missing entries triggers an action to find where the packet drop happened within the system of Cable MSO controlled networks (explained more fully in the section 2.4). This information can also be used to send an alert(s) to notify the subscriber of dropped packets or expected packets that were never received by the HGW.

1.3. Identifying Data Loss at Major Network Segments and Triggered Actions

Figure 6 depicts a simplified Cable MSO network showing the major network segments between points A, B, C and D, where loss can occur as identified by the red crosses. Table 4 shows a matrix of triggered actions as a result of data loss in each major network segment. It is important to determine data loss from a specific network area for accountability reasons and speedy repair. Depending on the IoT type and SLA, a

subscriber notification of data loss may or may not be appropriate. Loss between any Subscriber's network device and the HGW is not part of the Cable MSO network unless HGW is malfunctioning. HGW errors are handled by existing policies and procedures and are outside this proposal.

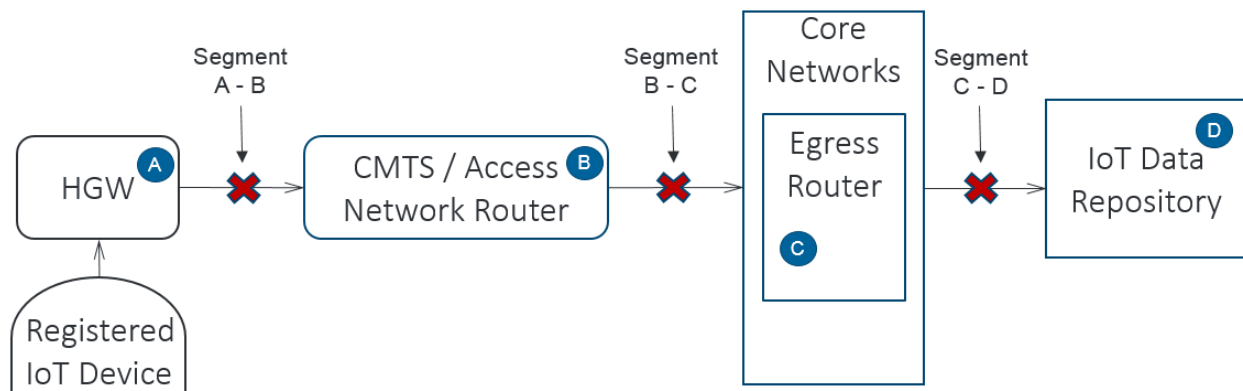


Figure 6 - IoT Data Delivery Network Architecture: Data Loss Segments.

Table 4 - Network Data Loss Action Matrix.

Network Segment	Location of Data Loss	Triggered Action
A - B	HGW to Access Network Router	Investigate path from HGW to AN Router
B - C	Access Network Router to Egress Router	Investigate path from AN Router to Egress Router
C - D	Egress Router to Repository	Investigate path from Egress Router to Repository

This proposal adds an additional capability from existing methods to find packet loss, and specifically from Registered IoT Devices only through the parsing of the database in the IoT Data Repository database.

Network Segment A – B:

Packet data loss on the Network Segment A – B can only be found from the comparison of the IoT Data Repository database to the HGW data store. As mentioned in the paragraph at Section 2.3, Network Point D, the IoT Data Repository database, periodically queries the HGW for a copy of the home premise IoT information store, and compares the HGW entries with its own copy to verify that no Registered IoT Device data was lost between the HGW and the Access Network Router.

Data loss can result from several sources:

- The HGW itself, such as a WAN interface defect that did not forward the data (or any data).
- An upstream Cable Modem defect, regardless of being an integrated or separate physical unit.
- A defect in the coax cabling, splitters or other components from the Cable Modem to the pedestal.
- Any component at the node, hub or headend, including the Access Network Router.

The general action triggered from loss in this network segment is to investigate the devices and network paths. More specifically, patterns are analyzed first to narrow the scope of investigation. An example is if one Registered IoT Device has no packets being sent, but other wireless devices on the customer's LAN have no issues with internet communications, the problem may be localized to the IoT device itself and the investigation should start with this device.

Network Segment B – C:

This segment can be complex, involving many separate networks and devices. Traditional network troubleshooting methods are used to find the issue and provide a fix. Most likely, non-IoT traffic is affected as well if transport problems are encountered on this network segment.

The general action triggered from loss in this network segment is to narrow the investigation scope, again, by using traditional troubleshooting methods.

Network Segment C – D:

More than one IoT Data Repository database may exist in actual implementations. The issues are either routing the data to a specific repository or finding the reason(s) that the repository won't accept or store the information.

The general action triggered from loss in this network segment is to determine if the problem source is network-based or is with the device(s) or database.

2. Operational Details

2.1. Registering Devices that Send Critical and/or Important Data

Before describing the packet flow details in this new architecture and in order that tracking and notification functionality can begin, critical and/or important IoT devices must first be identified and registered. The originating data source in this paper is the IoT device itself, but the treatment from critical and/or important data devices is not confined to IoT devices only – any device can be registered for tracking and notification functionality. Thus, while any device can qualify, only IoT devices are mentioned for reasons of simplicity. Assume that the IoT source device of the previous section is pre-registered with the Cable MSO and its data is defined as critical or important and therefore tracking the data from this device is desired by both the Cable MSO and the subscriber. In this paper, critical and/or important IoT devices are referred to as a 'Registered IoT Device'. Alerting services and data tracking is confined to registered devices only. Both critical and important IoT data are tracked and monitored, but service actions vary depending on the degree of consequence severity as further explained in section 2.4.4. The registration process, use of a subscriber portal, automatic identification and other operational details are out of scope for this paper. Also, for purposes of this analysis unless an IoT device uses cellular or telecom protocols for communications, the assumption is made that they connect wirelessly to a HGW and use the Cable MSO's infrastructure for at least partial transport of their significant data to its final destination.

2.2. Identifying an IoT Device and Gateway Pair

Tracking data from critical and/or important IoT data devices at both the entry and exit points in a Cable MSO's network and then recording those results is foundational to this proposal. Identification includes the IoT device and the Home Gateway pair that sent critical and/or important IoT data to the Cable MSO network. A method to identify this paired information is needed, but anonymity is also needed for privacy concerns. Cable MSOs should track and record the movement of critical and/or important IoT data from the ingress to the egress points in their networks, but not record the data itself. Most subscribers' data devices still use the IPv4 protocol. A key feature of this architecture is to convert all IPv4 traffic to the IPv6 protocol at the Home Gateway. There are several methods that provide this conversion; MAP-T or MAP-E being the preferred method used by several of the largest Cable MSOs around the world.

The reason to use IPv6 is to repurpose a field in the header called the ‘Flow Label’. The Flow Label field has a 20-bit length that defaults to all zeros. The HGW assigns an ID value to indicate the specific IoT device and gateway as a unique pair that is used to send the IoT data to the Cable MSO network.

A unique value to identify the IoT device/gateway pair can be derived using at least two methods:

1. Hashing – A hash value that fits within the 20-bit limit of the Flow Label header.
2. Mapping – A simple map of a registered IoT device and the subscriber’s Home Gateway.

The ID calculation function is best performed by an OSS function, not the local HGW. This is done during the IoT Registration process, again, outside the scope of this paper. Both methods above ensure anonymity even if a Flow Label with an ID is accidentally exposed publicly or internally because the information cannot be traced to a subscriber or device without the mapping key or hash algorithm. Once an IoT/HGW pair has an assigned unique ID, that value is placed into the IPv6 Flow Label field of any packet transmitted from that device/HGW pair. The 20-bit Flow Label field is sufficient to provide enough unique IDs such that a single CMTS can serve over 40,000 subscribers where each subscriber can have up to 24 Registered IoT Devices. This ensures scalability for the near term but there are methods existing today that can be used to ensure ID extensibility with no practical limit, such as using IPv6 extension headers.

2.3. Tracking an IoT Device/Gateway Pair ID

As explained in section 1.3, there are three points at which the Unique ID (IoT device/Gateway pair information) is copied and sent to the IoT Data Repository database as described in Table 5:

Table 5 - Devices that Use the IPv6 Flow Label Header.

Device	Purpose
Home Gateway	Records IoT device data received at the home gateway
Access Network Router	Records IoT device data received at the Cable MSO access/core network
Cable MSO Egress Router	Records delivery of IoT device data from the Cable MSO egress point

2.4. Operational Details

2.4.1. Home Gateway Operational Details

The HGW identifies a data transmission from a Registered IoT Device, either directly or indirectly from an IoT hub or another device. If the HGW doesn’t have an integrated cable modem, a separate Cable MSO provided cable modem encapsulates the Ethernet frame for transport between the cable modem and the CMTS (or similar) using DOCSIS protocols. In this paper, we assume that the cable modem is integrated into the HGW device. The HGW then completes the following steps:

1. Identifies packets from a Registered IoT Device using its MAC address or other identifier.
2. Converts from a wireless protocol such as IEEE 802.11ac to a wired protocol such as IEEE 802.3 (Ethernet).
3. Converts IoT device data from IPv4 packets to IPv6 if needed.
4. Assigns a predetermined, unique ID value (IoT and gateway pair) into the IPv6 Flow Label header field.
5. Copies that IPv6 header, applies a timestamp and:
 - a. Stores this header information in non-volatile memory on the gateway.
 - b. Forwards this same header information to the IoT Data Repository database.

6. Forwards the complete packet with the modified IPv6 Flow Label to the Access Network Router.
7. Waits for a periodic query from the IoT Data Repository database (or other actor performing this verification step). After acknowledgement that data was successfully transported from the HGW to the IoT Data Repository database, the HGW data stored in step 5a above is reset and ready to be used for new entries.

2.4.2. Access Network Router Operational Details

After the IoT data packet exits the HGW, (or the cable modem or the ONU), it traverses the Cable MSO access network, and then terminates at the Access Network Router. At this point in the access network, the following operations are completed:

1. Identifies in-scope packets by detecting a non-zero IPv6 Flow Label value in the IPv6 header.
2. Copies the IPv6 header, applies a timestamp and:
 - a. Stores this header information in non-volatile memory on the Access Network Router.
 - b. Forwards this same header information to the IoT Data Repository database.
3. Forwards the complete packet with the modified IPv6 Flow Label using standard routing procedures throughout the Cable MSO's network infrastructure to the Network Egress Router.
4. Waits for a periodic query from the IoT Data Repository database (or other actor performing this verification step). After acknowledgement that data was successfully transported from the Access Network Router to the IoT Data Repository database, the Access Network Router data stored in step 2a above is reset and ready to be used for new entries.

2.4.3. Egress Network Router Operational Details

Standard Cable MSO routing transports critical important and /or important IoT data from the Cable MSO Access Network Router which eventually terminates to a router at the edge of the Cable MSO network. At this point in the core network, packets egress from the Cable MSO controlled networks and complete a hand-off to a non-MSO network such as the Internet or MPLS network. The following operations are completed at the Cable MSO Egress Router:

1. Identifies packets of interest by detecting a non-zero IPv6 Flow Label value in the IPv6 header.
2. Copies the IPv6 header, applies a timestamp and:
 - a. Stores this header information in non-volatile memory on the Egress Router.
 - b. Forwards this same header information to the IoT Data Repository database.
3. If the external non-MSO network uses these IP protocols:
 - a. IPv6 - then reset the IPv6 Flow Label in the header of the complete packet to a value of all zeros and forward to the adjoining network.
 - b. IPv4 - then convert the IP protocol of the packet from IPv6 to IPv4 and forward to the adjoining network.
4. Forwards the complete packets to the adjoining non-MSO network using current procedures for normal operation.

2.4.4. IoT Data Repository database Operational Details

The IoT Data Repository database contains copies of the IPv6 Flow Label headers and timestamps from three network points that have handled the packets issued from registered IoT device transmissions; the HGW, the Access Network Router, and the Egress Router. This is described in sections 1.3 and 2.3. The

purpose of the IoT Data Repository database is to store the following records of the IoT data transmissions:

- Initially generated and received at the Home Gateway
- Received at the Access Network Router
- Received at the Egress Router
- Successfully handed off to the non-MSO adjoining network

A typical IoT Data Repository database structure example is shown in Table 6 where the IoT/HGW ID values are hexadecimal and time is represented as Unix Epoch Time values:

Table 6 – Example of IoT Data Repository database entries.

Entry	IoT/HGW ID	HGW Time	Access Network Time	Egress Router Time
1	B0301	1 531 179 199.501	Not Available	1 531 179 200.691
2	40A4E	1 531 179 200.519	Not Available	Not Available
3	665B2	1 531 179 198.637	1 531 179 198.660	1 531 179 199.112
4	665B2	1 531 179 798.243	1 531 179 798.651	1 531 179 799.145
5	665B2	Not Available	Not Available	Not Available

Any missing timestamp entry in the IoT Data Repository database indicates missing data at that collection point; HGW, AN, or Egress router. If a notification SLA is active, a lack of entries from an IoT device indicates missing data and triggers one or more alerts to the subscriber. Therefore, this table serves as a missing data detection point that triggers proactive alerting, which in turn can provide the customer with notice of the possibility their IoT devices are not receiving signal and cannot function as intended.

An example of missing data and corresponding action responses from Table 6 is now described:

Entry #1 has no timestamp value delivered from the AN collection point, but has a value from the Egress router collection point. This missing data is inconsequential and no action is taken because the IoT data was received into and egressed from the Cable MSO's system of networks.

Entry #2 is concerning because it represents data loss within the network segment from the HGW and the AN collection points. A triggered action would include investigating these network elements:

- If data is missing from many subscribers terminating at the same AN router, that AN router, AN router interface or physical media attaching to the AN router interface are investigated. As previously stated, this type of data loss most likely affects all data from all sources and not just IoT data. Therefore, other alarms and procedures would most likely detect and address this outage.
- If data is missing from one subscriber only, faulty components could include the HGW itself, the cabling between the HGW and the CM (if they are separate physical devices), the Coax/Fiber cabling between the CM and the pedestal and every component between the pedestal and the AN.

Entries 3 - 5 are an example of missing data from an IoT device sending regular and periodic messages once every 10-minutes. An SLA is in place where proactive alert notifications are sent if the IoT device doesn't send data as expected within the stated time period. Entries 3 & 4 represent normal and expected messages from the IoT device. Entry #5 is expected, but failed to be recorded at any collection-point.

Alert notifications would be sent to the subscriber and perhaps a 3rd party monitoring agency as defined in the SLA.

Registered IoT Device data are classified as one of three different types: Critical, Important or ‘of interest’ and might have corresponding SLA service packages labeled as Gold, Silver and Bronze. Cable MSO triggered actions such as notifications will vary depending on the class of data and SLA agreement. For example, missing data from a medical pulse monitor can be critical with only a few minutes to respond before health is threatened or even death results. This data class is ‘critical’ and results in immediate notifications to the subscriber and also health professionals and emergency providers. By contrast, an important message would be a power outage to an IoT-monitored large freezer could result in only an informative notification to the subscriber. The contents in the freezer may be unaffected for many hours during a power outage. This important information class is higher than ‘of interest’ but not critical. The class titled ‘of Interest’ is data that is not critical or important, but the subscriber wants Cable MSO monitoring and notification messages sent.

Lastly, the IoT Data Repository database can be archived as desired or rewritten after a suitable time period as the Cable MSO desires. Database reliability is met through common practices for redundant server/storage that are used today.

Comparison with Alternative Methods

Complex solutions exist today to track data, video and voice, specifically Lawful Intercept (LI) [4] for legal monitoring purposes, however, this level of effort is expensive to administer and may be affected by governing restrictions under multiple industry standards [5] [6] and those initiated by a judicial/administrative legal order.

LI begins as an unexpected legal order from the state or federal government judicial or administrative branch. This is therefore an unplanned and reactive request. A subscriber’s IP address is first identified, and then manual administration is performed to mirror a copy of those IP Packets (Voice, Video and Data) to the government agency that made the demand. Manual administration is again needed remove the packet mirroring. The Cable MSO are subject to restrictions under law regarding copying, redirecting or storing these packets for purposes other than securing, maintaining, and otherwise delivering the underlying services. Using the LI solution for IoT devices would result in significant architectural changes and also copying and storing the entire packet, including sensitive payload data which makes this a heavyweight solution and introduces many privacy problems that would need to be solved.

By contrast, the data delivery assurance method described in this paper plans data tracking in advance for each Registered IoT Device and is governed by subscriber SLAs. It is lightweight in comparison to LI because only packet headers are stored and it is also fully controlled by the Cable MSO. It also enables new subscriber services to pay for the setup and operational costs and become profitable. Because only IPv6 headers are copied and stored, there are no privacy issues which results in benefits without risks.

Table 7 - Comparison between the LI and IoT Data Delivery Assurance Methods

Method Name	Benefits	Costs
IoT Data Delivery Assurance Method	<ul style="list-style-type: none">• Relatively simple to implement• Relatively low implementation cost• Generates on-going subscriptionss• Enables Cable MSOs to introduce new services• Controlled proactive agreement w/subscriber and 3rd party vendors (SLA)• Minimizes potential liability issues due to data packet loss• Allows deeper Cable MSO integration and involvement with subscriber's IoT devices	<ul style="list-style-type: none">• Requires the installation and maintenance of additional network functionality:• HGW, AN & CN header copies and timestamps• IoT Data Repository database
Lawful Intercept (LI) Method	<ul style="list-style-type: none">• None for the Cable MSO	<ul style="list-style-type: none">• Complex administration• Privacy issues• Reactive service• Costly• No Cable MSO benefit• Not designed for tracking data• Create data storage

Conclusion

Explosive growth of IoT devices is expected to continue well into the next decade (e.g., Figure 1 and Figure 2). Subscribers have already embraced the use of IoT devices as conveniences, and are accepting and using a new class of IoT applications to monitor and transmit critical and/or important data in various areas such as personal health, safety, etc. If the IoT device's critical and/or important data is dropped anywhere in the network between the IoT device and the Cable MSO egress router or the 3rd party monitoring provider across the Internet, then there can be significant consequences to the subscriber.

In this paper, a novel and simple data delivery assurance method was presented to resolve this problem. The method is based on first registering the IoT device and the HGW pair in the Cable MSO database. The registered IoT IPv6 label header, which has a unique ID, is then used to track the IoT data packet flow through the network. In addition, the registered IoT IPv6 label header is timestamped at various key network elements as the packet is transmitted from the HGW to the Cable MSO egress router and stored at the IoT data repository database. If an IoT data packet is lost in the Cable MSO network, a matrix of triggered actions is enacted at the identified network segment. Depending on the IoT type, data type and SLA, the subscriber may receive a notification for the loss of IoT data packets. The IPv6 header reuse is a transparent function to both the IoT devices and the network elements beyond the Cable MSO network infrastructure. This ensures that non-MSO devices and network elements function normally when this solution is applied.

There is great potential for new service offerings associated with the use of IoT devices that transmit critical and/or important data. Furthermore, strong partnerships with external 3rd party monitoring

agencies can result in proactive data loss notifications that are sent to both the subscriber and the monitoring agency for faster resolution and to lessen the impact of missing data. Liability concerns associated with unfounded fault attribution can be mitigated, if not eliminated through the use of this tracking method.

Alternative methods such as the LI method are costly, complicated to administer, and are not designed for tracking the IoT data packets through the Cable MSO network. In contrast, the data delivery assurance method is relatively simple to implement and administer, and can even prove profitable depending how the Cable MSO may offer such additional functionality to the subscriber or other parties, such as health monitors.

Wireless cellular technology is currently being used in many cases for IoT data transport. However, cellular carriers do not have the tracking and alerting capabilities or the granular focus to support a per IoT device SLA contract. Adopting the IoT Data Delivery Assurance method will potentially enable the Cable MSOs to compete with the wireless cellular carriers as an alternative low-cost solution to guarantee the delivery of critical and/or important IoT data from the home or the business to the 3rd party monitoring companies.

Abbreviations

AN	Access Network
AP	Access Point
CM	Cable Modem
CN	Core Network
CMTS	Cable Modem Termination System
CoAP	Constrained Application Protocol
DOCSIS	Data Over Cable Service Interface Specification
EU	End User
HGW	Home Gateway
ID	Identification
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IoT	Internet of Things
LI	Lawful Intercept
M2M	Machine to Machine
MAC	Media Access Control
MAP-T	Mapping Address and Port using Translation
MAP-E	Mapping Address and Port using Encapsulation
MPLS	Multi-Protocol Label Switching
MSO	Multiple System Operators
OSS	Operations Support Systems
ONU	Optical Network Unit
SCTE	Society of Cable Telecommunications Engineers
SLA	Service Level Agreement
US	United States

Bibliography & References

- [1] Internet Engineering Task Force (IETF) RFC7252, the Constrained Application Protocol (June, 2014).
- [2] IoT Number of Connected Devices, Statista.
<https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide>
- [3] Smart Speaker Penetration among US Wi-Fi households, Forbes.
<https://www.forbes.com/sites/johnkoetsier/2018/04/11/smart-speaker-penetration-just-exploded-50-in-3-short-months/#744e086b4fbf>
- [4] Lawful Intercept Overview, Cisco.
https://www.cisco.com/c/en/us/td/docs/routers/10000/10008/feature/guides/lawful_intercept/10LIovr.html
- [5] Lawfully Authorized Electronic Surveillance, TIA/EIA/J-STD-025A.
<http://cryptome.org/esp/45-jstd025a.pdf>
- [6] PacketCable Electronic Surveillance Delivery Function to Collection Function Specification, PKT-SP-ES-DCI-I02-070925.
<https://www.forbes.com/sites/johnkoetsier/2018/04/11/smart-speaker-penetration-just-exploded-50-in-3-short-months/#744e086b4fbf>

Automation of Virtual CCAP to Reduce OpEx and Enable New Revenue Streams in the Access Network

A Technical Paper prepared for SCTE•ISBE by

Michael O'Hanlon

Principal Engineer

Intel Corporation – Network Platforms Group

michael.a.ohanlon@intel.com

Eric Heaton

Platform Solutions Architect

Intel Corporation – Network Platforms Group

eric.d.heaton@intel.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Active Power Management per Server	5
Software Infrastructure for Active Power Management	10
Savings Through Orchestration	12
Calculating Potential Efficiency using Consolidation	14
Consolidating Fixed, Mobile, and Enterprise onto the Same Infrastructure	17
Conclusion.....	19
Bibliography & References.....	20

List of Figures

Title	Page Number
Figure 1 - Primetime driving peak network demands	3
Figure 2 - Moving access functions to standard COTS servers	4
Figure 3 - Test server and benchmark components	6
Figure 4 - Wall power measurements for vCCAP system as clock frequencies are varied	7
Figure 5 - Wall power measurements based on varying clock frequencies	7
Figure 6 - Mapping vCCAP performance to CPU Frequencies	9
Figure 7 - Example of automated frequency scaling for vCCAP data plane	10
Figure 8 - Power measurements mapped to vCCAP throughout demand	11
Figure 9 - Software elements and logical control over the hardware.....	13
Figure 10 - Consolidation of vCCAP workloads.....	14
Figure 11 - Consolidation factor calculations mapped to 24-hour demand curve	16
Figure 12 - Power measurements mapped to vCCAP through demand	17
Figure 13 - Demand curves for fixed, mobile, and enterprise traffic over a 24-hour period	18

List of Tables

Title	Page Number
Table 1 - Calculating expected bandwidth per SG as frequencies vary	9
Table 2 - Throughput for consolidated vCCAP data plane downstream instances	14

Introduction

Network utilization is dependent on time, varying according to the hour, day, or week and reflecting activity spikes due to special events (e.g., sports games, news conferences). Figure 1 shows a typical profile of aggregate subscriber traffic throughout the day. This is just one example, but the key thing to note is that the demand fluctuates over some time period. In this case, one can see that there is approximately an eight-hour period of prime-time network activity with 16 hours of relatively low usage. In fact, primetime growth is outpacing average traffic growth, and thus driving CapEx spending for new access equipment to meet this demand.




Figure 1 - Primetime driving peak network demands

It is important to consider whether current network infrastructure has the capability and flexibility to adapt to these changing needs and conditions. Can it provide the lowest cost-per-bit, while always meeting the real-time demand of users? This paper provides a road map of options to help with the planning and deployment of a next-generation access network that takes these requirements into account.

Fixed appliances used to deliver network functionality (switches, CCAPs, EPCs, firewalls, etc.) are hardcoded with certain features and capabilities despite a clear underutilization of their compute, network, and storage resources during large parts of the day. However, with the emergence and maturity of SDN and NFV, the network architect can take a more intelligent path and design a flexible system reactive to the needs of both users and operators.

In this case, network infrastructure can be seen as a flexible entity with behavior and parameters that can be optimized based on real-time technical and business needs. In other words, the infrastructure will have a “state” at any given time that you can control and manage with the right hardware and software, as discussed in the following sections.

This paper specifically focuses on optimizing power usage for a virtual Converged Cable Access Platform (vCCAP) data plane VNF running on standard COTS servers, but this research is applicable across any type of network function. Figure 2 shows the general network transformation from purpose-built network appliances to virtual software functions running on a common server-based infrastructure.

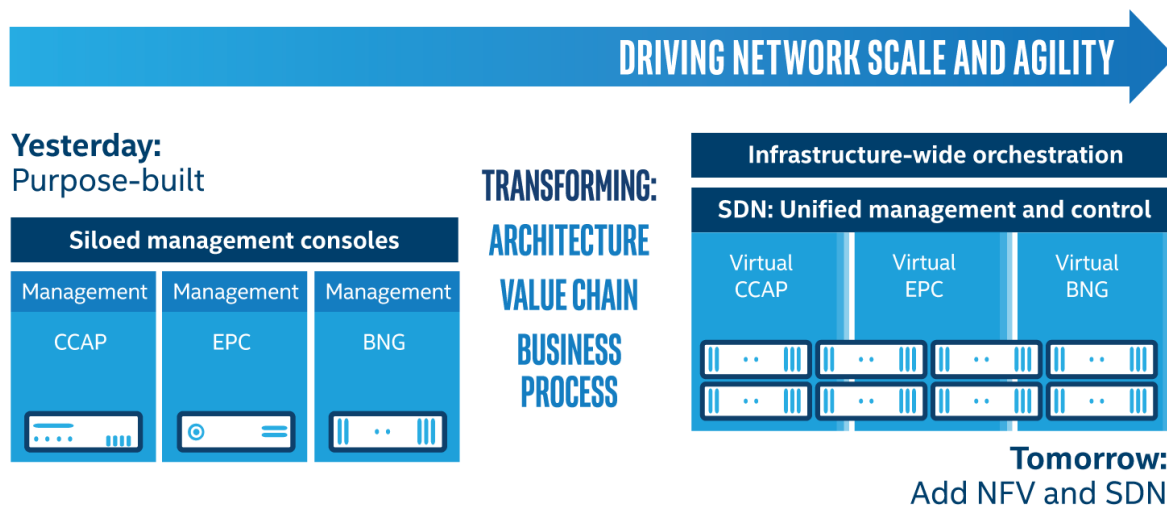


Figure 2 - Moving access functions to standard COTS servers

The power measurements discussed below can be understood in two ways: first, as literal savings in operational costs to pocket, and second, as a proxy for headroom in the infrastructure resources to perform other tasks. Further, the effort has pushed the state of the art for demonstrating best practices in hardware and software to realize the most efficient use of one's access infrastructure for the vCCAP or any other type of network function.

Our research provides a proof point of how a modern, container-based network functions virtualization infrastructure (NFVI) can be used to scale up or scale down various vCCAP operational parameters of the VNF itself or the platform it runs on through the collection of platform and network telemetry; a framework for decision-making through static or machine-taught policy engines; and, ultimately, the system automatically recognizing and reallocating resources to best meet its functional and operational requirements.

The learnings from this work will allow an operator to understand the potential of their SDN- and NFV-based network infrastructure to enable overall greater business agility and reduce network TCO. In fact, this paper presents a continuum of options that can be deployed in order to maximize the value of CapEx spent to upgrade from the legacy appliance model to one based on agile software.

- The first section discusses power management features native within COTS servers and techniques that can be used to save power when demand is low.
- The next section evolves this basic approach, allowing intelligent orchestration layers to make use of workload consolidation across a pool of servers to bring average server power down to the absolute minimum.
- The last section look to using this downtime in one access application to do other revenue-generating work, in particular it focuses on the potential for looking across multiple fixed, mobile, and enterprise needs with the goal of maximizing total resource utilization, thus maximizing the value of the investments of the operator.

Equipping your SDN- and NFV-based access network properly can provide the capability to:

1. Pay as you grow where labor costs for upgrades are high—for example, installing based on future-forward requirements and waiting to activate until the right business case is defined
2. Reduce OpEx by pocketing power savings when demand is lower
3. Use these savings to support more system maintenance and/or security (e.g., equipment failure detection or prediction, optimized redundancy schemes, security scans)
4. Enable new commercialized and next-generation services, such as VR/AR, smart cities and homes, autonomous driving, and IoT, on the same server infrastructure

Active Power Management per Server

In order to be able to take advantage of opportunities in NFVI for cost savings and/or to have the flexibility to otherwise use it to deploy and run new workloads, key NFV features must be part of the solution. To illustrate these elements and show how they can be used to optimize the needs of a particular network or business, this paper looks at how to minimize the fixed and dynamic costs of running a vCCAP data plane on an individual server.

It is important to note that standard commercial off-the-shelf (COTS) servers running NFV software already include a suite of power management tools. These can be used to increase or decrease the clock frequency of many hardware elements in the system to put them in lower power modes or turn them off altogether. While these capabilities are generally available, they are not always fully utilized to reduce OpEx. This paper focuses on the server's general ability to change the core and uncore frequencies of the CPU as it will give the operator the greatest “bang for the buck.” The core frequency generally applies to the cores themselves¹ (ALU, FPU, etc.) and the L1 and L2 caches, and the uncore frequency applies to shared resources, such as the LLC, integrated memory controller interfaces, and a few other tightly integrated internal units.

Initial benchmarking takes a look at vCCAP data plane performance measured as throughput against the AC wall power consumed by the server. This enables the modelling of a system based on dynamic performance and, conversely, power demands. For example, as bandwidth needs go down, compute, memory, and network elements of the system can run slower and still keep up, and with those slower clock frequencies, one will see a proportionate reduction in power.

Figure 3 shows the main components of the servers used to run these tests. This is just one sample server configuration out of many possible in the market. The key items to note are that the server has two Intel® Xeon® SP processors, each with 20 cores with a default frequency of 2.4 GHz for both the cores and the uncore logic. Each core is running one instance of the vCCAP data plane VNF and handles one Service Group (SG) with its data traffic coming in on one of the twenty-four 10 GbE ports available in the system. The maximum throughput per core (i.e., per SG) when the system is running at the default

¹ Note that individual core frequencies can be varied independently of each other allowing different cores to run at different frequencies. For the purpose of this paper all cores are set to the same frequency.

frequencies² for the configuration is 6.4 Gbps. Note that more than one core could be used to saturate each 10 GbE port, but it was not necessary for this testing. Details of the vCCAP dataplane software implementation and test environment can be found in the published paper: “Maximizing the Performance of DOCSIS 3.0/3.1 Processing on Intel® Xeon® Processors.”

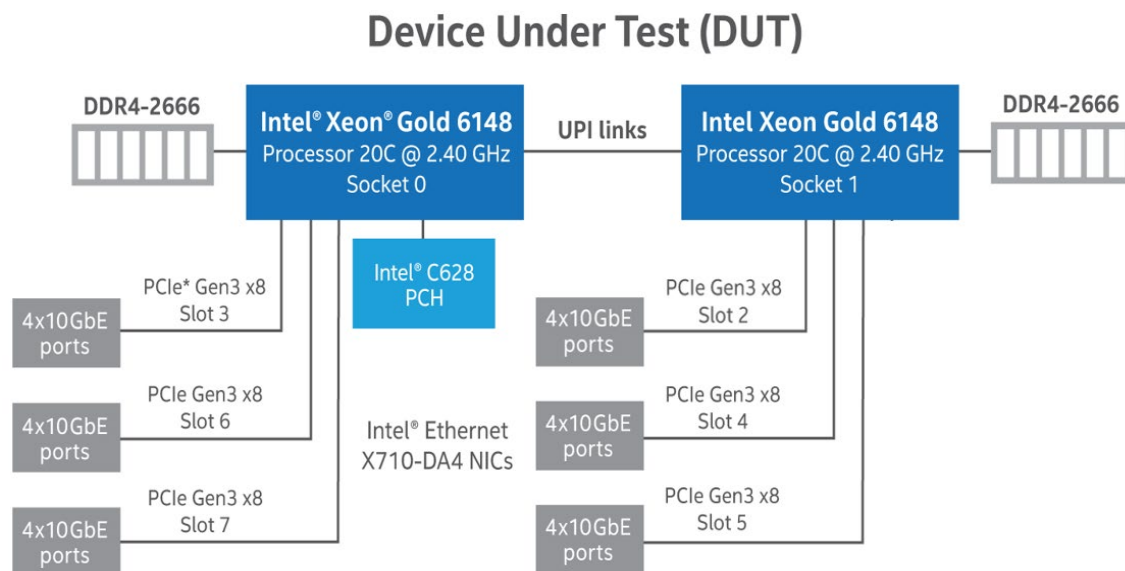


Figure 3 - Test server and benchmark components

Figure 4 shows server power measurements as one adjusts the core and uncore speeds when running 24 vCCAP instances. The core frequency is shown on the x-axis and the AC wall power measured for the system is the y-axis. Two different lines are used to represent the measurements, while the uncore frequency was held at either the default of 2.4 GHz (orange) or reduced to its minimum of 1.2 GHz (blue).³ In other words, the chart shows how much active power is required for the given server to pass the maximum amount of traffic possible per core at the specified clock frequencies with zero packet loss for all 24 service groups. There is clearly a linear relationship between the clock frequencies and the power consumption of the server.

² The default frequency is also known as “Base” frequency. It is possible to achieve greater performance by increasing the frequency of the cores above the default frequency using turbo modes. This feature and its application to NFV is outside the scope of this paper.

³ The uncore frequency can be adjusted anywhere in the range 1.2 Ghz to 2.4 Ghz.




Figure 4 - Wall power measurements for vCCAP system as clock frequencies are varied

Figure 5 summarizes the power measurements made for the 2RU server described above while varying the clock frequency of three different entities: (1) the cores running the vCCAP data plane workload; (2) the cores running other applications; and (3) the uncore.




Figure 5 - Wall power measurements based on varying clock frequencies

The left side of Figure 5 shows the idle power for a given server when not running any applications and taking advantage of minimum core and uncore frequency settings. More advanced settings that could increase power savings but also affect the responsiveness of the system (e.g., to start up and instantly run a new workload) were not used. For our purposes, this measurement represents the baseline or “static” power per server. Note that while this research focuses on power savings available through frequency scaling of the CPUs, there are other components in the test system (memory, NICs, etc.) that contribute secondary levels of power. A future analysis could add these into the optimization model.

Moving right from the optimal idle power measurement, the other power measurements are bunched based on how many service groups are being handled by the server for that test. For example, the next four bars show the power measured in different frequency permutations when only one service group is being handled, and then the next four bars show the measurements for 12 service groups, and so on. The right-side bar in each of these bunches is the default power of the system (when the core and uncore frequencies are at their defaults). Conversely, the left-side bar shows the lower power possible when these frequencies are dialed down to their minimums.

As vCCAP data plane instances are loaded and running, there is a linear increase in the “dynamic” power of each system to account for packets being received by the network interface controller (NIC), sent to the cores and/or memory, processed, and then sent out to another NIC port. In this way, Figure 5 shows how the wattage demands increase across all clocking permutations.

This data establishes a couple of things. First, even the most basic power management features, like adjusting clock frequencies, do have a tangible effect on the power used by a server and therefore the OpEx of the system. Second, there is a maximum and minimum amount of power each server will consume for a given workload, depending on how fast various elements in the system are clocked. Of course, if the clock frequency of the cores running your vCCAP data plane is reduced, it will handle less throughput. But what is the derating factor and how do we map this all back to meeting the real-time demands of the network?




Figure 6 - Mapping vCCAP performance to CPU Frequencies

Figure 6 shows how throughput per SG per core was affected as the different clock frequencies varied. Table 1 below the graph takes the calculations further and summarizes the power per SG and throughput per SG across all core and uncore frequency permutations that were part of the testing.

Table 1 - Calculating expected bandwidth per SG as frequencies vary

Core Frequency	Uncore Frequency	Power per Service Group	Service Group Bandwidth
1.0 GHz	1.2 GHz	10 W/SG	2.7 Gbps/SG
1.8 GHz	1.2 GHz	13 W/SG	4.7 Gbps/SG
2.4 GHz	1.2 GHz	14 W/SG	5.5 Gbps/SG
1.0 GHz	2.4 GHz	13 W/SG	3.3 Gbps/SG
1.8 GHz	2.4 GHz	15 W/SG	5.3 Gbps/SG
2.4 GHz	2.4 GHz	17 W/SG	6.4 Gbps/SG

To select a specific example from the data above: if you reduce the uncore frequency from the default of 2.4 GHz down to 1.2 GHz (the green line in the chart) and also reduce the core frequency from 2.4 GHz to 1.0 GHz, then you can expect the vCCAP data plane VNF running on that core to handle about 2.7 Gbps of traffic. To put it another way, if the throughput demand of a given SG is only 2.7 Gbps, you can reduce your uncore frequency from 2.4 GHz down to 1.2 GHz and the core frequency from 2.4 GHz to 1.0 GHz. This reduces the server power needs by 40 percent—from 411 watts to 243 watts.

Consider that 40 percent power reduction across 10, 100, or 1,000 such servers at a given location making up the access infrastructure and it adds up to considerable cost savings!

Software Infrastructure for Active Power Management

In order to realize these savings, there needs to be software in the system that can automatically adjust the aforementioned frequencies in response to real-time system behavior. This will require the solution architect to define a set of Key Performance Indicators (KPIs) for the network; choose the system telemetry that best represents these KPIs; define policy and the associated actions engine to maintain the KPIs within a desired range; choose the right tools for the job; and figure out how to automate the process going forward.

Most Linux* distributions will include many of the tools you need to scale core and uncore frequencies at runtime, but you may have to implement new logic within your applications or create your own “glue” software at a higher layer to take advantage of them for maximum effect. For example, Figure 7 shows how a data plane application running in a Kubernetes-based NFVI can use the Data Plane Development Kit’s power management library to automatically detect opportunities to save power, while still meeting the required latency and throughput demands of the network operator.

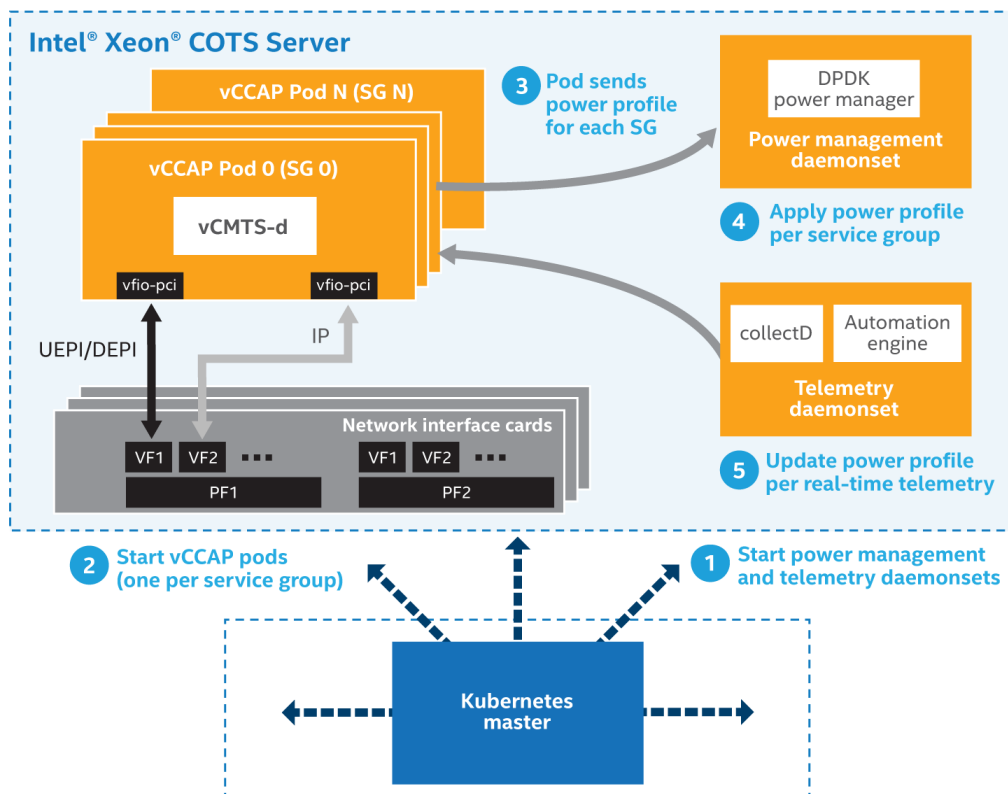


Figure 7 - Example of automated frequency scaling for vCCAP data plane

This example shows all the elements discussed earlier to implement a network infrastructure that can respond to real-time demands: system telemetry collection; an engine to make decisions based on that data; and then a harness to be able to execute those decisions with minimal to no operator input. Here, Kubernetes-based container orchestration and management infrastructure is used to deploy all of these elements onto a COTS server. Once they are in place, these elements are able to understand what is

happening in the system and adjust the core and uncore frequency of the platform according to a power management ruleset.

Looking back to our original demand curve, by deploying the frequency scaling techniques discussed above, a power curve similar to the one shown in Figure 8 can be achieved. The dotted red line at the top is the power consumed by the server running fully loaded at the default clock frequencies, and the green line is the power measured using optimal frequency scaling to accommodate the demand. There is a large amount of savings possible in this particular example, evident in the gap between the dotted red and solid green lines highlighted with the large arrow. Again, in those off-peak times, the server power usage is about 40 percent lower than the maximum. And over the full 24-hour period, the total savings accrue to approximately 33 percent relative to the default settings.




Figure 8 - Power measurements mapped to vCCAP throughout demand

The flat part of the green line representing the optimized power using frequency scaling techniques (approximately from the 0:00 hour to the 14:00 hour), shows a minimum server power (i.e., about 250W when all 24 vCCAP data plane instances are active) coming into play for these lower-demand parts of the day. It is beyond the scope of this paper, but knowing that the true minimum server power as shown in Figure 5 is about 176W⁴, there are other power management opportunities in the platform. For example, the vCCAP data plane software used in the benchmarking assumes that maximum performance is expected all the time and thus generates a lot of work polling for new packets on the network interface, whether they are actually there or not. Of course, in lower-demand parts of the day, the network driver

⁴ 176W is idle power of a system when no services are running the system, but cores and uncore are at base frequency. As stated previously, further power reduction is possible using other features outside the scope of this paper.

could be configured to reduce the amount of polling it does or moved to an interrupt-driven mechanism to further reduce system power.

These additional efforts would be rewarded with a possible further 30 percent reduction in the power usage (i.e., from 250W to the ideal of 176W). Alternatively, with the view that these power measurements are a proxy for excess resources in the system, the operator could decide to take advantage of this gap by running other applications on these servers “for free.”

In short, the data above definitively shows that there are real operational savings to be had if the SDN- and NFV-based solution for vCCAP or any other VNF has the hooks in place to frequency scale different parts of your system in response to demand. In fact, with an understanding of the particular demands of your network, along with the particular performance curves for the desired VNFs, one can calculate the OpEx savings of the system and drive some of that investment back into more powerful servers up front.

Further, Moore’s Law continues to bring down the fixed power costs of Intel® architecture-based servers, generation over generation, allowing more complex and intelligent power management features to become standard in the resulting hardware. In other words, as the performance per watt of the server CPUs and associated chipsets increases and new power management features within the silicon are developed, the platform idle power for newer equipment will be naturally lower.

The next part of this paper shows that there are even more ways to save power and/or use idle compute for running complementary applications by looking beyond the capabilities of an individual server and taking a pooled approach to network infrastructure through the use of smart orchestration tools.

Savings Through Orchestration

The next strategic approach starts by thinking of the network infrastructure as a pool of resources that can be managed in real time and not just stand-alone appliances to be individually controlled. To this end, one employs a full suite of software infrastructure and tools that allow the complete orchestration and management of all network functions and applications in such a way that in periods of low demand, workloads can be consolidated onto fewer and fewer running servers. This allows the infrastructure to be optimized for the lowest possible operational power and/or creates the space to run complementary applications. This approach can be seen as an evolution of the one discussed in the previous section or pursued independently, as one starts moving from network appliances to SDN and NFV.

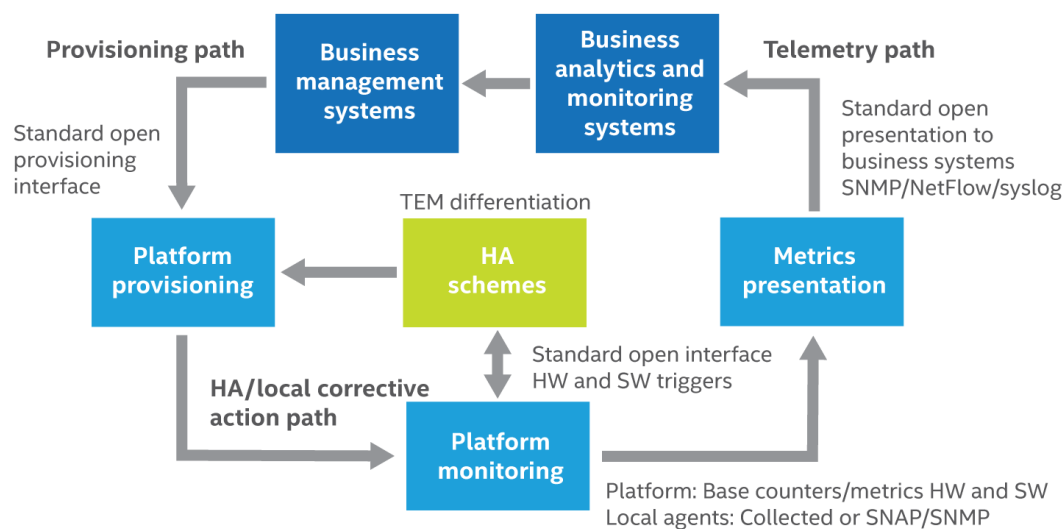


Figure 9 - Software elements and logical control over the hardware

Figure 9 shows the software elements required to create this type of environment—real-time telemetry collection, monitoring and analytics engines, business management and policy engines, and various action engines—holistically across all servers being used for access and edge service. With these elements in place, a fuller version of SDN and NFV is realized, where servers and attendant hardware are seen not as individual entities, but as a truly homogeneous pool of compute, network, and storage resources.

The model builds upon the calculations covered in the previous section, but focuses on reducing fixed power costs per system by consolidating work onto fewer servers. Each server introduces a fixed minimum power cost when running workloads. The idea is to reduce the cumulative fixed cost by powering only enough servers to meet the demand at any given time of the day. Of course, this only works in deployments where more than one server can be dedicated to the applications of interest.

First, total server needs are identified based on peak network throughput requirements and thus create a “pool.” Next, the capability is enabled through software infrastructure to be able to fully move applications to any available server in the pool. Finally, a set of functions must be added to be able to detect when certain KPI thresholds are reached and then react per operator policy. When demand is low, this type of system allows application consolidation onto the minimum number of servers to still meet demand and fully shuts down any that are not used (saving 100 percent of the power they would use just to be “on”).

To illustrate this point, Figure 10 shows that when demand is high for a particular virtualized application, like a vCCAP, workloads may need to run exclusively on three different servers to achieve the necessary peak performance. However, as real-time demand for that application drops, it may be possible to consolidate application instances onto one server. This degree of control allows the overall system to operate at power levels that almost exactly mirror the demand curve.




Figure 10 - Consolidation of vCCAP workloads

Calculating Potential Efficiency using Consolidation

Expanding the model starts with an assumption that all SGs will be serving the exact same throughput requirement for a given time of day. While this is unlikely, it makes the model calculations much simpler and illustrates the potential gains of workload consolidation. Taking a time-of-day example at 5:00 a.m., we see from the traffic capture data that each SG needs to support an aggregate bandwidth of 1.31 Gbps. This 1.31 Gbps represents 20.5 percent of the maximum possible per SG per core.

In theory, a core should be able to handle four service groups at 20.5 percent of peak demand with ease given the cumulative demand is 82 percent of peak. However, consolidating more than one application (i.e., vCCAP data plane instance) onto a core adds overhead due to context switching and for low-level resource sharing (e.g., cache). Consequently, there is a reduction in the per-core throughput. Table 2 shows the performance degradation measured as up to five vCCAP data plane instances are deployed to a single core.

Table 2 - Throughput for consolidated vCCAP data plane downstream instances

vCCAP Instances Per Core	Throughput Per Instance	Total Throughput Per Core	Performance Degradation per core
1	6.40 Gbps	6.40 Gbps	N/A
2	2.72 Gbps	5.44 Gbps	15.0%
3	1.56 Gbps	4.68 Gbps	26.9%
4	1.04 Gbps	4.16 Gbps	35.0%
5	0.72 Gbps	3.70 Gbps	42.2%

After taking the overhead into account for the original 5 a.m. example above, a single core can support up to three of those 1.33 Gbps SGs at the same time. Figure 11 repeats the calculation and maps these “consolidation factors” to the whole 24-hour network demand curve. The larger the consolidation factor, the fewer servers are required to meet the network demand and, thus, the greater amount of power savings possible.

In reality, each SG may have different throughput needs at any given time. In this case, the ideal solution would have the telemetry layer monitor how the servers are tracking demand over some timescale, and then have a machine learning-driven algorithm solving for the problem of packing SGs into the minimum amount of servers necessary. To return to our early morning example, at 5:00 a.m., if two SGs require the aforementioned 1.33 Gbps but two others require 2.74 Gbps, then at least two cores will be needed. In this particular case, each core will be allocated one 1.33 Gbps SG and one 2.74 Gbps SG, such that the maximum throughput required for all SGs per core does not go above the 6.4 Gbps maximum. All of these calculations can be handled in real time by a utility or an automated decision-making engine built into the orchestration software layer.

So while frequency scaling provides some very compelling power savings on the individual server level (as described in the previous section), each server will retain a minimum power requirement on the order of 176 watts^{5,6} simply to be “on.” By having a pooled view of resources being controlled under the same management domain, the number of servers active to deliver a particular service, like vCCAP, can be scaled down to allow for great savings in average power usage per SG. In this way, the system effectively breaks the minimum power barrier at the individual server level by amortizing the fixed costs of power supplies, memory, NICs, etc., by packing more SGs per core (and hence needing less servers to do the job).

⁵ Minimum power varies from system to system and depends on several factors including but not limited to selected CPU, memory quantity and size, plugin cards, and storage.

⁶ Lower idle power is possible where advanced power features are available and used.

Platform consolidation to achieve ideal power




Figure 11 - Consolidation factor calculations mapped to 24-hour demand curve

For instance, if you were to consolidate four servers down to one server running at maximum frequency, one server might be maxing out its power profile, but there are also three other servers now turned off and contributing zero watts to the overall power draw of the system. The most savvy software infrastructure should actually implement algorithms to save power at both the individual server level and pool level in order to deliver maximum value for the infrastructure.

Getting back to the original demand curve and using the simplified assumption of homogenous SG needs and the calculations above, average server power requirements for a vCCAP data plane were plotted over a full day of traffic (see Figure 12). Again, the dotted red line is the power of the servers running the full vCCAP load with no power management enabled, the green line is the power measurements for the frequency scaled case described in the previous section, and the dotted blue line is the measurements when the fully orchestrated consolidation scheme described in this section is used. This last approach allows the network operator to truly tune the operational expenses represented by power to a minimum, while still meeting the needs of the users.




Figure 12 - Power measurements mapped to vCCAP through demand

The previous section introduced the frequency scaling opportunity for power savings at the individual server level; this is highlighted with the top blue arrow in Figure 12. This section took a “pool of servers” view toward the goal of saving power (or reusing the compute represented by that power), and broke the per-server minimum power barrier to achieve up to 85 percent savings in the lowest demand periods of the day. This is highlighted with the bottom dark blue arrow in Figure 12.

This work represents the start of what is possible and it is expected that the industry will continue to bring down the TCO of NFV-based infrastructure as more telemetry-gathering, decision-making, and automation layers are refined and added to deployed solutions. There will be up-front costs to develop or buy these new capabilities, but the data above shows that it will be made up many-fold over the lifetime of the equipment.

Consolidating Fixed, Mobile, and Enterprise onto the Same Infrastructure

As discussed, the power savings outlined in this paper can be seen as literal OpEx savings for providing the electricity to run the virtualized access infrastructure. The other view of the power savings metric is that it is also possible to reuse the spare compute, networking, and storage resources to run other access or enterprise applications on the same servers. In other words, by transforming the headend or central office to a distributed data center, a network operator can realize the full vision of network functions virtualization shown in Figure 2, where a common infrastructure can be used to support whatever

functions are demanded by the operator and users in real time. There are several different approaches for this.

Some applications may be run opportunistically at any time if they do not have particularly high technical or business demands. For example, back-end machine learning-assisted analysis of network or user telemetry data. Others may have particular and unwavering demands of their own (similar to vCCAP), and thus can use the same type of demand curves and power/performance calculations described in this paper to harmonize how they share compute and other resources.




Figure 13 - Demand curves for fixed, mobile, and enterprise traffic over a 24-hour period

For example, Figure 13 shows that generally fixed, mobile, and enterprise workloads have different time-of-day demands. These usage measurements came from internal research for a Next-Generation Central Office (NGCO) that aims to support a mix of all of these services on the same COTS server infrastructure as an alternative to deploying parallel fixed function appliances with uncoordinated management facilities. In this way, expanding the intelligent orchestration and management concepts discussed in the previous section to comprehend more than one application at a time will allow the same equipment to be dynamically reused across all of these workloads, providing a cost-effective, flexible, and future-looking approach to network facility architecture.

Of course, Figure 13 shows that while there are times where the demand is complementary (i.e., toward the start and end of the day), there is a bulk of time in the middle where at least two of the applications have high requirements at the same time.

This presents an area for further study, as supporting multiple demanding applications may require breaking up the notion of “one big pool” and instead utilize sub-pools dedicated to running only one type

of mission-critical application. Or a sub-pool may be designed to run any number of instances of application one and two, but not application three. These decisions will be based on the characteristics of each application, such as whether all applications require access to common hardware elements, have demand curves that coincide (or are mirror images), etc.

Regardless of the particular implementation, it is clear that if the servers can be shared effectively using intelligent orchestration solutions, then inefficiencies can be driven out of individual server use and thus allow for fewer capital expenditures up front to support this wide range of services.

The details of implementing these strategies, and therefore the server needs, can be determined beforehand using theoretical or empirical data and then codified via a hardcoded policy engine. Or, an automated decision-making and policy-modifying engine can be deployed that uses machine learning algorithms to optimize the behavior of the system in real time. Strategies may also be modified, based on the availability of certain hardware features. For example, workloads with real-time requirements may initially be segregated onto a special sub-pool of servers, but it may be possible to use a single pool if the CPU features can provide resource determinism for workloads. The value of static versus automated decision-making is left for further study.

The beauty of SDN and NFV is that this is all defined in software, so that new sources of telemetry can be enabled and new decision models deployed. Essentially, all system operation can be managed in a flexible, agile manner, if the right amount of intelligence is employed in the solution.

Conclusion

SDN- and NFV-based solutions in the access network promise benefits over legacy hardware appliances in the realms of flexibility, manageability, and scalability. However, this paper highlights that while these systems might have the potential to deliver on these promises, not all solutions are created equal.

Using the vCCAP data plane VNF as a representative workload for other types of access technologies that could run in COTS servers, this paper outlined a continuum of options to reduce power usage as demand rises and falls over a given 24-hour period. At a minimum, the solution should take advantage of the frequency scaling of the CPU cores and uncore logic, as they are the largest contributors to both power and performance of the system for this type of workload. Savings per server can be on the order of 33 percent overall, with a peak savings of ~40 percent relative to the default server configuration!

In addition, if the solution adds intelligent orchestration and automation frameworks to the NFVI that can autonomously determine opportunities to consolidate the vCCAP (or other VNF workloads) onto the fewest servers possible to meet real-time demand, then an additional average of ~21.5 percent savings can be unlocked, with a peak of 45 percent for a total of ~57 percent lower power usage on average over 24 hours and upwards of ~85 percent in times of minimal demand. The power per SG can track demand very closely and thus provides the lowest TCO for the equipment.

The power it takes to run the access equipment can also be seen as a proxy measurement for the ability to run other workloads on the server. This allows an operator to take better advantage of the fixed costs of the equipment and/or create a foundation for new services on the same infrastructure. To this end, an evolution of the consolidation technique would account for not only a single workload (e.g., the vCCAP), but instead look across all the access and enterprise application needs of the network operator in order to take full advantage of the equipment at all times. In this way, a truly intelligent solution could find the maximum flexibility and savings across all infrastructure requirements.

By understanding the power and performance impact of specific features that can be made part of the NFV-based access network, operators can make better decisions for designing and deploying next-generation infrastructure to support the ever-increasing data throughput needs of their users over time and be able to nimbly respond to competitive pressures in a cost-effective manner. The ultimate aim is to unlock the potential of SDN and NFV to improve the user and operator experience, save costs, and create a foundation for new workloads and services in a world that requires constant evolution.

Bibliography & References

1. <https://www.theatlantic.com/charts/H1tALGE4>

Estimated results reported above may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information about benchmarks and performance test results, go to www.intel.com/benchmarks.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© Intel Corporation

Automation Opportunities for Subscriber Management in Cable Television

A Technical Paper prepared for SCTE•ISBE by

Greg Nicholson
Chief Operating Officer
NuTEQ Solutions, provider of GOCare
gnicholson@nuteqsolutions.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	4
1. Subscriber Management – A Continuum	4
2. Automation – Systematic Migration of Customer Interactions	5
2.1. Phase One	5
2.1.1. Event: Outage / Function: Select Outage Notification	6
2.1.2. Event: Service Call / Function: Real Time Logistics	6
2.1.3. Event: Service Suspension / Function: Billing Mitigation	6
2.1.4. Event: CPE Management / Function: Self-Service	6
2.2. Phase Two	7
2.2.1. Event: Upsell / Function: Intelligent Targeting	7
2.2.2. Event: Customer Care / Function: Personalization	7
2.3. Advanced Services	7
3. Outcomes – Modeled and Measured	7
3.1. Risk Model.....	7
3.2. Impact Model.....	8
3.3. Return on Investment (ROI) Model	8
4. Measured Performance – Implementation Data	9
4.1. Call Center Impact	9
4.2. Missed Appointments.....	9
4.3. Billing Collections	10
4.4. Targeted Promotion	10
4.5. Reduction In High Friction Experiences	11
5. Architecture of ASM	11
5.1. Platform	11
5.2. Scope	12
5.3. Workflow.....	12
5.4. Ontology	13
5.5. Knowledge Base	13
6. Implementation – Building A New Customer Channel	14
6.1. Planning	14
6.2. Make Or Buy	15
6.3. Delivery Mechanism	16
7. Evolution of Capabilities	17
Abbreviations	18
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 - Subscriber Activities	4
Figure 2 - Behavior and Mitigation Measures	5
Figure 3 - Call Center Impact.....	9
Figure 4 - Missed Appointments Impact	10

Figure 5 - Billing Collections Impact.....	10
Figure 6 - ASM Platform Abstract	11
Figure 7 - Workflow of Subscriber Initiated Inbound Inquiry (P2A).....	12
Figure 8 - Workflow of Event Driven Operator Initiated Campaign (A2P)	12
Figure 9 - Sample First Level Ontology Categories and Event Sequence	13
Figure 10 - Generic Knowledge Base for Event Management	13
Figure 11 - Make Versus Buy.....	16

List of Tables

Title	Page Number
Table 1 - Expected Profit At Risk	8
Table 2 - ASM Impact Opportunities	8
Table 3 - Impact Range Per System Size	8

Introduction

The lifecycle of a subscriber relationship is composed of identifiable and controllable interactions. Automation models can incorporate all events that compose the experience through a comprehensive business logic and ontology. Systems can elicit or preempt subscriber behaviors through defined measures designed to control outcomes. This paper describes the promise of automated subscriber management (ASM), the migration process, system design considerations, target outcomes and results of operator implementations. ASM can deliver improved customer satisfaction and retention, remove costs, promote and deliver service enhancements. Properly designed and implemented, an automated platform will continuously improve cable operator performance.

Content

1. Subscriber Management – A Continuum

The continuum of customer interactions is known and addressable. Operator business models are based on a lifecycle of delivering services over an expected duration. When subscribers perceive the value of the product exceeds its price, the relationship is maintained over the expected period. Perceptions include indirect contributors that create friction, or additional cost - altering the value/price relationship and affecting the expected outcome. Through subscriber management, the contributing factors to maintaining the expected customer lifecycle are addressed.



Figure 1 - Subscriber Activities

Subscriber management addresses a finite set of manageable variables. Services are provided over a known infrastructure; subscribers consume and pay for the product. Problem instances create friction in the relationship resulting in financial costs; negative impact on value perception and outcomes such as termination. When expected performance is compromised, subscribers will behave in a predictable manner according to available options.

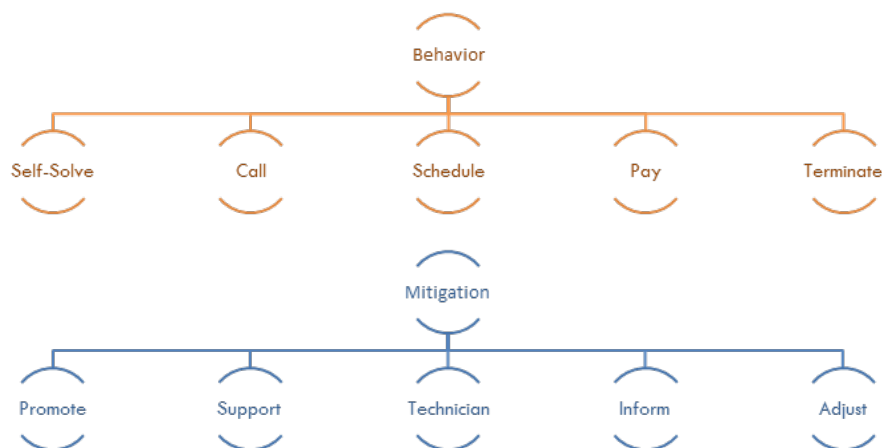


Figure 2 - Behavior and Mitigation Measures

Operator mitigation measures are also discrete - primarily the provision of instructional, temporal, and financial information; enhanced through personalization and timeliness.

The effectiveness with which these dynamic, controllable and measurable variables are managed determines the customer experience and resulting lifetime value of the subscriber relationship.

2. Automation – Systematic Migration of Customer Interactions

Operators have invested significantly in traditional means of subscriber management. Call centers and service fleets address subscriber needs. This support structure reliably handles subscriber events such as 100% monthly calls, 7% monthly technician service call rate, 10% missed truck rolls/repeats, and 15% late pay. Significant expenditure is dedicated to these services, but satisfaction ratings remain low and monthly subscriber termination averages 2.5%.

Within the modern plant, network monitoring and detection capabilities are comprehensive and accessible. The data, systems and capabilities required to automate subscriber management exist. Every element of the operator – subscriber relationship is monitored, tracked, and archived. Business rules, problem mitigation strategies, marketing promotions, and resolution procedures are determined. Human resources are committed but the frequency, scope and scale of activity limits fast, customized and optimal results.

The business logic of customer interactions can be codified into software systems. Automation can improve scope, scale and efficiency; and act in real time - handling millions of simultaneous interactions personalized to the individual. Algorithms can mine and monitor data to discover correlations, predict events and reliably preempt high friction experiences through contextual communications.

Automated subscriber management (ASM) can improve customer satisfaction, retention and profitability. The migration from existing processes is described.

2.1. Phase One

Initial candidates for automation include routine events that drive a high volume of frictional interactions. Examples include outages, service calls, billing events, and troubleshooting. Described are generic events, the automated subscriber management function and projected results.

2.1.1. Event: Outage / Function: Select Outage Notification

ASM discovers outages through integration with network monitoring software, or inbound notifications from subscribers and parses the affected segment through mapping, such as per affected node. Historical data indicates likely individual subscriber behavior and constructs mitigation measures based on business rules and preexisting content. Informational notifications are delivered in real time to select subscribers regarding resolution at a predetermined frequency or based on continued inbound inquiries. Business rules can incorporate financial considerations and adjust activity accordingly. Preemptive notifications reduce high friction experiences for subscribers. Informational needs are eliminated, resulting in removal of call center volume.

2.1.2. Event: Service Call / Function: Real Time Logistics

ASM interfaces with field management, service scheduling and subscriber data. It requests, confirms or reminds the account holder at a frequency determined by business rules and review of historical data. It also monitors real time status of service technician and delivers schedule options to the subscriber awaiting arrival and communicates change request to field management and service technicians. Upon completion (e.g. closed work order), it polls the subscriber to confirm issue resolution and reports exceptions to field management and service tech while on premise. It allows the subscriber to engage the technician on their timetable and terms and reduces wait time, friction, missed appointments, repeat truck rolls, and call center activity.

2.1.3. Event: Service Suspension / Function: Billing Mitigation

ASM monitors billing data for late payment and engages automated mitigation measures based on business rules and preexisting content, such as payment scheduling. It accepts inbound inquiries from subscribers; constructs and delivers content regarding account status, payments due, scheduled suspension, and payments received. It also mines customer historical data for indicators of likely termination and creates and communicates customized alternatives based on business rules. Communication activities deliver faster collections, avoidance of call center calls, and reduction of high friction subscriber experiences. Business logic can incorporate pricing and plan adjustments to impact profits and preempt account suspension or subscriber termination.

2.1.4. Event: CPE Management / Function: Self-Service

ASM integrates with network monitoring systems to identify equipment problems or receives inbound notification from subscribers. It correlates network and premise equipment measurements with the subscriber problem and creates resolution measures based on business rules and preexisting instructional content. It also reviews historical data for the account. It instructs subscriber activities (cycle), provisions automated procedures through network operations center (NOC) control systems (reset), interfaces with monitoring software to confirm issue resolution, and polls the subscriber for outcome. Business rules can escalate to advanced resources such as visual engagement (how to) or provision service calls to address complex situations. Enabling self-service supports avoidance of high friction experiences, reduction in call center volume, and faster time to resolution.

The above cases illustrate high impact opportunities to identify and interrupt a **cascading cycle** of problem instances leading to costly, high friction experiences and avoiding reduced perception of value/price which may contribute to a termination decision.

2.2. Phase Two

Further candidate interactions include revenue generation through upselling products and services. Across all activities and events, the opportunity exists to ‘wow’ customers with intuitive insights.

2.2.1. *Event: Upsell / Function: Intelligent Targeting*

ASM mines the account and NOC systems for conditions that qualify for a promotional offer, according to predetermined business rules. (Example: speed or quality of service (QOS) complaints compared to network capacity improvements.) Polling can further determine subscriber preferences, price sensitivity and demand. Per criteria, curated content is delivered. Upon affirmative response the ASM provisions order to account and billing system and manages authentication and confirmation. Resulting upgrades, enhanced value/price perception, elimination of high-friction solicitations, and avoidance of traditional promotional costs improve lifecycle profitability.

2.2.2. *Event: Customer Care / Function: Personalization*

ASM integrates with the customer account and historical data and marketing campaign strategy and content. It then polls selected subscribers with opt-in offers for recommendations, reminders, special offers and personalized content. It creates subscriber preference profiles, monitors systems, and mines account data for conditions that meet marketing criteria or subscriber requested content. It also delivers personal notifications, polls subscriber on utility of content and adjusts based on feedback. As subscribers opt-in for personalization features, opportunities for merchandising of services or content for new revenue generation increase, and the system delivers increased value perception and loyalty.

2.3. Advanced Services

Automation systems can enhance the functionality, productivity and value of advanced, data-intensive services such as **security, in-home medical monitoring, and programmable smart homes.**

In summary, the modern systems available to the operator can monitor infrastructure, equipment and subscriber data to maintain a real time customer profile. Automation delivers preemptive information, striking the perfect balance of software-provisioned mitigation management and human interaction.

3. Outcomes – Modeled and Measured

Automation of subscriber management can deliver improved customer retention and financial performance for the operator. Presented are impact models and performance data from ASM implementations.

3.1. Risk Model

The table below depicts profits of a subscriber lifecycle, the cost of common subscriber behaviors in a non-automated environment, and impact (loss avoidance) of service upgrades and retention.

Table 1 - Expected Profit At Risk

Event	\$ Net Profit	% Net Profit
Subscriber target lifecycle at 5 Yr \$100 ARPU at 40% net profit	\$2,400	100%
Subscriber contacts call center once/month at \$6/contact	(\$360)	(15%)
Subscriber requests one service visit per year at \$100/roll	(\$500)	(20%)
Subscriber terminates one year earlier than target	(\$480)	(20%)
Subscriber avg. 10% more likely to upgrade service by \$20/month	\$120	5%
Profit Impact Opportunity	\$1,460	60%

3.2. Impact Model

ASM impact opportunities and financial gains to a 1M-subscriber system are illustrated.

Table 2 - ASM Impact Opportunities

Event	ASM Impact	Description	Annual Gain
Call Volume	20% Reduction	100% monthly call rate reduced by 20% at \$6 cost/call	\$14,400,000
Missed Appointment	80% Improvement	100% annual rate with 10% missed at \$100/roll, reduce missed by 80%	\$8,000,000
Customer Retention	10% Improvement	25% annual churn reduced by 10% at \$480 profit/year	\$12,000,000
Upgrade	5% Take Rate	5% customers speed upgrade by \$20/month	\$12,000,000
Total			\$46,400,000

3.3. Return on Investment (ROI) Model

One and five-year cumulative financial gains are projected for a range of system sizes.

Table 3 - Impact Range Per System Size

# Subs	100,000	1,000,000	5,000,000	10,000,000
One Year Impact	\$4,640,000	\$46,400,000	\$232,000,000	\$464,000,000
Five Year Impact	\$23,000,000	\$232,000,000	\$1,100,000,000	\$2,320,000,000

4. Measured Performance – Implementation Data

Background: Results are based on operator trials. Activation of the subscriber base grew to 25% over the measured period. Activated subscribers were provided a self-service menu for inquiry/response and operators delivered event triggered mitigation messages via simple messaging service (SMS). Results are period-to-period changes versus non-activated subscribers.

4.1. Call Center Impact

When mitigation measures were deployed against identified triggers of subscribers placing a call to the call center, the operator reported an **average reduction in actual call volume of 26%** year to year 2017/2016. Importantly, 70% of activated subscribers stopped calling altogether. As shown in the trending graphic, impact on call center reduction increases as automation begins to change the way subscribers interact.




Figure 3 - Call Center Impact

4.2. Missed Appointments

When mitigation measures were deployed against identified triggers of missed appointments, the operator reported an **80% reduction in missed appointments** over the course of a year for activated subscribers. The 80% improvement virtually eliminates missed appointments for activated subscribers.




Figure 4 - Missed Appointments Impact

4.3. Billing Collections

When mitigation measures were deployed against identified segment of subscribers that historically pay late, the operator experienced a **50% reduction in the collection period** - the time between subscriber entering late pay status to payment.




Figure 5 - Billing Collections Impact

4.4. Targeted Promotion

In a system trial, analytics were conducted leading to segmentation of a subscriber base targeting accounts with a predicted propensity to upgrade their service to a higher speed/bandwidth offering. Business rules were based on system characteristics (availability based on recent construction) and not historical

purchasing or other private subscriber data. Messaging was delivered by SMS to subscriber phones. The automation platform delivered a **6% take rate on the promotional offering** within two days versus past direct mail conversion of <1%. The system completed automated billing and provisioning functions through an operational support systems (OSS) interface.

4.5. Reduction In High Friction Experiences

Operators noted the removal of high-friction experiences such as complaint calls, missed appointments and repeat billing inquiries. The on-demand provisioning of useful account information received positive feedback from responding subscribers. Although every interaction provides subscribers the option to opt-out of the service, few chose that path.

Reduced calls translate 1:1 to removal of high friction experiences for subscribers.

In summary, automation of subscriber management applied in the field has improved operator performance through reduced call volume, improved appointment attendance, reduced truck rolls, faster collections and higher promotional take rates. The removal of high friction experiences is inherent in problem avoidance.

5. Architecture of ASM

5.1. Platform

An automated services workflow supports both inbound inquiry/response and event driven, operator generated communications through a common platform architecture.

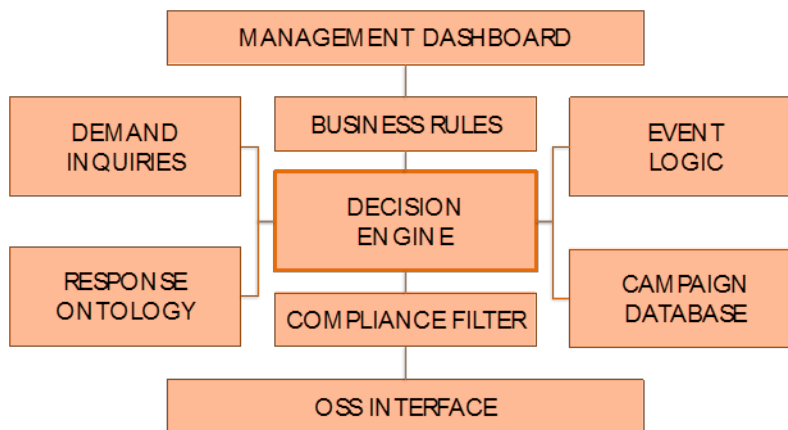


Figure 6 - ASM Platform Abstract

Decision Engine: monitors and collects inputs for activity generation. Inputs are solicited through an OSS systems interface, directly from inbound subscriber inquiries or autonomously generated from a status monitor of candidate events that meet predetermined thresholds in the business rules.

Business Rules: govern action taken for inbound inquiries and system generated qualifying events; built from a detailed review of business practices, processes and strategy.

Demand Inquiries: inbound requests generated by subscribers are assigned to the response ontology for manufacture of response content, amended with OSS-derived data.

Event Logic: qualifying events are assigned to event logic ontology for manufacture of outbound content per the campaign database, amended with OSS-derived data.

Compliance Filter: necessary rules and processes that support security, privacy and Federal Communications Commission (FCC)/carrier regulations related to utilized messaging mechanisms and communication channels.

Management Dashboard: a repository of event elements and transactions, activity tracking, outcomes and impact reporting, serving as the user-interface for platform management.

By abstracting the workflow and architecture the system is designed for maximum flexibility. Systems interfaces are defined through an application layer. Considerations for scale, adaptation and extension are inherent in the centralized logic, shared functions and resources of the platform.

5.2. Scope

An adaptable platform will support the chosen scope of ASM. Business requirements determine the complexity of the system. Selected and defined events and mitigation strategies require workflow, ontology and knowledge bases available to support automated decision-making and activity.

5.3. Workflow

ASM addresses events originated internally, through subscriber requests, or other sources. An abstraction of the person to application (P2A) and application to person (A2P) process is described in Figure 7 and Figure 8.



Figure 7 - Workflow of Subscriber Initiated Inbound Inquiry (P2A)

When an inbound inquiry is made, such as a simple text inquiry generated from a subscriber phone, business logic is applied, and the response ontology assembles predetermined content relevant to the inbound request, amended with system data; resulting in a response event. Information is delivered through chosen mechanisms.



Figure 8 - Workflow of Event Driven Operator Initiated Campaign (A2P)

When a threshold is met in the status monitor, an event is created, and business logic engaged. Target subscribers are qualified, campaign ontology enacted, and relevant content is amended with data drawn from supporting systems.

5.4. Ontology

ASM events are categorical entities with dependent activities. The system accesses capabilities, data repositories, content and feedback according to the ontological model. Sample categories and sequences are shown in Figure 9.

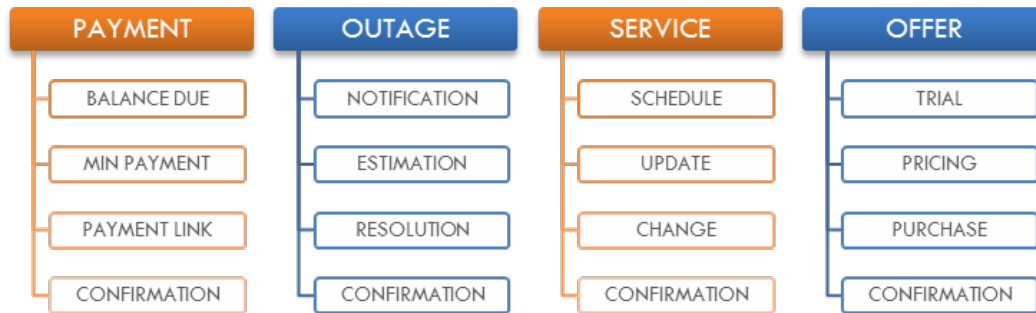


Figure 9 - Sample First Level Ontology Categories and Event Sequence

5.5. Knowledge Base

Decisions for each sequence are dependent on the knowledge base represented in the system, requiring uniform access from different data sources, such as those represented in the matrix in Figure 10.

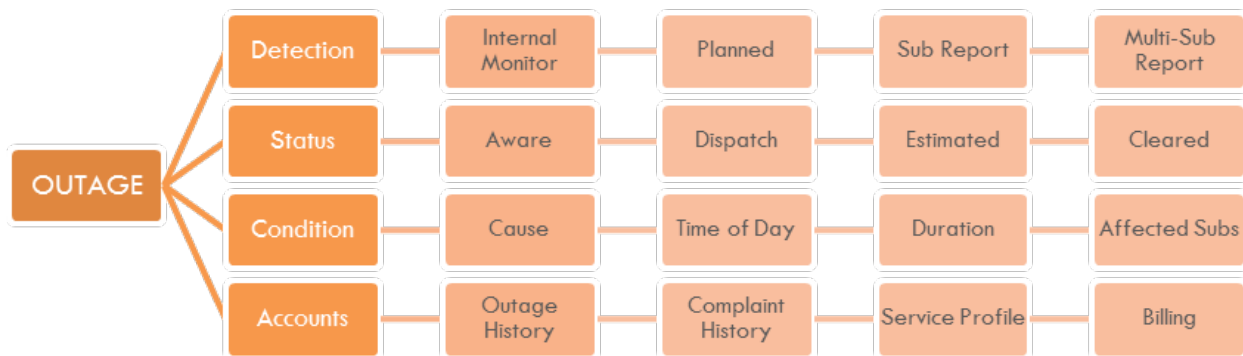


Figure 10 - Generic Knowledge Base for Event Management

Thus, the scope of decision-making is based on available data, process ontology and workflow. Routine events and mitigation strategies have predetermined content constructed from highly available systems and rules – for instance, notifying affected subscribers of an outage. Higher-level ontologies can tailor responses, such as providing prioritized treatment based on profile, history and rules.

Emerging technologies will drive a company's competitive differentiation through customer service. Software robots perform routine business processes and make simple decisions by mimicking the way that service agents interact. Companies can automate entire end-to-end processes, with humans typically only managing exceptions. Expect to see continued focus on automating repetitive rules-based tasks. (Forrester 2017)

Conceptual frameworks can incorporate subscriber satisfaction and retention analysis through direct polling or evidence-based correlation of hidden contributors. At scale and in real time, inferential insights can be presented for management consideration and performance optimization.

The capabilities of an automation suite depend on available resources. Real time performance data on network services and equipment exist for the most part. The addition of models and knowledge bases incorporating problem instances, consumer behavior, and operator mitigation opportunities represent a more complete ontology of subscriber management. Acting upon a codified representation of these scenarios generates results and measures that enable learning and performance improvement.

6. Implementation – Building A New Customer Channel

6.1. Planning

Preparing for automated subscriber management is a process involving disciplines of business strategy, marketing content development, software and data access design, regulatory compliance and financial measurement.

Design: A flexible, open, scalable architecture can incorporate new capabilities. Application programming interfaces (APIs), content databases and analytics applications should be adaptable and extensible; as systems, network and consumer premise equipment evolve.

Business Requirements: Requirements are derived from a logistics and business process review of subscriber interactions. Business logic includes rules and qualifications for system functionality. Logic governs qualification parameters, event descriptions, thresholds, activity triggers and content selection.

Ontology: The representation of relationships between events and actions depends on business logic and rules. This core element powers the automated decision-making process, such as mapping events to supporting subsystems.

Content: Content development will relate to the scope of business logic and ontology. The system will draw from a predetermined content database to generate mitigation measures or other subscriber communications. In a two-way system, ASM detects distinct inquiry nomenclature, relates response content, and amends with custom data based on account characteristics. For event driven outbound campaigns, content is selected based on the ontology of events that trigger mitigation measures to a qualified segment of subscribers.

Systems Availability and Data Access: ASM requires integration with operator systems according to a data extraction and API logic. ASM integrates with billing, scheduling, field services, marketing, analytics and other applications.

Compliance, Privacy and Security: Best practices must be applied. Delivery mechanisms will incorporate FCC and carrier regulations. Data exchange from systems carrying customer proprietary network information (CPNI) data must be properly parsed.

Reporting: Historical data is captured to support indexing of outcomes compared to past performance. Normalizing data for extraneous events supports effectiveness tracking, learning and optimization of the system.

Opt-In Management: A recruitment campaign facilitates activation of subscribers through advertising media or directly through CSR origination, advance registration via terms of service, or mobile SMS messaging.

Limitations: Technology capabilities limit the scope of automation. While process logic and ontologies can reliably manage distinct interactions, consumer psychology cannot be fully represented - requiring escalation mechanisms to engage human resources. Operational considerations include resistance to disinvest in existing processes.

6.2. Make Or Buy

Benefits of internal development include insights into the performance characteristics and subscriber dynamics of the system. Maintained, there will be no dependency on third party resources or external financial commitment.

Considerations for Make:

- Long-term commitment of talent from multiple disciplines.
- Committed engineering resources for maintenance and extension.
- Financial, marketing, operations and customer care oversight.
- Continuous tracking, reporting and optimization required to maximize benefits.
- Time to reach deployment may alter the ROI of a make v. buy decision.

Contrarily, outsourcing inherently represents a low-cost implementation; available content ontology; fast integration with disparate systems; existing security, privacy and compliance measures; and perspective on program success factors across multiple operators and extensive subscriber interactions.

Considerations for Buy:

- Broadband infrastructure and systems experience.
- Integration with common billing, OSS and field management software.
- Open and adaptable programming interfaces.
- Available process ontologies and content inventories.
- Demonstrable efficacy in implementation, compliance and impact reporting.
- Capabilities in security, privacy and FCC regulatory adherence.

Typical in the make/buy evaluation is time to market. In Figure 11, savings associated with an immediate deployment are compared to a twelve-month development delay.




Figure 11 - Make Versus Buy

Leveraging internal knowledge and resources to maintain maximum control and security, while leveraging vendor expertise and availability represents a hybrid combination of make versus buy.

6.3. Delivery Mechanism

Alternate communications channels are available. ASM can support email, set top, custom-built apps or standard phone-based SMS delivery; each with its advantages and challenges.

95% of Americans have a cell phone, 90% home-internet, and 90% email access. (Pew Research)

Email: This ubiquitous medium is highly available to subscribers. It is comprehensive in capabilities when hypertext markup language (HTML)-enabled, resides on computing devices with browser access to cloud services and can be customized to include additional interactive features. **Challenges include:** Frequency of engagement, lack of immediate notifications and reaction. Clutter and whitelisting also present engagement barriers.

Apps: Dedicated applications provide extensive functionality and leverage device capabilities. Rich features can extend to other operator objectives. **Challenges include:** Cannot automatically opt-in a subscriber base without download. App builds require native implementations and maintenance updates. Consumer interest has waned, and usage statistics show low engagement. Notifications must be enabled to support timely response.

Research Perspectives:

- **Majority of US consumers download zero apps per month.** (comScore)
- **Only 36% of apps are retained after one month, only 11% for a year.** (comScore)
- **77% of users never use an app again 72 hours after installing.** (comScore)

Set Top: Where available to subscribers, notifications delivered through the set top box appear in a familiar context. Utilization of the set top equipment is costless. **Challenges include:** Navigation imposing on a viewing medium. Subscribers may engage away from the television screen. Any footprint required on the set-top will present challenges of available processing capabilities.

Simple Messaging Service: SMS messages utilize the notification layer, which elicits immediate response. Consumer engagement is high and frequent. Operators can activate subscribers through common business practices. **Challenges include:** SMS messages must comply with FCC and carrier regulations. Subscribers must provide their cell phone number and opt-in to a notification program. SMS is text-based, with a per message cost.

Research Perspectives:

- **Over 80% of American adults text, making it the most common cell phone activity.** (Pew Internet)
- **Text messages have a 98% open rate, while email has only a 20% open rate.** (Mobile Marketing Watch)
- **90% of all text messages are read in under 3 minutes.** (Connect Mogul)
- **It takes the average person 90 minutes to respond to email, 90 seconds to respond to a text.** (CTIA)

Operators may choose alternate or combined approaches to implementing an automated subscriber management system. Make versus buy decisions will be determined by resource availability, time sensitivity and cost.

7. Evolution of Capabilities

“Companies will anticipate needs by context, preferences, and prior queries and deliver proactive alerts, relevant offers, or content. They will become smarter over time via embedded artificial intelligence.”

Forrester: 2017 Customer Service Trends: Operations Become Smarter And More Strategic

Automation of subscriber management is an investment in the future. Initially an ASM platform addresses routine, high-volume interactions. Focused efforts reduce financial risks and create opportunities to provide new utility to subscribers. Subsequent phases rely on adaptation of business rules and content to increasingly reliable audience profiling, segmentation and management based on the relative state of each subscriber in the continuum of service. Business logic and ontology can incorporate new event variables and mitigation measures dynamically. Each interaction is programmatically tallied, harvesting data for further learning and improvement.

Contemporary automation opportunities target identifiable and deterministic processes where conditions can be monitored and decision support data is available. Deployments have demonstrated efficacy. Capabilities of data mining, analytics, and machine learning are enabling autonomous, intelligent management systems. On the experimental frontier, ASM can engage subscribers and discover new ways to improve the relationship.

The subscriber experience will improve accordingly. Today, operators can eliminate the majority of high friction interactions through automated provisioning of just-in-time information and support. As systems evolve, subscribers will enjoy the convenience of transparency and control; and experience smart,

personalized service. The economic impact to operators will expand from efficiency to profit production and subscriber retention.

Investing in automation will ensure the constancy of innovation in the broadband services industry.

Abbreviations

A2P	application to person
API	application programming interface
ARPU	average revenue per unit
ASM	automated subscriber management
CPE	consumer premise equipment
CPNI	customer proprietary network information
FCC	Federal Communications Commission
HTML	hypertext markup language
NOC	network operations center
OSS	operational support systems
P2A	person to application
ROI	return on investment
SCTE	Society of Cable Telecommunications Engineers
SMS	simple messaging service
QOS	quality of service

Bibliography & References

Pew Research Center, Mobile Fact Sheet, February 2018.

Pew Research Center, U.S. Smartphone Use, April 2015.

comScore, The 2017 U.S. Mobile App Report, August 2017.

Mobile Marketing Watch, SMS Marketing Wallops Email, July 2018.

Connect Mobile, Texting Statistics, March 2013.

TextMarks, 6 Benefits of Text Messaging, February 2015.

Forrester, 2017 Customer Service Trends, January 2017.

Cable's Role in the 5G Evolution

A Technical Paper prepared for SCTE•ISBE by

Erik Gronvall

VP Strategy and Market Development

CommScope

501 Shenandoah Dr

Shakopee, MN 55379

952.403.8691

erik.gronvall@commscope.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
The Need for Wireless Densification.....	3
The Need for a More Efficient Solution	4
Power	5
Backhaul.....	6
Site Acquisition.....	7
Conclusion.....	8
Abbreviations	8
Bibliography & References.....	8

List of Figures

Title	Page Number
Figure 1 - Cisco Visual Networking - Mobile Data (Cisco, 2018).....	3
Figure 2 - Wireless Network Migration.....	4
Figure 3 - Revenue vs Traffic in Wireless Network.....	4
Figure 4 - Low Power Small Cells.....	5
Figure 5 - Power Consumption in HFC	6
Figure 6 - RAN Backhaul Diagram.....	6
Figure 7 - RAN Front, Mid, Backhaul Methods	7
Figure 8 - Small Cell Siting Methods.....	7

Introduction

For years, mobile network operators (MNOs) have tried to deploy small cells to boost coverage and increase capacity. Deployments have been limited because of the costs involved in putting sub-5-watt radios everywhere, and they end up going back to the macro site and increasing capacity there. The cost issues involve power, backhaul, and real estate. However, with 5G there will be a need to move to a more densified network and MSOs have all the components available to them to assist the MNOs in deploying these networks. This paper will discuss the challenges and trends from the mobile networks and how MSOs can take advantage of their existing assets to enable the next wireless generation.

The Need for Wireless Densification

Mobile traffic continues to grow at an accelerated rate, and each generation of wireless technology has fueled this growth. Every 10 years there is a new generation in wireless networks. As we are entering the 5G era, it offers traditional MSOs an opportunity to participate in this next generation. Each “G” has offered new applications to consumers and businesses. 2G was focused on digital voice, 3G on mobile browsing, and 4G on mobile video. 5G is promising three main benefits: enhanced mobile broadband, Internet of Things (IoT), and ultra-low latency. The other promise each generation has had is an investment in the physical network. 4G drove fiber to the cell sites to support the mobile backhaul requirements, and while MSOs were able to participate in providing this backhaul, the 5G network offers even more opportunities in supporting the wireless rollout.



Figure 1 - Cisco Visual Networking - Mobile Data (Cisco, 2018)

Just like every generation of wireless technology 5G will place new requirements on the network. The macro cell network does not look sufficient to supply capacity, latency or connections for the 5G network. The only method of solving this is to densify the cell sites. Architecturally this means more cells at the building and street level throughout cities and residential neighborhoods.

Much of the new wireless spectrum used for 5G will be at a higher frequency to allow for greater capacity, however it will also necessitate smaller cells due to the shorter distance the wavelengths will travel effectively. There will also be change at the macro sites. Centralized and cloud RAN (radio access network) will drive more equipment at some sites and more fibers between sites. Regardless, the next generation of wireless network will run on fiber and will present an opportunity for those that have fiber and the ability to use it.




Figure 2 - Wireless Network Migration

The Need for a More Efficient Solution

Wireless network operators are facing a key problem as it relates to operating the wireless network. The amount of data provided has continued to grow at an exponential rate while the average amount paid for wireless access has only increased slightly.




Figure 3 - Revenue vs Traffic in Wireless Network

This drives a need to increase efficacy in deploying wireless services. There are three main challenges in deploying a cell site:

- Power,
- Backhaul, and
- Site Acquisition

Each of these have prevented the widespread use of small cells in the past and must be more economically solved for 5G to succeed in the future. This paper will take a closer look at each of these areas and how MSO networks can solve these challenges.




Figure 4 - Low Power Small Cells

Power

The powering needs for small cell solutions vary based on the size and desired performance of the cell. Larger “Small Cells” will require full metered drops from the power grid at the cost of thousands of dollars. These sites will support multiple bands of wireless spectrum and sectors. However, there is another set of small cells that will require less than 120 watts.

MSOs with their HFC network are well positioned to provide power to these types of small cells. Typically, 15-amp service at 90 VAC is available, and industry Pareto analysis shows an average usage of only 7-8 amps. On average that leaves 600 watts of unused power, more than enough for wireless APs whether Wi-Fi or LTE/5G small cells distributed along the plant, which may operate at lower than 50 watts each. This alone would significantly help in the economics of deploying small cells. For the larger “small cells” the MSO community has more experience obtaining and managing power in the OSP than any other type of network operator.




Figure 5 - Power Consumption in HFC

Backhaul

The new CRAN deployments bring different requirements on the backhaul network as the base band unit is being moved from the site to centralized locations. This centralized RAN, creates three types of backhaul with different demands: front, mid and backhaul.



Figure 6 - RAN Backhaul Diagram

The RU/AAU is the radio unit, the DU is the distributed unit of the baseband controller, and the CU is the centralized unit of the base band controller. 5G allows for splitting the stack of the baseband into two units like remote PHY. Generally, there are four methods of deploying cell sites, each with different advantages and disadvantages.

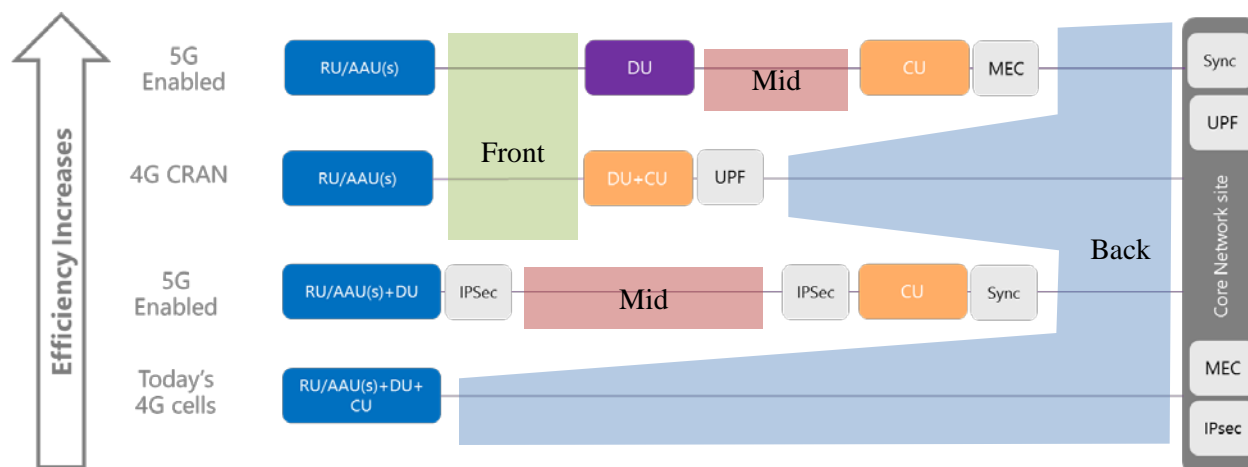


Figure 7 - RAN Front, Mid, Backhaul Methods

While dark fiber can be used in each area and is often the preferred method for MNO's, other technologies can be leveraged with the HFC network to provide front, mid and backhaul to small cells. WDM technology can be used to optimize the fiber used in front, mid and backhaul. The MSOs are more familiar with WDMs and able to track the wavelengths better throughout their network. In some midhaul applications PON or DOCSIS™ can be leveraged. Even more applications for backhaul can use PON or DOCSIS™. The MSO community has fiber and coax in appropriate places to be able to provide the connectivity to cell sites.

Site Acquisition




Figure 8 - Small Cell Siting Methods

The last challenge in deploying small cells is placing the radio and antenna. MSOs have access to many of the poles and other sites. There are solutions that place radios and antennas on the top of poles. There are also smaller radios that are capable of being placed on strands simplifying the deployment process further. In locations where the plant is underground, other solutions can be leveraged to provide the vertical height for mounting antennas. Considerations include integrated light poles and active equipment cabinets. Either way, the MSO community has vast experience in working with municipalities in siting equipment in the OSP.

Conclusion

The drive to 5G is underway, but the MNO's will need help solving the economic problems presented by the densification of the wireless network. The networks operated by the MSO community contain many advantages in deploying small cells for the 5G and 4G densification. The three main challenges in deploying small cells are power, backhaul, and site acquisition. Each of these can be provided using existing plant or provided by the MSO community.

Abbreviations

AP	access point
HFC	hybrid fiber-coax
SCTE	Society of Cable Telecommunications Engineers
RAN	Radio Access Network
MSO	Multiple System Operator
PON	Passive Optical Network
RU/AAU	radio unit in wireless network
DU	distributed unit of the baseband controller in a wireless network
CU	centralized unit of the baseband controller in a wireless network

Bibliography & References

Cisco. (2018, July 1). *VNI Global Fixed and Mobile Internet Traffic Forecasts*. Retrieved from <https://www.cisco.com/c/en/us/solutions/service-provider/visual-networking-index-vni/index.html>

Capacity and Technology Considerations in DAA Backhaul Deployment Strategies

A Technical Paper prepared for SCTE•ISBE by

Fernando X. Villarruel

Architect

Cisco Systems

5030 Sugarloaf Pkwy, Lawrenceville GA 30044

770-236-1385

villarf@cisco.com

Martin Mattingly

Technical Solutions Architect

Cisco Systems, Inc.

5030 Sugarloaf Pkwy, Lawrenceville GA, 30044

770-236-1338

mattinm@cisco.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Background	4
Architecture Comparison.....	5
1. Backhaul Option (A), “Direct Connect”	6
2. Backhaul Option (B), Field Aggregation Router (FAR)	7
2.1. FAR, Initial Uplinks.....	7
2.2. FAR, Long Term Uplinks.....	7
3. Backhaul Option (C), Muxponding	8
4. Physical Trunking and Distribution.....	8
5. Componentry Comparison	9
6. Architectural Comparison.....	10
6.1. Configuration	10
6.2. Usage of optical signals	10
6.3. Implementation.....	11
6.4. Uplink Bandwidth	11
6.5. Multicast and Unicast Bandwidth	12
6.6. Converged Access	13
7. Coherent Optics	13
8. Traffic Engineering and Backhaul Capacity	14
8.1. Estimating Uplink Bandwidth Capacity	14
8.2. Engineering Backhaul Capacity	15
8.3. Connecting the Uplink.....	19
Conclusion.....	20
Abbreviations	21
References.....	21

List of Figures

Title	Page Number
Figure 1 - High Level Remote PHY Architecture	4
Figure 2 - Remote PHY Backhaul Architecture Comparison	6
Figure 3 - Transition from Analog Transmission to Remote PHY	9
Figure 4 - Sample bill of Material for Field Router/Muxponder	9
Figure 5 - Architecture Options Comparison Pictograph	10
Figure 6 - Physical Connectivity of Remote PHY System	12
Figure 7 - DOCSIS Service Groups and Multicast Bandwidth.....	13
Figure 8 - Utilized Bandwidth per CCAP Chassis	14
Figure 9 - Provisioned versus Utilized Bandwidth per CCAP Chassis	15
Figure 10 - Subscriber Usage Over Time, 40% CAGR	16
Figure 11 - Aggregate Derived DOCSIS Service Group Capacity Over Time.....	17
Figure 12 - Backhaul Capacity Needed for Field Aggregation Router, 12 and 24 RPD's Subtended.....	18
Figure 13 - Most Likely Scenario for Needed Backhaul Capacity, 12 and 24 RPD's Subtended.....	18

Figure 14 - Initial Uplink Configuration	19
Figure 15 - Secondary Uplink Configuration	19
Figure 16 - Third Uplink Configuration	20

Introduction

This white paper compares the benefits of several architectural options for the Distributed Access Architecture (DAA) backhaul in the context of bandwidth growth over time. Some of the specific topics covered include the networking and optical implementations needed to address DAA backhaul, routing and TDM framing in hubs and HFC nodes, methodology for estimating needed capacity, concurrency, implementation of direct detect and coherent optics, and cable's new-found synergies with standard bodies. After reading this white paper, the cable operator will be able to compare the various architecture options and decision-making process for the deployment and long-term evolution of DAA.

Background

In the world of DAA there are at least two types of digital endpoints, remote PHY (RPHY) and remote MAC-PHY, (RMAC-PHY.) From the networking perspective, they both have the same function and they are both point to multipoint signaling that terminate at 10 Gbps endpoints. With a closer look, there are certain efficiencies concerning multicast that favor R PHY. Thus, in this paper we use RPHY as our main example, pointing out relevant differences with RMAC-PHY when necessary.

Figure 1 shows a summary for the RPHY architecture. It also helps us to define the Converged Interconnect Network (CIN) as the Ethernet/IP network between the packet cores and the RPHY nodes, drawn as the shaded cloud. We note that DOCSIS and other packet core payloads are encapsulated within an L2TPv3 pseudo-wire and then in IP and Ethernet layers. This process makes any subscriber specific data effectively invisible to networking elements thus making it fully addressable via switching and routing principles. It is also important to note that the packet cores and RPHY nodes have network facing Ethernet client interfaces, which mean that the most direct path for transmission is via client-to-client connections, not unlike what could be used in non-access networking such as in home or data centers applications.

The DAA backhaul is a subset of the CIN, in particular it is the physical link and method for aggregating multiple DAA endpoints. Backhaul methods then are distinguished by connectivity, packet processing and transport options. In practical terms, it is the way we leave the hub or headend and connect to a networking element that combines the signaling to and from RPHY nodes.

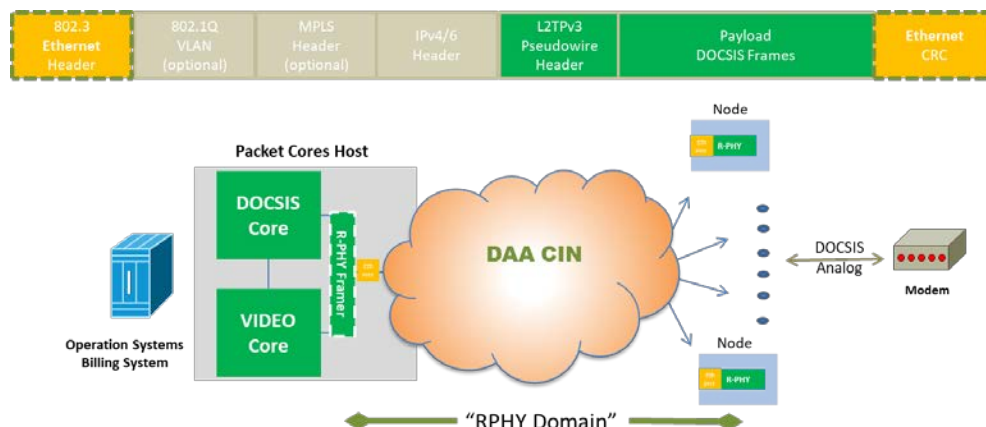


Figure 1 - High Level Remote PHY Architecture

Architecture Comparison

There are many variations in how to design the DAA backhaul. In Figure 2 we present several generalized options. Focusing on commonalities, we note that the northbound elements are a collection of packet cores, which execute subscriber management policy and create the data plane necessary. These can include the DOCSIS packet core, multiple video packet cores including broadcast, video on demand and switched digital video. There are also support packet cores that include such elements as out of band signaling and HFC RF monitoring tools. Generally, there could also be other service cores such as broadband network gateways (BNG) for PON or mobile. These packet cores can create point to point or point to multipoint sessions where signaling could be unique or shared as it passes through an initial series of hub routers towards their final destination, RPD's in the field. In particular from Figure 2, we note the existence of a layer of "core routers" whose job is to coalesce the signaling directly from packet cores. This routing layer uses 100 Gbps connectivity, with typical forwarding capacity nearing 1 Tbps with typical port count of 36 ports that can be used as either uplinks or downlinks. The packet cores themselves have direct 100 Gbps connectivity or are facilitated by an extra layer of routers with 1 or 10 Gbps connectivity for their uplink and 100 Gbps connectivity in their downlink.

Following the 100 Gbps connectivity router there is the existence of a dense 10 Gbps connectivity aggregation router. This router has 100 Gbps connectivity on the uplinks and 10 Gbps connectivity on the downlinks. Typically, these routers have two or four 100 Gbps connections along with 40 or 48, 10 Gbps connections. We call this the aggregation router because this router is the last logical connection at the hub or headend as the signal enters the outside plant. This multilayered approach of the remote PHY CIN is very similar to the spine – leaf architectures that are now prevalent in data centers and useful to facilitate the evolution to virtual packet cores. We note that between the 100 Gbps and dense 10 Gbps routers there could be photonic network as the packet cores might not be located at distribution hubs. Photonic equipment allows for a TDM aggregation of signals for transport of large bandwidths. We will discuss this type of transport gear again later.

These items account for commonalities. We now look at the differences in the options for the backhaul of Remote PHY Devices (RPD's).

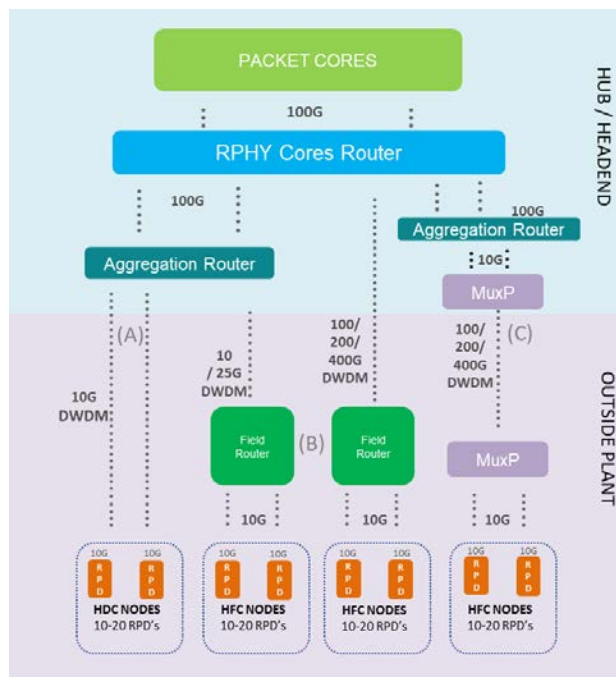


Figure 2 - Remote PHY Backhaul Architecture Comparison

1. Backhaul Option (A), “Direct Connect”

The connection labeled (A) in Figure 2 is part of an architecture description generally called “Direct Connect.” We call this Direct Connect because there is a direct connectivity between the dense 10 Gbps aggregation router and the 10 Gbps RPD endpoints. The main quality of this option is its simplicity and availability. In most cases, a collection of signals will be incident on one fiber along with an accompaniment of a Mux/Dmux to manage them. This architecture leverages 10 Gbps DWDM ZR optics, which is a description for optics that are wavelength specific within the 100GHz channels, as described by (G.694.1). The one variance is that the ZR optics used have to be thermally hardened to exist within the RPD enclosure that exists in the outside plant. This adds some complexity and some cost over the otherwise commodity structure of 10 Gbps ZR optics.

In situations where fiber is taken deep into the outside plant, the number of 10 Gbps endpoints can be 10 to 20 times the number of nodes that exist without DAA. As such, the number of DAA endpoints translate to a corresponding large number of DWDM wavelengths per trunk fiber. In Figure 3 we cover in more detail the evolution of the trunk fiber and DAA endpoint in a way that would be typical in end to end architectures.

It is worth noting that after the transition to DAA, we typically see a large disparity between the provisioned bandwidth per DAA endpoint and the actual utilization of bandwidth by the end user. In instances where multiple DAA endpoints are logically clustered together to create “service groups”, we often see that the 10 Gbps connection to each DAA endpoint is significantly under-utilized. In a later section we will also explore how obtaining a better understanding of provisioned versus utilized bandwidth can be an effective tool when estimating backhaul capacity.

2. Backhaul Option (B), Field Aggregation Router (FAR)

In Figure 2, we show two connectivity options labeled (B). Both have in common the introduction of an active networking element in the outside plant, a Field Aggregation Router, (FAR). The FAR has the task of facilitating several packet-processing functions to its subtending RPD's. From a topology perspective the FAR fits well at the same location where once was an analog node that now spawned multiple connections to RPD's. Southbound, there is 10 Gbps connectivity from the FAR to RPD's that are typically no further than 1 or 2 km away. This creates an opportunity for using lower cost 10 Gbps LR optics, which could be further de-rated for less than 2 km. These optics can be as low as one fourth the cost of the 10 Gbps ZR discussed in the previous section. Also, note that because of the short distance and the low link budget the necessity for thermally hardened optical components changes in scope and should not be a considerable cost adder.

Northbound of the FAR this architecture brings several dimensions of flexibility, and in order to appreciate the flexibility we explore the connectivity options available to the FAR by design. Routers are built around silicon chips whose inputs and outputs (I/Os) are well defined but with inherent flexibility. For example, a typical router in the service provider space might have a collection of 10 Gbps I/Os or 25 Gbps I/Os, all with equal access to the forwarding plane. How these I/Os are used is up to the designer of the router, with many variations possible. These transmission lines can be combined or down rated, when accompanied with the right media access control, (MAC). Four 25 Gbps lines can be combined to facilitate a 100 Gbps signal, or a 25 Gbps line can run 10 Gbps, or a 10 Gbps line can run 1 Gbps. In the case of the FAR, this allows for a range of options in northbound connectivity. We will see in a later section that pairing this flexibility with the expected capacity over time allows for pay as you grow scenarios.

2.1. FAR, Initial Uplinks

In Figure 2, the left portion of (B), we see that the FAR can have uplinks in speeds of 10 or 25 Gbps that leverage the existence of first generation hub aggregation routers. This allows a fine-tuned way to address the capacity needed for the uplink over time. In detail, this means purposing some of the 10/25 Gbps transmission lines of the FAR for the purpose transmitting the uplink, where the rest of the 10 Gbps, or 10/25 Gbps lines can be used for downlinks. Note, this type of uplink connectivity only makes sense to do with a few ports and maybe only at the beginning of the lifetime of the FAR, which should be a decade or more. This also makes sense as a transitional step if initially investments have already been made for aggregation routers in a hub. In this case, the introduction of the FAR enables fewer ports on the hub aggregation routers, as we will see in the backhaul capacity section.

2.2. FAR, Long Term Uplinks

In Figure 2, the right portion of (B), we see the FAR can have uplinks in speeds of 100 Gbps, and above. These uplinks represent the natural evolution of optics for higher transmission bandwidths and allows for the long-term transmission of signals to the FAR. These uplinks are made to address the challenges mentioned earlier in the direct connect section, where many lambdas were needed to address a collection of endpoints. In this case, the full bandwidth needed for all the RPD endpoints subtended by the FAR can be addressed with one lambda. From the perspective of the FAR, the usage of 100, 200, or 400G on one lambda can still employ the same routing fabric, if sized accordingly and accompanied with the necessary collection of transmission lines in the design. For example, 4 x 25 Gbps transmission lines facilitate a 100 Gbps optical PHY interface. If available, 8x25 Gbps transmission lines can facilitate a 200 Gbps optical PHY interface. Moving beyond first product implementations, the expected eventual addition of 50 Gbps transmission lines will facilitate higher bandwidths with even more simplified connectivity.

3. Backhaul Option (C), Muxponding

In Figure 2, option C, we see the insertion of a TDM framing layer between hub networking gear and the RPHY endpoints. Note that this practice would be new in the cable access but has been practiced in long distance optical transport for a long time. In the case of the cable infrastructure for DAA, the insertion of this transport mechanism makes perfect sense when hubs are collapsed to a more centralized location and the needed bandwidth between the hubs and head ends is on the order of many hundreds of Gigabits. Transport platforms currently have of up to 1 Terabit, with interfaces of up to 400 Gigabit, (Microsemi, 2017).

The common framing mechanism used for TDM solutions is called Optical Transport Network (OTN). This solution in essence takes in different client signals, possibly even with varied rates, and without any examination or manipulation of packets it stitches these different signals in the time domain and puts them on a signal at a much faster speed, a process that is reversed on the other side. We call the platforms that execute this function a muxponder. In practice, a common example would be a muxponder that would take twenty 10 Gbps signals and output a 200 Gbps signal.

4. Physical Trunking and Distribution

Figure 3 shows a practical approach for the evolution to DAA, particularly from what are typical starting points. In the scenario labeled “Analog Hub”, where the end to end signaling terminates at a hub, there are four downstream wavelengths and two upstream wavelengths used on a single fiber between the hub and legacy analog node. This is a byproduct of having an internally segmented node. In this type of deployment, the fiber is generally “point to point” from the hub to a physical legacy node and includes several spare fibers within the same sheath. In transition from analog to digital, one of the spare fibers can be used to provision 20 to 40 DWDM wavelengths to facilitate point to point DAA.

In the scenario labeled “Analog Secondary Hub”, two fibers carrying 16 analog wavelengths each from the Primary Headend are de-multiplexed to four groups of four wavelengths each. Each group of four wavelengths is used for downstream and upstream transmission between the Secondary Hub and the legacy analog node.

From Figure 3, we see that the transition to DAA can have several impacts to the connectivity. First there can be a replacement of the cores to a more central location where the connectivity from the new core position, like a primary headend or a data center, has to be accounted for. Because of the scale of the signal it is conceivable that the best solution for this link is in terms of TDM, OTN framing, allowing for multi-hundred Gig rate signals. Also, this implies that the hub location is a networking point. Which is the launch point towards a final aggregation point the FAR, facilitating the nature of replicated and concurrent signaling typical in cable plants.

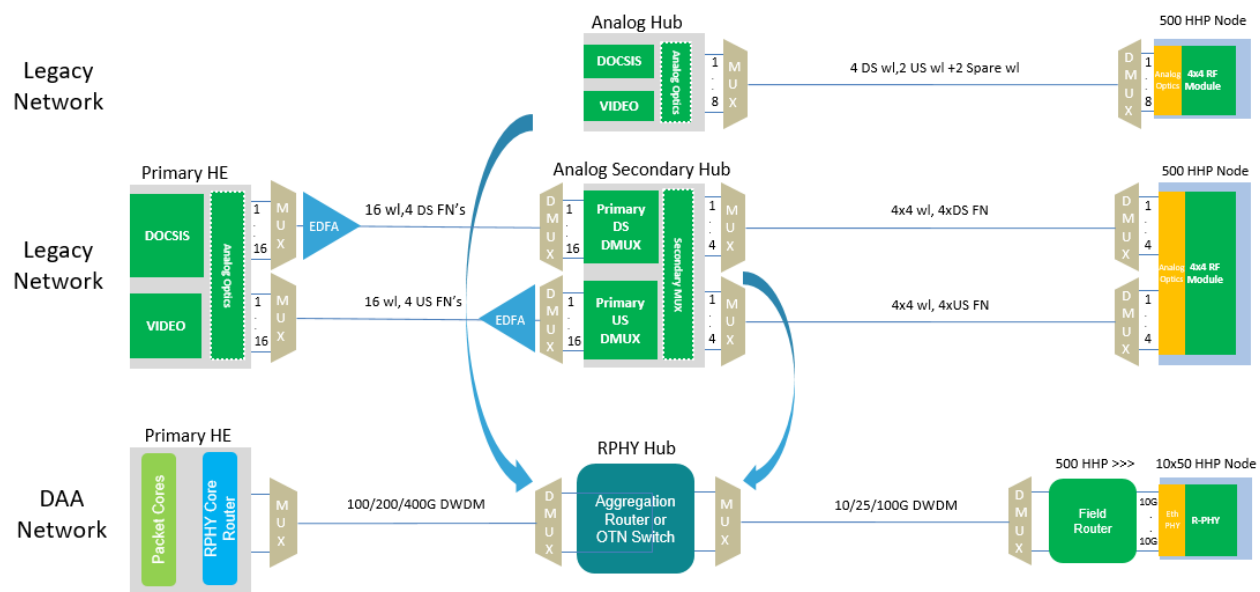


Figure 3 - Transition from Analog Transmission to Remote PHY

5. Componentry Comparison

Coming from the cable access world, we typically do not have a native understanding of what goes into these new remote digital technologies. Figure 4 represents the high-level componentry that make up the digital parts of both a router and a muxponder. Within the diagram we can see that there are some high-speed uplink optical inputs and accompanying PHY, lower speed optical downlinks and their accompanying PHY, and most importantly a function specific ASIC or FPGA along with a robust processor. The size of the silicon, the number of available gates drives its substantive differences in power consumption and functions. Interestingly, in the case of routers there is now a vast set of robust, third party off-the-shelf options to choose from that the industry refers to as merchant silicon. The silicon within the muxponder with functions such as an OTN framer and mapper can also function as an OTN switch, with increased gates and power. Finally, accompanying the ASIC or FPGA is a CPU processor complex that functions with the control plane and runs the operating system of the product.

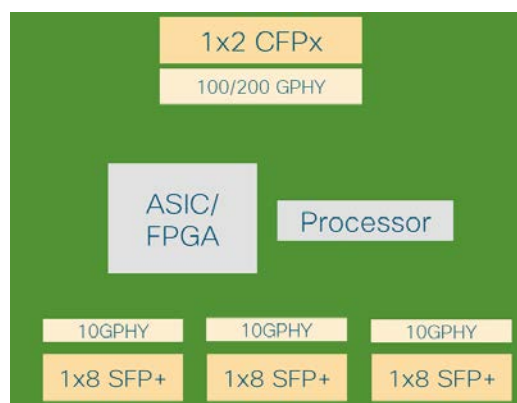


Figure 4 - Sample bill of Material for Field Router/Muxponder

6. Architectural Comparison

Figure 5 below, shows a comparison of features (in green) and challenges (in red) for each of the presented architectures. While we have already mentioned some items on the lists, it merits to compare them together by topic. Note that there is no one solution that is perfect for all situations but knowing how to evaluate them is a useful tool as these options make their way to market.

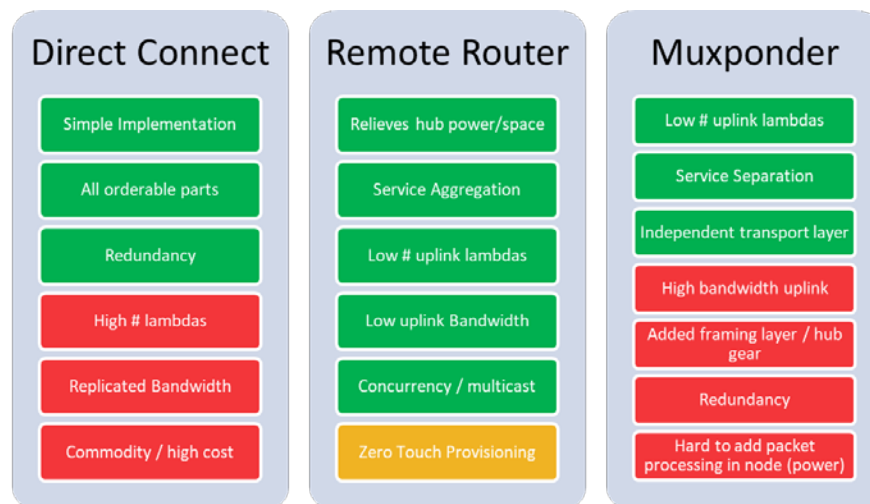


Figure 5 - Architecture Options Comparison Pictograph

6.1. Configuration

With regards to implementation, there is no doubt that the most straightforward option is direct connect. While adding another active device in the field adds complexity, it should be no more effort than configuring a remote PHY device. Since connection to packet cores already necessitates zero touch provisioning at scale, there is no new technology in packet processing being implemented here. Ultimately, any solution must be plug and play, with no settings intervention needed on site. Further, the target for actives in the field should require no manual interaction at all, even at a central location. The target should therefore be toward full automation of the service with general profile guidelines set by the network operator.

6.2. Usage of optical signals

It is worth considering the number of trunk lambdas used by each option. In the case of direct connect, there are as many lambdas needed as there are DAA endpoints. In the case of a FAR or muxponder there can be one lambda per trunk fiber servicing between 12 and 24 DAA endpoints. This allows the trunk fiber in a WDM environment to be used for other services. Note that this does put an added challenge on the description for optical uplinks. They must be able to overcome the passive losses of added WDM equipment while also being able to operate in a wavelength specific environment. These two items are not a given for high throughput long distance optics and are discussed further in section 7.

On the downlink side, there is an opportunity to use lower cost 10 Gbps optics that can easily adapt to the environment. Based on the lower cost of short reach optics, the total cost for a solution that introduces

remote aggregation should be much less than direct connect. This is an achievable goal, as we will show in section 8.3.

6.3. Implementation

The physical aspect of implementation is also worth considering. Note from Figure 5 that for all cases there is a transition from a large uplink signal to a breakout distribution for DAA endpoints. In the direct connect solution, this transition happens once at the hub as signal leaves to the access plant. In the case of the muxponder, this transition happens twice where the infrastructure for 10 Gbps connectivity happens both at the hub and at the node. In the case of FAR there is an option to do one of two things: In a pay as you grow scenario, as we will show in section 8.3, the uplink to the FAR can use signaling in terms of 10 Gbps, as needed, allowing the use or reuse of the 10 Gbps layer at the hub. Note that options for dense routing gear can include options of 10 or 25 Gbps, which in cases where the optics allow, the uplink can be in terms of 25 Gbps making the time for the reuse longer. On the other hand, the uplink connectivity to the FAR can be in terms of 100 Gbps or more. This application then skips over the extra 10 Gbps aggregation layer in the hub, going directly from the cores router to the FAR. This is a savings in physical and carbon footprint. Effectively, the aggregation layer in the hub is moved directly to the node.

As we saw previously, the OTN layer is a whole separate logical function that is done independently, and in addition to the routing layer that will also be in use. There are products that aim to combine these two functions so that from the outside it looks like one “box” is doing both. This approach will benefit connectivity, but there is no way of getting around the fact that muxponding and routing are two distinct functions that will evolve and be implemented separately. Combining them in one box negates the benefit of treating these networking functions in their own time and availability of scale.

6.4. Uplink Bandwidth

The bandwidth aligned with direct connect and muxponding solutions is equivalent to the physical connectivity of the remote PHY system, as seen in Figure 6 below. Without needing to know the functionality of the packet cores, one can deduce the bandwidth sizing of the backhaul link.

Note both the direct connect solution and the muxponder solution have to transmit as much (or more) bandwidth as is determined by the number of connections to the DAA endpoints. In the case of direct connect the necessary throughput is 10 Gbps times the number of DAA endpoints. In the case of the muxponder, it is 10 Gbps times the number of endpoints rounded to the nearest multiple of framing speed. For example, suppose the framing speed is 100 Gbps, then for 12 DAA endpoints at 10 Gbps, the throughput would have to be two connections at 100 Gbps or one connection at 200 Gbps composite. If the number of DAA endpoints is 20, the throughput would be the same 200 Gbps. At 24 endpoints, the throughput would have to be 300 or 400 Gbps, depending if the framing speed is at 100 or 200 Gbps multiples. As we will see in section 8, there can be a large difference, for a long period of time, between the connection speed at the DAA endpoints and how much bandwidth is being used in the uplink. Thus, part of looking for the right solution is the ability to predict the amount of bandwidth that will be used for the backhaul over time.

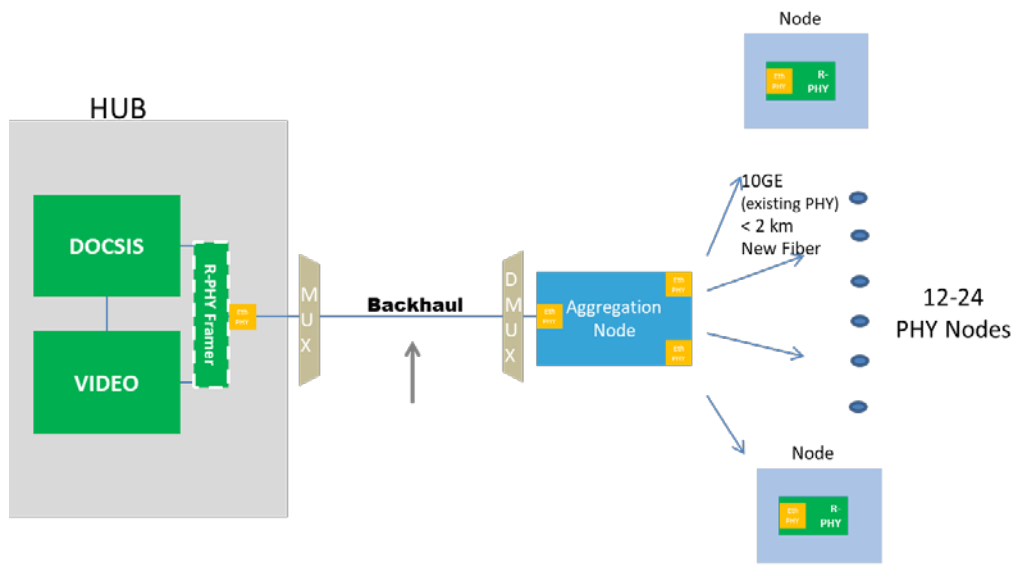


Figure 6 - Physical Connectivity of Remote PHY System

6.5. Multicast and Unicast Bandwidth

An additional consideration is how each solution handles multicast and unicast bandwidth to a group of RPD's. One notable feature of the packet cores used by MSOs is their ability to address their subscribers with sets of common bandwidth. Figure 7 represents the logical connectivity of a DOCSIS packet core, which has at its disposal several unique service groups, and the relationship those service groups have with RPD's. It is not necessarily a one to one correspondence, and when it is 1xN each logical DOCSIS service group can service multiple RPD's. In practice we have seen that the number of RPD's per SG can be as high as eight, though it is more common that we see four RPD's or less sharing unicast DOCSIS bandwidth. A similar relationship exists when considering how unicast video content is shared across multiple RPD's. Although the operator could maintain a one to one relationship between video and DOCSIS service groups, we typically see a 1xN relationship whereas there are typically twice as many RPD's per video service group. This practice is generally in an effort to share unicast video content across a larger base of homes, thus minimizing the changes to the existing back office infrastructure. This is not the case for multicast video content. In theory, one set of multicast video channels could serve an entire headend or hub service area. From a more practical perspective, the number of ad insertion zones, thus unique copies of multicast content dictates how multicast video content is shared across multiple service groups. Since a single FAR would rarely, if ever, span multiple ad zones, it is fair to assume a single set of multicast video channels would typically serve all the RPD's connected to a single FAR.

In the case of the muxponder, there is no method to differentiate replicated versus unique bandwidth. The field aggregation router on the other hand by its very nature has the ability to differentiate packet relationship between DAA endpoints and the packet cores which can be multicast or unicast. This has the effect of significant savings in the overall backhaul bandwidth capacity needed and opens considerable options for managing the use of optics over time, as we will see in section 8.3.

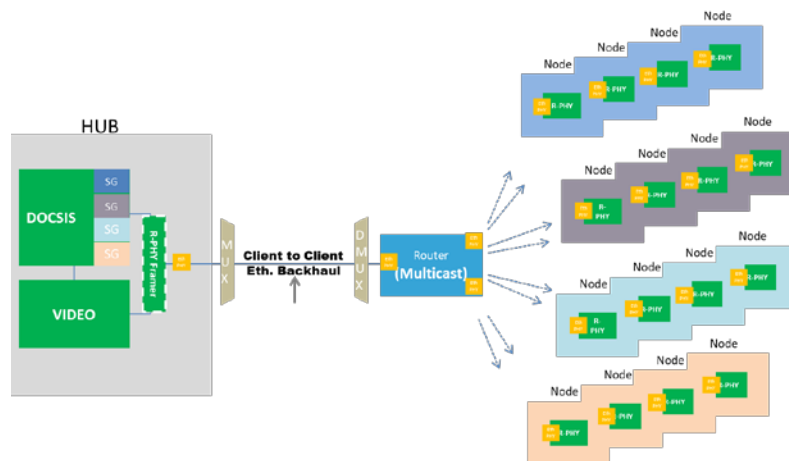


Figure 7 - DOCSIS Service Groups and Multicast Bandwidth

6.6. Converged Access

It is also worth considering that these solutions, in the context of an access network, could converge multiple services beyond what is typically used in Cable, for example mobile or PON. All the solutions here can address this need but there is a difference depending on whether the services are to be aggregated and transmitted (like in the FAR) or transmitted on the same medium but kept separately (like in direct connect or the Muxponder.) In this consideration there is no one answer, but there is a matter of preference for network engineers. This certainly means that there is space for a muxponder solutions that should be investigated, particularly for services that do not replicate bandwidth and run at line rate by contract.

7. Coherent Optics

For reasons of enabling an aggregation element in the field, the MSO community has put considerable effort into making the transmission of 100 Gbps and beyond accessible for the cable access plant. Just recently, CableLabs released a specification for a 100 Gbps ZR solution that is of coherent technology and capable of DWDM channel specificity. This technology also has the ability to cover distances up to 80km and facilitates the use of a FAR, (CableLabs, 2018). There is now an effort to extend a definition to 200 Gbps that also facilitates the natural use of muxponders.

The connectivity of high throughput optics to the FAR falls under Ethernet client to client connections and has been recognized within a greater market opportunity for similar Ethernet signals. This need is also being addressed at the IEEE 802.3 “Beyond 10km” group (802.3, 2018). This effort is tracked very closely in aim to leverage the eventual evolution of ZR 10 Gbps links towards their next transition at 100 Gbps. This will also have the effect of drastically reducing the cost of these optics, ideally in the time frame that will be needed by MSOs.

On the transport side, with distances that can span hundreds of kilometers, there are also efforts that overlap the work being done at IEEE and CableLabs, because the speeds that overlap at 100 and 200 Gbps. For more information see the work being done at OIF, (Forum, 2018), and at OpenRoadm (OpenRoadm, 2018).

8. Traffic Engineering and Backhaul Capacity

8.1. Estimating Uplink Bandwidth Capacity

Understanding and quantifying the difference between provisioned bandwidth and utilized bandwidth is a useful tool when estimating backhaul capacity. For the purposes of illustration in this section, we quantify downstream “provisioned bandwidth” as the DOCSIS bandwidth provisioned per service group. In comparison, we quantify downstream “utilized bandwidth” as the peak bandwidth per service group measured at the WAN port of the CCAP chassis.

In Figure 8 , we see a collection of data that represents several hundred CCAP chassis across a major North American market. The graph illustrates that 78% of the chassis “utilize” between 5 and 15 Gbps of bandwidth as measured at the WAN interface of the CCAP chassis. In an effort to err on the high side, we’ve elected to use 20 Gbps per CCAP chassis for the comparison to “provisioned bandwidth” per CCAP chassis. Using 20 Gbps per chassis as our representative estimate therefore ensures that we’ve captured data from 99% of the chassis within the sample network.




Figure 8 - Utilized Bandwidth per CCAP Chassis

In order to compare the utilized bandwidth to provisioned bandwidth per chassis, one simply needs to take the provisioned bandwidth per service group multiplied by the number of service groups per chassis. Using 32 DOCSIS 3.0 channels plus a 192 MHz of DOCSIS 3.1 OFDM, we could estimate the provisioned bandwidth at approximately 3 Gbps per service group. Using 100 service groups per chassis for our comparison, we can also estimate about a 15:1 ratio between provisioned and utilized bandwidth as represented in Figure 9 below.




Figure 9 - Provisioned versus Utilized Bandwidth per CCAP Chassis

8.2. Engineering Backhaul Capacity

In the same manner as done in the section 8.1, we use a bandwidth consumption approach to project the backhaul capacity needed to a FAR over time. There are three type of data points to consider when projecting the backhaul capacity. One is the projection of actual subscriber usage over time, another is the aggregate usage of the subscribers being addressed by a packet core service group, and finally the composite usage for the subscribers being addressed by the FAR.

Based on empirical field data, 200 Mbps per Service Group is referenced as the “utilized bandwidth” in Figure 9 above. However, within the backhaul capacity modeling we use 400 Mbps per Service Group as an initial value, to err on the high side. The utilized bandwidth per Service Group is expressed on a per subscriber basis in Figure 10 below and is the initial value used to calculate backhaul capacity. We have also found that a compound annual growth rate (CAGR) of 40% is a representative value for consumption growth, for most MSOs in all parts of the world. The growth from 4 Mbps per subscriber in 2018 at 40% CAGR is therefore shown in Figure 10.




Figure 10 - Subscriber Usage Over Time, 40% CAGR

The usage per subscriber is meaningful in context of its grouping within the servicing packet core. In this example, we use DOCSIS as the packet core driving bandwidth usage to its service groups. In cable engineering circles we refer to these groupings in relation to the topology they are serving, the capability of subscribers addressed by a DAA endpoint, households passed (HPP), and the related take rate, which is the percentage of subscribers actually using the service. We note that in fiber deep applications we see topology arrangements such that each remote PHY is set up for 50HHP. In addition, we note that the typical take rate of service is about 50% so only 25 of the 50 possible customers is using the service. Nevertheless, we use 100% penetration in our backhaul estimates in order to err on the high side.

The subscriber size for the logical service groups of DOCSIS varies. This is driven by many factors, but overall it can range from 25 subs, creating a one to one correspondence between service group and remote PHY device, to several hundred subscribers having a one to many correspondence between DOCSIS service group and RPD's. In practice, we see up to eight RPD's per DOCSIS service group, with a common number being four RPD's per DOCSIS service group. The linear addition of bandwidth from its subtending subscribers gives the usage for the DOCSIS service group. Figure 11 shows the aggregate DOCSIS service group capacity for various subscriber densities, in terms of RPD's and homes passed, noting that 100% penetration HHP is in fact the number of subscribers.

One interesting note here is that the physical interface of DAA endpoints, including RPD's of current generation is 10 Gbps, which from Figure 11 shows the time that capacity runs out according to how many subscribers are being serviced. In the case of 400 subscribers that capacity runs out early in the 2023 time-frame, while if you have only 50 subscribers that capacity runs out much later, nearing 2030. This relationship between service group granularity and physical line side capacity of RPD's is important to understand and can help drive decisions for how the DOCSIS groupings and physical data rates of DAA endpoints will evolve. For example, it might be advantageous to have plans in the next decade to reduce number of subscribers towards one to one correspondence between DOCSIS service groups and RPD's. It is also possible that a 25 Gbps RPD will be feasible before the end of the next decade and thus not necessarily force a one to one correspondence between DOCSIS service groups and RPD's. These type of calculations facilitates those decisions.

Figure 11 also allows us to understand the capacity of peak bandwidth per customer. While the subscriber usage is a factual usage value assigned for traffic engineering, there is also the provisioned bandwidth that a Service Provider allocates on a per Service Group basis. The provisioned bandwidth that is shared across

the DOCSIS service group, along with the statistical nature of subscriber usage, is what allow customers to peak beyond their individual usage allowance. Note for example the 400 HHP Service Group, where on day one the provisioned bandwidth for the whole service group is 3 Gbps. These 3 Gbps at any point in time, as given by the statistical nature of usage can be available anywhere within that group of subscribers. Naturally, this is the job of DOCSIS, to facilitate this statistical nature of usage. It would be a mistake to put this job on the network itself. In fact, if the network capacity were calculated in this form, with the given CAGR the throughput capacity would grow to unrealistic values very quickly, making a necessary change in physical interfaces beyond 10Gbps very soon for all cases.

Note: for completeness Figure 11 also includes the framing overhead for Ethernet, IP, remote PHY pseudowire and DOCSIS concatenation. The overhead for a PSP frame with 5 DOCSIS segments and MTU size target at 2000 bytes is calculated to be 3.7%.




Figure 11 - Aggregate Derived DOCSIS Service Group Capacity Over Time

The FAR serves a collection of endpoints which themselves service a collection of subscribers within a number of DOCSIS service groups. In essence, the data in Figure 11 is linearly added according to the number of service groups and thus subscribers passed through the FAR. Figure 12 then shows the backhaul capacity needed on the uplink of the FAR over time. In its lifetime the FAR might service a number of RPD's, Figure 12 also shows curves for 12 and 24 RPD's in service. For completeness, Figure 12 includes a 1.25 Gbps addition to all users for a broadcast video tier, from the video packet core.

Of greatest interest are the actual values of the capacity needed over time as they drive both the connectivity design of the FAR and the optical solutions. With that in mind if we estimate the lifetime of the FAR between nine and eleven years then we see that capacity solutions needed will be within 100 Gbps, making that the obvious long-term bandwidth target for first product implementations. We also expect that towards the end of the decade there would be use for solutions that expand beyond 100 Gbps. The other item to note is that for the early years of the FAR the uplink capacity can be fully addressed without 100 Gbps solutions. We discuss this management of solutions in the next section.




Figure 12 - Backhaul Capacity Needed for Field Aggregation Router, 12 and 24 RPD's Subtended

Note that Figure 12 includes in it the premise of high estimates of very generous 4.0 Mbps usage per subscriber and 100% penetration per DOCSIS service group. The data in Figure 12 therefore gives targets that are shortened in time and larger in capacity than what is more likely to happen, and thus it is an error on the high side as previously indicated. Figure 13 however, is plotted with a starting point of 3.0 Mbps of usage per subscriber and a penetration rate of 75%. Figure 13 is a more likely scenario of what backhaul capacity may be expected. Note that the need for throughput beyond 100 Gbps is further delayed, making the case that maximal capacity at 100 Gbps for the FAR is more than sufficient.

Also, note that backhaul capacity calculations like the ones presented below allow us to calculate the forwarding capacity of the routing fabric. For example, if the maximal capacity needed on the uplink is 100 Gbps, then the downlink capacity needed will be equal, so a 200 Gbps forwarding capacity for a routing chip is sufficient for non-blocking operation.




Figure 13 - Most Likely Scenario for Needed Backhaul Capacity, 12 and 24 RPD's Subtended

8.3. Connecting the Uplink

Note that the data we have shown gives an opportunity for a transitional approach to the uplink. A transitional approach might be beneficial in leveraging better cost or availability for the optics, or architectural flexibility as the network moves from classic analog to a fully formed digital plant. To that end we show the Figures below.

Figure 14 shows a possible initial uplink configuration. As we see from the modeled data, the first few years of the FAR the uplink can suffice with a few, in this case two, 10 Gbps connections. Leveraging for instance some of the open ports on the router if they are not already used for downlinks.

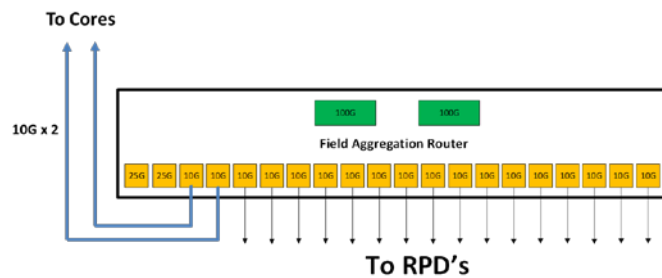


Figure 14 - Initial Uplink Configuration

Figure 15 , shows another possible configuration, where one 25 Gbps link can address the necessary uplink capacity for several years. It is important to note that the optical solutions for 25 Gbps are limited by distance, thereby this solution might not be available to all deployments, but there is a broad footprint within the MSO space where links from hubs to RPD's, and thereby the FAR are within 20 km range.

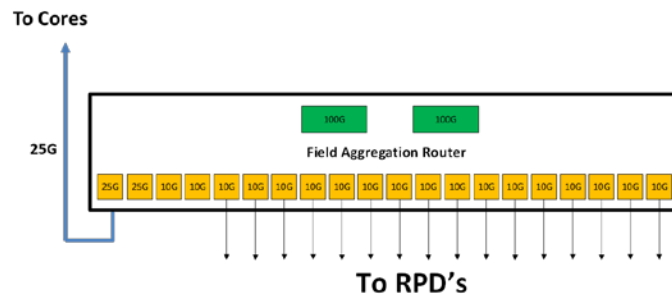


Figure 15 - Secondary Uplink Configuration

Figure 16 , shows the final transition leveraging the large bandwidth available at 100 Gbps. Note that in the router representation there are two 100 Gbps connections which allows for redundancy as needed.

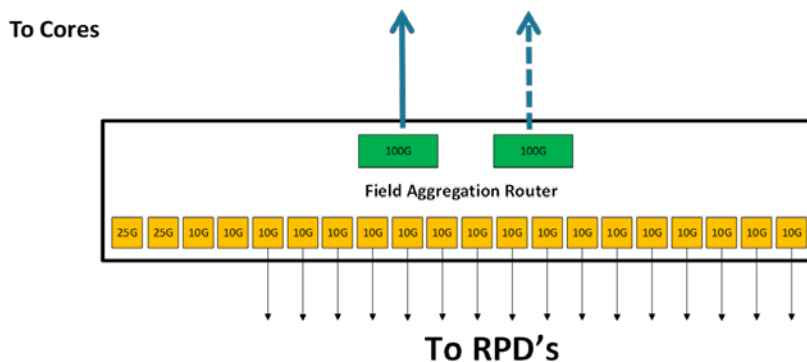


Figure 16 - Third Uplink Configuration

Note also that the Figures above give a method for a more granular approach to building the backhaul. For example, if in year 8-10 of the product an MSO has modeled usage to be below 120 Gbps, but generally over 100 Gbps. That connectivity then is straight forward by using a 100 Gbps connection and a handful of 10 or 25 Gbps connections. This could significantly minimize the cost of the optical connections while not taking away from the functionality of the FAR.

With regards to cost specifically, one of the main challenges to remote aggregation devices, like the FAR or the muxponder, is the cost of the uplink optics, particularly at the 100+ Gbps rates. The challenges of having optics that can generally work beyond 40 km and of wavelength specificity make the technology choices for high rate optics very limited, basically only to the use of coherent optical solutions. The drawback of coherent optical solutions is cost however. One way to limit cost then is to purchase these units as volumes and scales make them available. Another way to limit the cost of the coherent optics bought is to have them align with what will be realistically used, in this case not much more than 100 Gbps. In other words, if someone upsells 400 Gbps capability, there is very little value in that proposition as it might never be used in the timeframe of that product.

Conclusion

Distributed Cable Access Architectures have created a new form of backhaul market that did not exist just a few years. While the DAA backhaul has similarities to what is done for other non-cable access services (like Mobile backhaul, or PON backhaul) the nature of signal distribution for DOCSIS, video and other supporting packet cores allow for unique solutions. As the DAA backhaul is built, it would be best to do so in a cost-effective manner with a general toolset that already exists within the networking world.

This white paper has covered cost and technology comparisons for architectural options for the DAA backhaul in context of bandwidth growth over time. We have also detailed the optical and networking implementations needed to address DAA backhaul. Some of the specific topics that were covered included switching and routing in hubs and HFC nodes, distinctions in DOCSIS and video network transmission, and the applicability of subscriber usage when engineering the backhaul capacity. We also explored the effect of unicast and multicast content on provisioned bandwidth, applicability of OTN framing, implementation of direct detect and coherent optics, and relation to ITU and IEEE standard bodies.

After reading this white paper, the cable operator should have obtained the practical perspective necessary to compare the various DAA architecture options with the view towards deciding for deployment and long-term evolution.

Abbreviations

ASIC	Application Specific Integrated Circuit
BNG	Broadband Network Gateway
CAGR	Compounded Annual Growth Rate
CCAP	Converged Cable Access Platform
CIN	Converged Interconnect Network
CPU	Central Processing Unit
DAA	Distributed Access Architectures
DOCSIS	Data Over Cable Service Interface Specification
DWDM	Dense Wavelength Division Multiplexing
FAR	Field Aggregated Router
FPGA	Field Programmable Gate Array
Gbps	Gigabit per second
HHP	Households Passed
IP	Internet Protocol
L2TPv3	Layer two Tunneling Protocol version 3
MAC	Media Access Control
Mbps	Megabit per second
MSO	Multiple System Operator
MTU	Maximum Transmission Unit
OFDM	Orthogonal Frequency-Division Multiplexing
OIF	Optical Internet Forum
OTN	Optical Transport Network
PON	Passive Optical Network
RMAC-PHY	Remote Mac and PHY
RPHY	Remote - PHY
TDM	Time Domain Multiplexing
WDM	Wavelength Division Multiplexing
ZR	Long Range Optic, 80km.

References

- 802.3, I. (2018, May). *Beyond 10km Adopted Objectives, Study Group*. Retrieved from 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*.

CBRS Use-Cases With focus on Localized Indoor Mobile Access (LIMA), Mobility and Service Continuity

A Technical Paper prepared for SCTE•ISBE by

Rajat Ghai
VP Wireless & Open Networking
Technicolor
Sugarloaf Pkwy, Lawrenceville, GA
+1-508-360-0621
rajat.ghai@technicolor.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	4
1 Introduction.....	5
1.1 SAS, CBSD, ESC.....	6
1.2 Priority Access License (PAL).....	11
1.3 CBRS (3.5 GHz) relative to Unlicensed 5GHz.....	13
1.4 Key Industry Groups Affecting CBRS Deployment.....	13
2 CBRS Use Cases for MSOs	16
2.1 Inside-Out: Localized Indoor Mobile Access (LIMA).....	17
2.1.1 Inside-out Mobile Access economics.....	18
2.1.2 Licensed Assisted Access (LAA)	20
2.2 Outdoor Mobile Access.....	20
2.3 Private LTE Networks	22
2.4 Neutral Host Networks	23
2.5 Industrial IOT Networks.....	24
2.6 Fixed Wireless Access (FWA).....	25
2.6.1 Connect America Fund	25
2.6.1.1 Connect America Fund II	25
2.6.1.2 Connect America Fund II Phase 2	27
2.6.2 Broadband Technology Considerations.....	27
3 CBRS/LTE FWA Network Design & Deployment	29
3.1 CBRS Radio Dimensioning	29
3.2 Mounting CBRS to existing Macro Cell structures	30
3.3 Low Power CBSD-B and ODU.....	30
3.4 Quality of the Radio Link	30
3.4.1 Link Budget	30
3.4.1.1 Losses.....	30
3.4.1.2 Free-Space Path Loss (FSPL).....	31
3.4.1.3 Link Margin.....	32
3.4.1.4 Signal-to- Noise Ratio (SINR).....	32
Conclusion.....	33
Abbreviations	35
Bibliography & References.....	35

List of Figures

Title	Page Number
Figure 1 CBRS Spectrum Tiers.....	5
Figure 2: CBRS Band spectrum sharing system	7
Figure 3 ESC sensor network operation	9
Figure 4 Original exclusion zone scheme	12
Figure 5 New DPA based scheme	12
Figure 6 GWPZ map	13
Figure 7 Main CBRS standards bodies and their roles.....	14
Figure 8 CBRS standards bodies and their interactions with related industry bodies	15
Figure 9 CBRS participant network roles	17
Figure 10 CBRS Inside-Out Mobile Access	18
Figure 11 CBRS network economics	19
Figure 12 CBRS strategic outdoor mobile access	21
Figure 13 CBRS MVNO economics	22
Figure 14 CBRS private LTE networks	23
Figure 15 CBRS Neutral Host Networks.....	24
Figure 16 CBRS Industrial IOT networks	24
Figure 17 CAFII Award amount	26
Figure 18 CAFII target number of households.....	26
Figure 19 CBRS economics for green field broadband	28

List of Tables

Title	Page Number
Table 2	31
Table 3	32
Table 4	33

Introduction

Citizens Broadband Radio Service (CBRS), is a 150 MHz-wide shared spectrum in 3.55 GHz to 3.7 GHz band. FCC spectrum sharing policy creates an innovative way for a lightly licensed tiered access that creates dynamic sharing of spectrum in real time. As a one of its kind spectrum sharing concept, CBRS aims to combine the best of traditional licensed spectrum (LTE) and unlicensed spectrum (Wi-Fi) by combining the best of both technologies.

Specifically, for MSOs, the CBRS band offers them a path to deploying their own LTE network without making significant investments for licensed spectrum acquisition.

MSOs can thus strategically utilize CBRS for diverse use cases like:

- Local Indoor LTE mobile access (LIMA) to augment Wi-Fi coverage to control quality of user experience and offload MVNO costs.
- Leverage the HFC plant and deep fiber nodes to deploy CBRS small cells for outdoor mobile access and as mobile backhaul.
- CBRS based Fixed Wireless Access (FWA) technology to provide broadband access in areas that don't have cable access.
- Help enterprises or venue owners deploy CBRS based *private LTE networks* to beef up in-building wireless coverage and capacity.
- Create new business models like *Neutral Host Networks* and *Industrial IOT* using this band.

MSOs are positioned very favorably to leverage the economics of FCC's CBRS initiative to deploy mobile infrastructure very cost effectively for various use-cases listed above.

In this paper, we describe those use cases, with a specific focus on LIMA (Local Indoor Mobile Access), which we have identified as - by a large margin - the largest opportunity for MSOs to capitalize on CBRS based mobile coverage. Indoors is where there are the highest concentrations of subscribers, as well as majority of time spent on mobile devices by users. The paper provides details on how MSOs can innovate with a hybrid indoor Mobile and Wi-Fi service that seamlessly integrates with the macro cellular network as well as extend service continuity to the MVNO network.

Content

1 Introduction

In 2015, the U.S. Federal Communications Commission (FCC) established the Citizens Broadband Radio Service (CBRS) for, a one of its kind and the first ever, shared wireless broadband use of the 3550-3700 MHz band (also referred to as 3.5 GHz Band). FCC also released the first public notice on protection of pre-existing 3650 to 3700 MHz licenses which were utilized by Department of Defense (DoD), and other incumbents. In 2016 the FCC issued a second report on rule making, position on methodology adopted on protection of pre-existing 3650 to 3700 MHz Band Licensees, and conditional approval of Spectrum Access Servers (SAS) administrators. In 2017 finalization of spectrum rule, which are contained in Part 96 of Title 47 of the Code of Federal Regulation (CFR), referred to as Part 96 in this document. Rules making is still in process for private licensing of CBRS band but should be finalized in 2H-2019 with a (Priority Access License) PAL auction.

CBRS (also commonly known as the ‘innovation band’) was envisioned to support a 3-tier shared spectrum model that required a detailed architecture to be standardized. The FCC has created a three-tiered framework to facilitate shared federal and non-federal use of this band using automated frequency coordinators, known as Spectrum Access Systems (SASs), to coordinate operations between and among users in different access tiers. The CBRS has three tiers of users: Incumbents, Priority Access Licensees (PAL), and General Authorized Access (GAA) users.

Citing the intent of this shared spectrum principle as stated by FCC:

“The Citizens Broadband Radio Service takes advantage of advances in technology and spectrum policy to dissolve age-old regulatory divisions between commercial and federal users, exclusive and non-exclusive authorizations, and private and carrier networks.”




Figure 1 CBRS Spectrum Tiers

Given tremendous interest in this spectrum, the FCC decided to leverage the Wireless Innovation Forum ([Winnforum](#)) as a unifying alliance that brought government entities, regulators, service providers, industry associations and equipment providers together. Winnforum’s goal is to standardize a technology

neutral architecture that enabled a shared licensed access and protection for incumbents. As of end of 2017 the WINNFORUM has created 10 CBRS standards and many policies, databases and procedures.

While CBRS Spectrum and licensed shared access architecture are radio technology neutral, there was significant interest from LTE service providers, manufactures and standards bodies to integrated LTE in CBRS. In 2016 the CBRS Alliance was created to drive standardization of LTE use of CBRS in collaboration with Winnforum.

The FCC identified incumbent users for the CBRS band that fall into several categories:

1. The 3550-3650 MHz band is allocated to the Radiolocation Service (RLS) and the Aeronautical Radio Navigation Service (ARNS) (ground-based) on a primary basis for federal use. Both fixed and mobile high-powered DoD radar systems on ground-based, shipborne, and airborne platforms operate in this band. These radar systems are used in conjunction with weapons control systems and for the detection and tracking of air and surface targets. The U.S. Navy uses the band for radars on guided missile cruisers. The U.S. Army uses the band for a fire finder system to detect enemy projectiles. The U.S. Air Force uses the band for airborne radar Station Keeping Equipment throughout the United States and Possessions to assist pilots in formation flying and to support drop-zone training.
2. The 3600-3650 MHz band is also allocated to the Fixed Satellite Service (FSS, space-to-Earth) on a primary basis for non-federal use. Use of this FSS downlink allocation is limited to international inter-continental systems and is subject to case-by-case electromagnetic compatibility analysis. The Commission has licensed primary FSS earth stations to receive frequencies in the 3600- 3650 MHz band in 35 cities.
3. The 3650-3700 MHz band is also allocated for terrestrial non-federal Wireless Broadband Services. Such service is authorized through non-exclusive nationwide licenses and requires the registration of individual fixed and base stations. All stations operating in this band must employ a contention-based protocol. Base and fixed stations are limited to 25 watts EIRP per 25 MHz. Mobile and portable stations may operate only if they can positively receive and decode an enabling signal transmitted by a base station; airborne operations are prohibited.

1.1 SAS, CBSD, ESC

The core principle of CBRS is dynamic spectrum sharing in a tiered access. For that, a real-time spectrum coordination mechanism has been created to facilitate the required spectrum sharing. The spectrum coordination architecture for CBRS is based on a distributed system. At the top of the hierarchy is the FCC database which centralizes spectrum allocation. The next tier is the Spectrum Access System (SAS). The SAS is a third party certified vendor offering SAS services. The next tier is the sensor network referred to as the Environmental Sensing Capability (ESC). The ESC system detects and communicates the presence of a signal from an Incumbent User to an SAS to facilitate shared spectrum access. The next tier is the SAS user network which interaction with the SAS for PAL and GAA usage.

A block diagram of the CBRS Band spectrum sharing system is given in Figure 2: CBRS Band spectrum sharing system. At the heart of the system is the Spectrum Access System (SAS). It is the gatekeeper that takes information from the FCC Database, other SASs, Environmental Sensing Capability (ESC), and the CBRS Broadband Service Devices (CBSD). Then it applies the FCC rules to allocate Frequency and Power resource to each of the CBSDs.

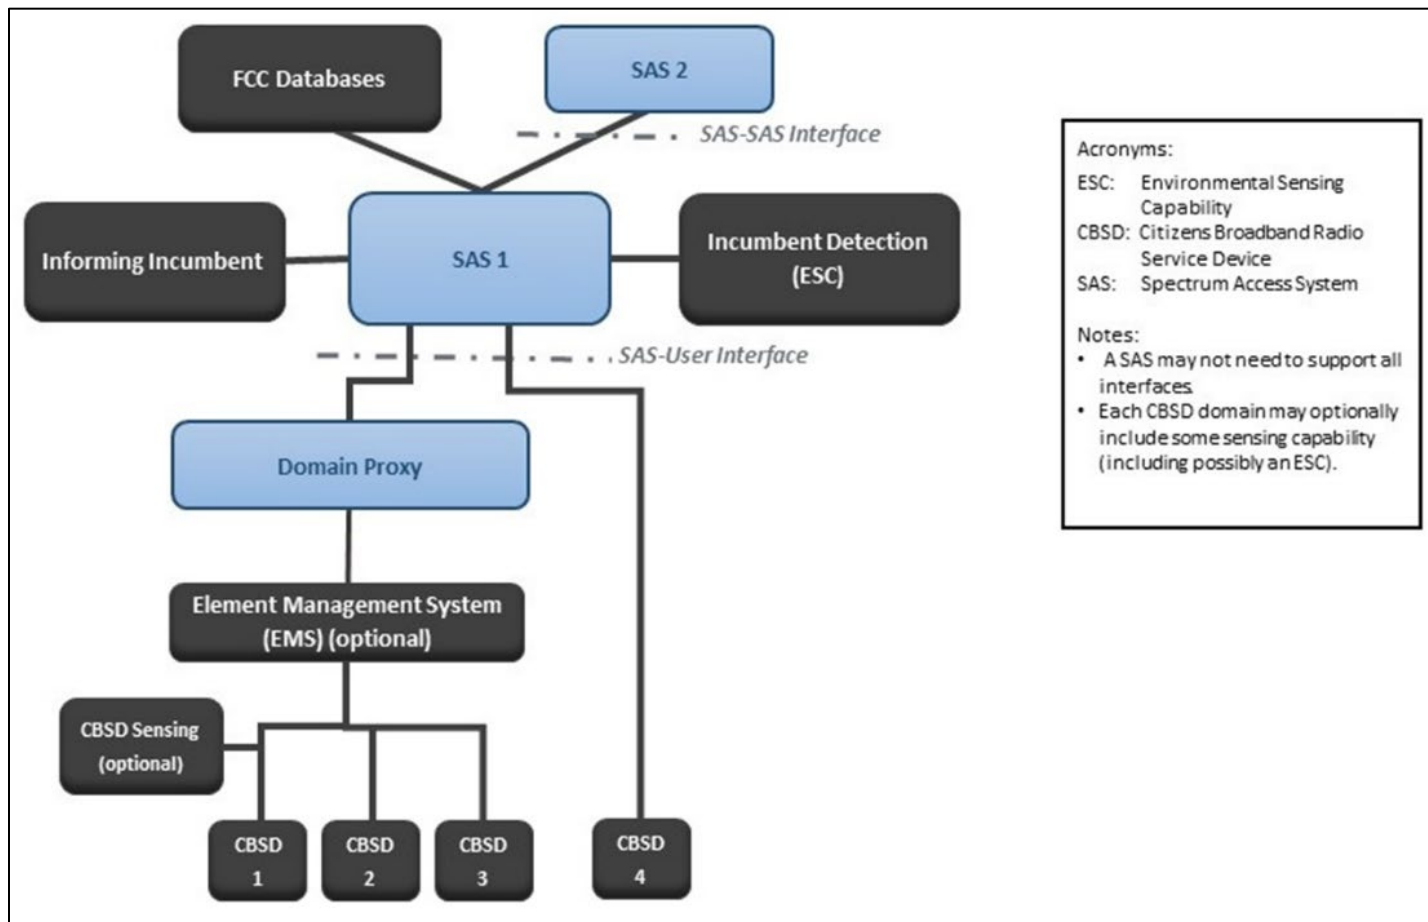


Figure 2: CBRS Band spectrum sharing system

- **FCC Database** is administered by the FCC. It is the repository of tracking information on CBSDs, Incumbents, and PAL licenses. FSS Earth Station incumbents must register with FCC yearly and must include geographic location, antenna gain, horizontal and vertical antenna gain pattern, antenna azimuth relative to true north, and antenna elevation angle. The SAS will communicate with the Database to get relevant information for CBSDs and Incumbents for its geographic area.
- **SAS Operators** are tasked with implementing the HW, Network, and operating the SAS according the FCC rules. SAS Operators must be certified by the FCC.
- **SAS-SAS Interface** will enable communication between SAS's to effectively administer rules across neighboring geographic areas. A group of CBSDs owned by a network operator may extend across SAS's thus requiring coordination between SAS's.
- **ESC** is required before SAS can enable CBSD to operate near DoD incumbents. Without ESC availability an Exclusion Zone of 80 km radius around federal radiolocation sites must be maintained. CBSD are not allowed to operate inside the radius of an exclusion zone. Exclusion Zones are converted to Protection Zones when one or more ESCs are used by the SAS. CBSDs may be authorized within these Protection Zones when ESC reports no incumbent operation. Within 60 seconds after the ESC communicates signal detection from DoD system in given area,

the SAS must either confirm suspension of the CBSD's operation or its relocation to another unoccupied frequency.

- **Informing Incumbent** block enables information gathering and communication between SAS and Incumbent equipment directly. Sensitive DoD incumbents will not be included.
- **Domain Proxy** is included in the SAS system to facilitate communication between groups of CBSDs and the SAS. A SAS operator may make this option available to groups of CBSDs to offload some of the bookkeeping functions and service CBSDs efficiently.
- **Element Management System** is optional block network operator can implement to centralize communication to the SAS network while also offloading and simplifying the individual CBSDs.

The SAS serves as an automated frequency coordinator across the CBRS band. Though the Spectrum sharing concept is similar to UNI-II DFS, the architecture for spectrum sharing in CBRS is very different than that of DFS. CBRS has a frequency coordination model wherein the centralized SAS nodes perform frequency coordination controller function and manage spectrum along with the CBSDs, while in DFS the spectrum coordination is done entirely by the Wi-Fi Access Points that support the UNI-II band. The role of the SAS is to protect the incumbents (higher tier users) from those beneath and optimizes frequency use to allow maximum capacity and coexistence for both GAA and PAL (Priority Access users). It provides dynamic allocation and management of spectrum resources that fall into 3 tiers:

- 1) Tier-1: Protects the incumbents such as DoD / Navy, Fixed Satellite Stations (FSS) which there are about 30 sites in US, and legacy license holders.
- 2) Tier-2: Priority Access operations receive protection from GAA operations. Priority Access Licenses (PAL), are defined as an authorization to use a 10 MHz channel in a single census tract for three years, except in the first auction bidders can request automatic renewal after the first three years for a total of six years. Priority Access License (PAL) licenses will only be assigned in up to 70 megahertz of the lower portion of the 3550-3650 MHz portion of the band. A single PAL holder can only get assigned total of four separate 10 MHz channels (40 MHz aggregation) in each census tract. If a PAL is not in active use, then it reverts to GAA use. It is expected that PAL licenses will be awarded through reverse auctions in 2018. There are still some open items, the most contentious topics are:
 - a. PAL license term: 3 yrs. vs 10yrs.
 - b. PAL block size (75k Census Blocks vs 3142 counties vs 404 PEA).
- 3) Tier-3: General Authorized Access (GAA) which allows opportunistic use of the full 150 MHz CBRS band on a shared basis. GAA has no expectation for interference protection. GAA users must not cause harmful interference to and must accept interference from PAL and Incumbent Users operating in accordance with the rules. GAA deployments will be gated by CBSD certification which is anticipated for Q4'18.

While indoor and outdoor CBRS base station devices (CBSDs) can be assigned to either GAA or PAL, more indoor GAA deployments are expected until ESC certification and PAL auctions get finalized.

The SAS maintains a database and tracks a host of information needed to execute its function.

1. Information of all CBSDs in it controls; Tier status (Tier 1, 2, 3), Geographical location and antenna height within 50 meters (horizontal) and 3 meters (vertical). Such geographic coordinates is reported by the CBSD to the SAS at the time of first activation (e.g. from a power-off to power-on condition). CBSDs also report their location to the SAS within 60 seconds of a


change in location exceeding the accuracy requirement. SAS uses this information to determine frequency availability and maximum power limits for CBSDs.

2. Incumbents in the geographic area under its control. The incumbents can be DoD radar, FSS or WISPs (Part 90).
3. It queries data from the FCC database for list of incumbents in the geographic area under its control.
4. It also audits/logs ESC information related to incumbent activity.

The FCC has set aside a transition period of 5 years, or the expiration of the incumbent FSS license, whichever is longer. During the transition period the incumbent users can operate as they normally do under the old part 90 rules and get full protection from interference of the CBRS network deployments. CBRS deployment in this 50 MHz portion of the band (3650 to 3700) will not start until the rules are finalized. The Incumbents who are eligible to be grandfathered are those that were registered with the FCC Universal Licensing System (ULS) as of April 17, 2015; the date the FCC issued new rules creating the CBRS band. After an incumbent grandfather period has expired they can continue using the band 3650-3700 band or the wider 3550-3700 band for that matter but must follow the new part 96 CBRS rules.

CBRS Devices (CBSD) are certified radio base stations that radiate CBRS 3.5GHz band. Before transmitting the CBSD must contact the SAS. The SAS assigns authorized CBSDs to specific frequencies, which may be reassigned by that SAS. There are two categories of CBSD. Category A CBSD is a lower power base station indoor, and outdoor (6 meters or less in height) and can be self-installed. Category B CBSD is a higher-powered base station for outdoor use only and must be installed by certified installer. Lastly, there are the End User Devices (EUD) that connect to the CBSD and are under the attached CBSD control (e.g. Power Control).

Protecting incumbents (DoD, Navy) is important aspect of sharing this spectrum. As described the sensor network is key aspect of protecting DoD and Navy use of the spectrum. It has been estimated that use of this spectrum is less than 1% of the time but is critical to national security. Shown below in Figure 3 is a depiction of how the sensor network will operate.



Basic geometry of the reciprocal-propagation monitoring approach

Figure 3 ESC sensor network operation

As per the Part 96 rules, a CBSD has 300 seconds to cease transmission and move to another frequency range or change its power level once it receives a command from a Spectrum Access System (SAS) alerting it to a federal system emitting an interfering signal nearby. Per FCC Part 96 *“Within 300 seconds after the ESC communicates that it has detected a signal from a federal system in a given area, or the SAS is otherwise notified of current federal incumbent use of the band, the SAS must either confirm suspension of the CBSD's operation or its relocation to another unoccupied frequency, if available.”*

There are specific radio emission requirements such as transmit power, for each category of device as follows from FCC Part 96.

CBRS Device	Geographic Area	Output Power (dBm/10 Mhz)	Max EIRP(1) dBm/10MHz	Max Conducted PSD (dBm/10 Mhz) (2)	Height Limit
End User Device	All	n/a	23	n/a	n/a
CBSD Cat. A	All	24	30	14	<6 Meters
CBSD Cat. B	Non-Rural	24	40	14	
CBSD Cat. B	Rural	30	47	20	

- (1) Where an FCC rule specifies limits in *radiated* terms such as EIRP or ERP, the limits apply to the maximum emission that would be observed by a linearly polarized measurement antenna. For radiated measurements, the maximum need be performed only over two polarizations for the receive antenna—horizontal and vertical.
 - a. If one of the transmitter outputs is a 90-degree phase-shifted replica of the other and the phase centers of the two antennas are co-located (as would be the case when creating a circularly polarized transmission using linearly polarized antennas), then the each of the two EIRPs or ERPs (total or spectral density) must individually be below the limit
- (2) Where an FCC rule specifies limits on antenna-port conducted power or *conducted* power spectral density (PSD), the rule applies to the total power or PSD delivered to the two antennas (i.e., the sum of the two powers or PSDs).

Another deployment limit is Received Signal Strength Limits within a PAL and GAA users. Part 96 states *“For both Priority Access and GAA users, CBSD transmissions must be managed such that the aggregate received signal strength for all locations within the PAL Protection Area of any co-channel PAL, shall not exceed an average (RMS) power level of –80 dBm in any direction when integrated over a 10 megahertz reference bandwidth, with the measurement antenna placed at a height of 1.5 meters above ground level, unless the affected PAL licensees agree to an alternative limit and communicate that to the SAS.”*

Aggregate Interference Consideration

For non-federal-government protection, it considers aggregated CBSDs within 40-150 km, depending on protected entity and type of CBSD. Federal government protection distances are still being finalized as of May 2018.

Aggregate interference calculation must ensure a result that is at least as conservative as a Monte Carlo method defined in the Requirements, where, essentially, the random variable is the Irregular Terrain Model (ITM) reliability factor (for the ITM model) or the situation-dependent log-normal distribution (for

the eHata model). ITM model is always used in rural markets. For zone-based protection of non-federal incumbents, the aggregate interference is computed across a standard 2" grid (even arc secs in lat/lon).

For protecting (dynamic) federal incumbents, a move list is generated per channel. Generally speaking, CBSDs are rank-order by their impact on interference, and the fewest number/greatest contributors to interference are the ones targeted for re-accommodation to mitigate predicted interference in a channel when federal incumbents are active.

The DoD has divided the offshore area to roughly 200 Km off coast, and roughly 50 areas called Dynamic Protection Areas (DPAs). Each DPA is monitored by one or more ESC sensors. When the federal incumbent activity is detected in the DPA, the entirety of the DPA is protected from aggregate interference to a pre-defined level. CBSDS that may impact interference in the DPA are reconfigured accordingly. DPAs may be used to protect some inland sites.

1.2 Priority Access License (PAL)

It is expected that PAL will currently be auctioned in 2H 2018 or 1H 2019 and will provide the holder 3+3 years initial license term on a per census block (~74,000) then 3 years thereafter. However, the FCC is considering changing some aspects of the PAL tier per (Docket 17-258) such as considering expanding the license term, expanding geographic area, modifying auction rules etc., but none of the other rules in Part 96. Only 70 MHz is available for PAL and this is carved from the lower block (3550-3650 MHz). Current FCC part 96 rules allow each PAL to be authorized for 10 MHz channel, and up to 7 license holders in any given census block.

An important consideration for PAL license holders which will utilize CBSD Cat B base stations is the Exclusion Zones areas (ntia.doc.gov/category/3550-3650-mhz). *"Exclusion Zones shall be maintained for an 80 km radius around the federal radiolocation sites listed in 47 CFR 90.1331 and 47 CFR 2.106, US 109. These Exclusion Zones shall be maintained and enforced until one or more ESCs are approved and used by at least one SAS, in accordance with §96.67. Thereafter, Exclusion Zones shall be converted to Protection Zones."* The Exclusion Zone for DoD, Navy is the blue line in the diagram below which was reduced after NTIA studies in 2015/16 (*NTIA Report 15-517*).



Figure 4 Original exclusion zone scheme

As of May 2018, the FCC has temporarily waived static exclusion zone restrictions based on a new DPA (Dynamic Protection Zone) scheme based on DPA enabled SASSs. DPAs and DPA enabled SAS function is shown below:

- Offshore region is divided into “Dynamic Protection Areas” (DPAs)
- Each DPA is monitored by one or more ESC sensors
- When federal incumbent activity is detected in a DPA, the entirety of the DPA is protected from aggregate interference to a pre-defined level
- CBSDs that may impact interference in the DPA are reconfigured accordingly
- DPAs may be used to protect some inland sites



Figure 5 New DPA based scheme

There are also Wireless Protection Zones to protect incumbent fixed wireless operators, that are provided protection until their license sunsets by 2020 to 2023. This for 50 MHz of spectrum that falls in the 3650 – 3700 MHz band.

- Grandfathered Wireless Protection Zone (GWPZ) highlighted in red are protected until their licenses are sunset by 2020-23 timeframe
- Protected in all or portion of 3650-3700 MHz spectrum
- Protections based on aggregate interference, propagation model, and 2" grid per spec
- All CBSD as far as 40 Km from the GWPZ are considered



Figure 6 GWPZ map

1.3 CBRS (3.5 GHz) relative to Unlicensed 5GHz

LTE deployment in the CBRS band can be augmented with LTE in unlicensed 5GHz spectrum which has made significant progress with the support of LTE-U in release 12, LTE Assisted Access (LAA) release 13 and LTE Wi-Fi link aggregation (LWA) in release 13. LTE-U, LAA and LWA however require licensed anchors. The development of MulteFire which does not require a licensed anchor can be considered by NSPs who don't operate Licensed (exclusive use) spectrum.

Unlicensed 5GHz band in US is available for FWA point-to-point and point-to-multipoint which has some potential application. 5GHz UNI-3 and UNI-4 defined in FCC Part 96 can provide a number of enhancements to FWA deployment such as Mesh networking of Small Cells, and point-to-point backhaul.

1.4 Key Industry Groups Affecting CBRS Deployment

Many Standards bodies as well as industry groups are working together and alongside the FCC for development and deployment of services in the CBRS band. The two main active bodies doing most of the work are

- WinnForum
- CBRS Alliance

The responsibilities of these two bodies are described and compared below:

Comparing the WInnForum and CBRS Alliance

WInnForum



- Official SDO with multiple committees
- Spectrum Sharing Committee (SSC) handles FCC Part 96 rules (CBRS) & working closely with US Government
- **Technology Neutral** (many members support LTE, but other members support WiMAX & proprietary technology use in the band)
- Developing SAS, CBSD & ESC requirements; Security methods; SAS-CBSD & SAS-SAS protocols; Certification tests for: CBSD & SAS
- **Specifying Coexistence across multiple technologies in Release 2 (not yet approved)**

CBRS Alliance



- Focus on LTE technology in the CBRS Band. Builds upon and compliant with WInnForum Standards
- **Developing technical specs to support LTE deployments of Private Networks, Neutral Host Networks, Multi-Service Operator Networks, etc**
- **Focus on LTE coexistence**
- Addressing LTE commercialization, business and marketing issues
- Broad range of members: manufacturers, operators, verticals and more

Coexistence work is ongoing in both the WInnForum and CBRS Alliance

Figure 7 Main CBRS standards bodies and their roles

These two industry bodies interface with many other bodies e.g. Incumbents, Govt., Standards etc. that pertain to various technologies being developed for the CBRS band. As an example, the CBRS alliance works closely with 3GPP. 3GPP organization is a standards body that is responsible for creating specifications related to LTE.

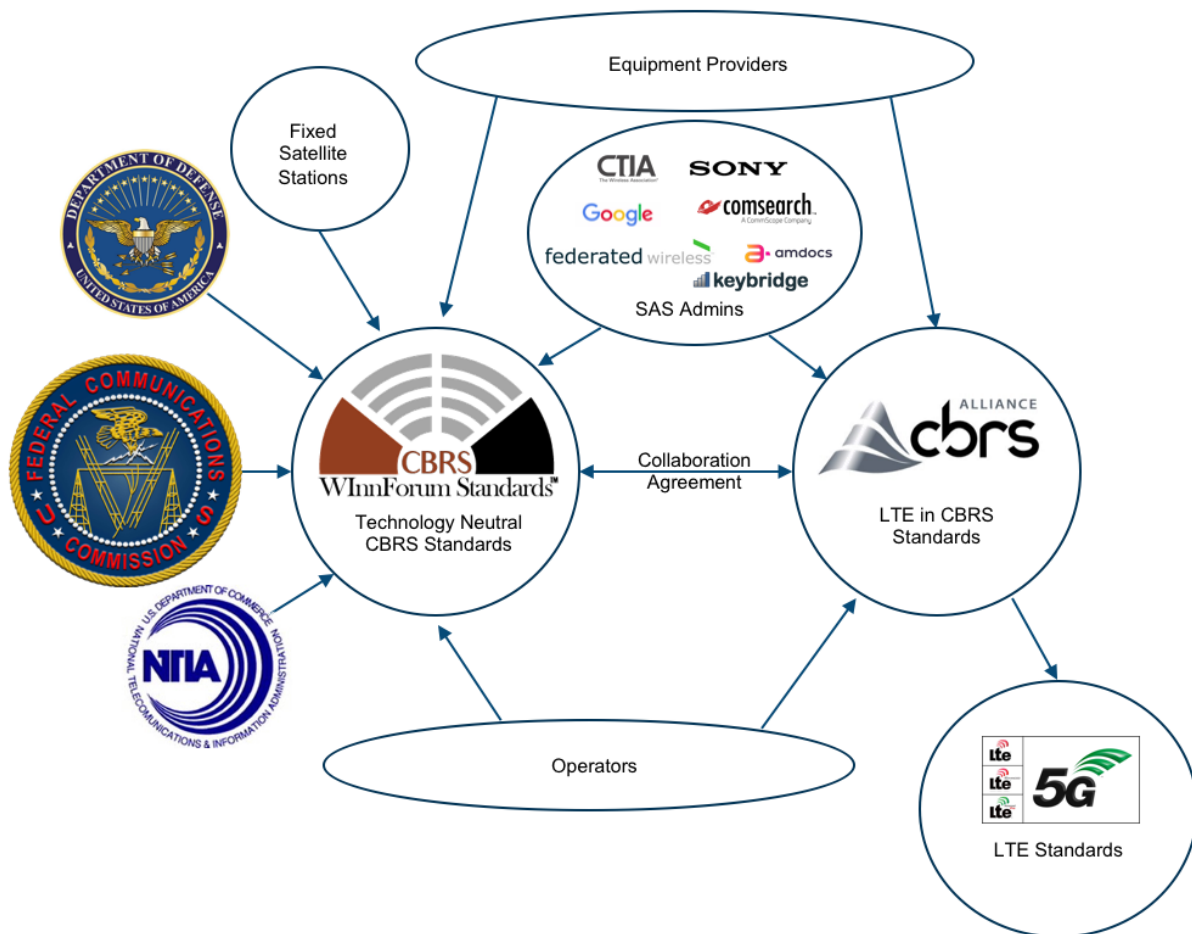


Figure 8 CBRS standards bodies and their interactions with related industry bodies

2 CBRS Use Cases for MSOs

Historically, licensing spectrum for exclusive use (Licensed Spectrum) has been very expensive to acquire, costing billions of dollars and representing a majority of the cost of a mobile wireless network. Such an upfront and sunk-in spectrum cost (again in billions) was considered cost prohibitive as it created a barrier to entry for non-traditional cellular operators.

The FCC's choice of (first ever) spectrum policy innovation, through creation of shared spectrum rules for CBRS, significantly lowers the barriers to entry for non-traditional wireless carriers. The flexible three-tier licensing framework lowers the barrier to spectrum and promotes success-based investment for new entrants. Due to significantly lower cost of PALs compared to exclusive use licensed spectrum costs, the FCC has leveled the playing field by democratizing LTE networks. MSOs specifically stand to gain substantially from this CBRS initiative.

Considering that a typical traditional Tier 1 US mobile operator holds, on average, about 130 MHz of licensed spectrum for exclusive use, in contrast, 150 MHz of favorable mid-band spectrum in the CBRS band is a significant resource for MSOs to provide LTE based mobile capacity to compete with traditional Mobile operators. CBRS thus offers cost-effective LTE solutions for both indoor and outdoor applications; opens up new use cases; and, encourages new revenue generating business innovations for MSOs.

CBRS Service Model

Shared Spectrum democratizes LTE network services where, depending of the business model, the LTE Radio Network may be offered only as a transit access or a vertically integrated mobile services like the ones offered by the cellular operators using licensed spectrum. Such decoupling of access and services makes it possible for the service providers to perform roles as shown in the figure below:

1. CBRS Network Operator (CNO)

CBRS Network Operator deploys a CBRS/LTE Radio Network at a venue or across a geographical footprint with the intention to provide mobile connectivity using LTE. Typically, a CBRS Network Operator does not have a direct business relationship with the end user or device. CBRS Network would have business relationship with the Participating Service Providers that have direct relationship with the end users. For example, if a mall owner deploys a CBRS LTE network in a mall to provide better indoor mobile coverage to mall patrons; the mall owner will allow access to subscribers of 'Participating Service Providers' that have business relationship with the mall owner. Participating Service Providers in this case are the Tier 1 Cellular operators.

2. Mobile Service Provider

Mobile Service Provider provides mobile services to end users. They have a business relationship with the end user and provide them valid SIM cards that let the end users get authenticated and authorized for mobile services. Mobile Service Provider may also assume the role of CBRS Network Operator (CNO) if it also operates a CBRS access network.

3. Subscriber

End user or device that is requesting mobile services.

Roles in CBRS Services

- **CBRS Network Operator**
 - Deploys a CBRS network with an intention to provide connectivity and/or enable services to Subscribers of participating service provider(s). Network operator may need to authenticate the subscriber, as well.
- **Service Provider (SP) Role**
 - Have service agreement, authenticates, authorizes, and provides services to subscribers,
 - MNOs, MSOs, or MVNOs
 - A Participating Service Provider (PSP) is a service provider offering services via the specific Neutral Host.
- **Subscriber Role**
 - Authenticated and authorized by one or more service providers.
 - Could be person or a device;

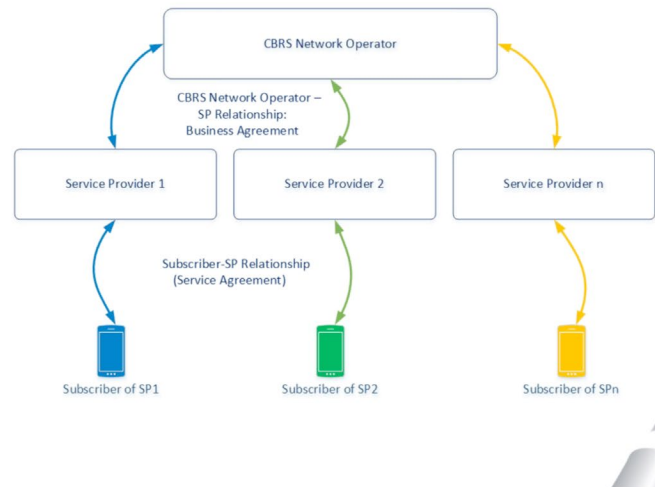


Figure 9 CBRS participant network roles

While there are a number of application MSO CBRS use cases, the major ones are covered here for review in this paper:

- **Indoor Mobile Access**
- **Outdoor Mobile Access**
- **Private LTE Networks**
- **Neutral Host Networks**
- **Industrial IOT networks**
- **Fixed Wireless Access (FWA)**

2.1 Inside-Out: Localized Indoor Mobile Access (LIMA)

Most industry statistics state that an average of **80-90%** ^[2] of mobile sessions happen indoors while the rest occur outdoors. MSOs currently have the lion's share of Fixed Broadband and an elaborate DOCSIS broadband capable network that reaches **85%** ^[1] of US residential and enterprise locations and offers High Speed internet connectivity of upto a theoretical max of 10 Gbps DL and 2 Gbps UL. Leveraging such a deep cable network, MSOs can employ a novel *inside-out* ^[3] strategy where the MSOs can build a massive LTE network by initially focusing on advanced wireless LTE solutions inside the residential and enterprise, and eventually expanding outdoors (hence *inside-out*). As an example, if an MSO has 10 million broadband cable subscribers, they could convert all the broadband subscribers into a 10 million LTE cell towers by incorporating CBRS/LTE base station (eNB) function in the DOCSIS modems. Such an *inside-out* strategy leverages deployment of indoor CBRS small cell radio in the home or business. A

small cell is a low power radio node that connects to a mobile Evolved Packet Core (EPC) and/or IP Multimedia Services (IMS). An indoor CBRS small cell provides indoor mobility initially then moving to outdoor.

Such an indoor mobile network will cover 80-90%^[2] of the mobile sessions on the indoor network, the remaining 10% of the mobile sessions originated by their subscribers outdoor 'on the go' may be offloaded to the MVNO partner network (i.e. a Tier 1 cellular operator partner).

Inside-Out Mobile Access

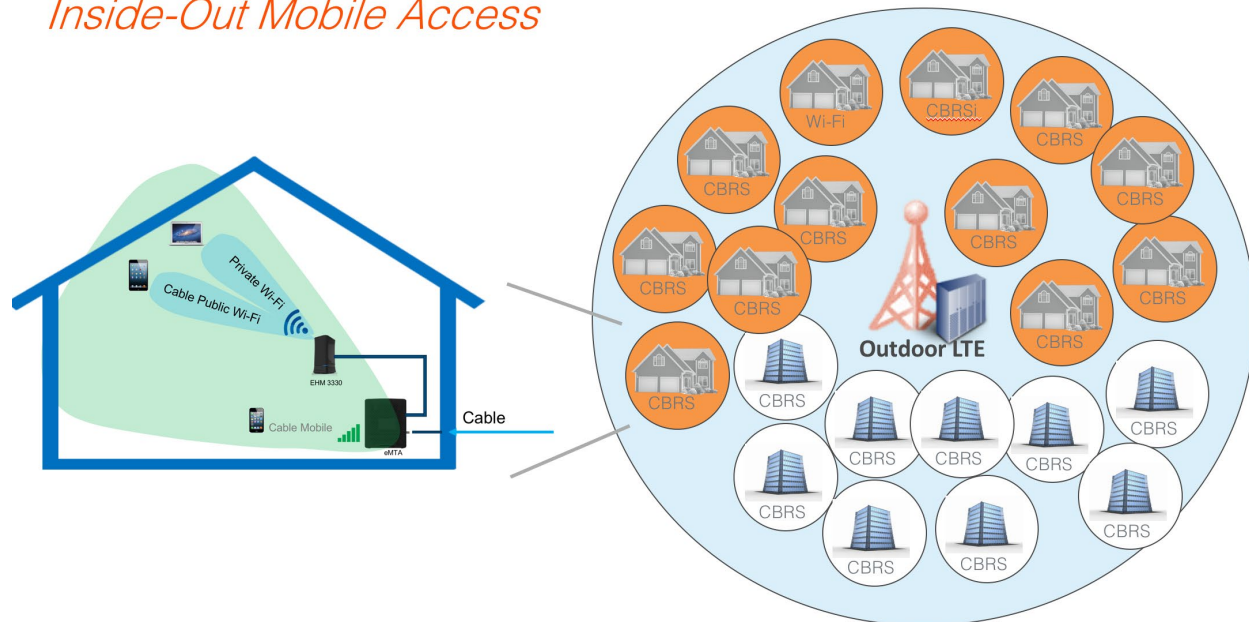


Figure 10 CBRS Inside-Out Mobile Access

Reduced CBRS shared spectrum acquisition cost along with an *Inside-Out* enables MSO to launch a *near* ubiquitous market wide LTE network at a fraction of the cost of a traditional macro cellular network AND at a fraction of the cost of being a pure MVNO. This can lead to an enablement of a very competitive mobile service offerings to compete with the traditional Tier 1 Cellular Operators and create a profitable and sustainable mobile wireless business for the MSOs.

2.1.1 Inside-out Mobile Access economics

Mobile wireless subscription services provide a great new source of revenues. In fact, users spend an average of \$45/month/device, hence a family of four spends around \$180/month on mobile subscription plans. *As an example, an MSO with 20 million broadband subscribers stands to gain \$43.2 Billion of TAM in its served market.* This is approximately 2x of what a household spends on average on residential cable broadband and Linear broadcast TV subscription combined, which is approx. ~\$80/month. MSOs already have a Triple play service offering (Voice/Video/Internet), and with a mobile wireless they can more than double their revenues by becoming a Quad play provider. Historically however, building and operating a wireless with traditional approach (Macro cellular) was a barrier to entry for MSOs. Firstly, there were huge spectrum acquisition costs associated with licensing wireless spectrum for exclusive use. Then, deploying a macro cellular network using 200+ foot towers, as well as hiring a trained workforce of radio network planners who could install and maintain such a complex radio network was a very difficult task as well. All in all, MSOs found it near impossible to enter and compete in wireless service market with a tradition network build approach^[5].

MVNO based mobile service gets MSOs into the mobile business quick, however the MVNO terms were negotiated years ago, with certain fixed costs between \$5 to \$10/GB (close to \$8/GB ^[7]), and it's likely that the MNO host operator wouldn't be eager to give better MVNO deal anytime soon due to competitive reasons. With an average mobile data usage of approx. 11.9 GB / month ^[6], a pure MVNO deal would require the MSO paying close to \$95/month per mobile device to their MNO host. Clearly, such a business model is not sustainable in the long run if the MSO unlimited plan subscriber (majority of the subscribers) only pay \$45/month to the MSO ^[8].

However, with CBRS/shared spectrum, an innovative *inside-out* strategy coupled with opportunistic Cable Wi-Fi offload and a modest outdoor MSO owned outdoor CBRS/LTE network in strategic densely populated areas (hot zones), the MSO can minimize the amount of wireless traffic that flows over the MNO host cellular network. This is the inflection point that creates a very powerful competitive advantage for the MSOs.

CBRS/LTE based Cable Gateways in particular (*inside-out*) provide the greatest advantage as MSOs now have close to 75% Cable broadband penetration in US ^[9]. This provides the opportunity for the MSOs to turn each one of such locations as a turnkey CBRS/LTE cell tower to create a near ubiquitous indoor LTE network to provide the much-needed critical mass of coverage for their mobile wireless network.

The chart below shows costs associated with delivering a wireless service for various network deployments types ^[10].




Figure 11 CBRS network economics

Given that 80% - 90% ^[2] of the wireless sessions originate indoors, it is highly likely that an MSO network planned around *inside-out* strategy could reach a traffic distribution pattern that, for a typical user, might look as follows:

- 50% of the monthly LTE data on MSO owned indoor CBRS/LTE network based on LTE enabled Cable Gateways; i.e. 5.95 GB/month/device

- 20% of the monthly LTE data on MSO owned outdoor CBRS/LTE network based on LTE small cells at their outdoor fiber optical nodes of their fiber deep HFC plant; i.e. 2.38 GB/month/device
- 10% of the monthly LTE data strategically offloaded to the MSO owned Cable Wi-Fi network; i.e. 1.19 GB/month/device.
- Rest, 20% of the monthly LTE traffic would use the MNO host's macro LTE network; i.e. 2.38 GB/month/device

For such a monthly traffic profile, the MSO can provide a profitable wireless service @ a cost of \$22.25 / month / device, which leads to a sustainable business since such a network costs only 1/4th of the, otherwise, cost of \$95/month/device associated with a “pure” MVNO business model.

2.1.2 Licensed Assisted Access (LAA)

Since licensed spectrum is a scarce resource, a number of radio innovations have been adopted by LTE standards to use unlicensed spectrum with LTE radio. The first innovation was LTE Unlicensed defined in 3GPP release 13.

LTE-U and LAA protocols will be utilized in the 3.5 GHz band. LTE-U and LAA are desirable technologies, because they will allow carriers to expand their capacities while still ensuring that they can rely on stable licensed spectrum for high quality service. Current versions of LTE-U and LAA operate with an anchor licensed carrier's channel and use carrier aggregation to integrate licensed and unlicensed spectrum, while utilizing coexistence mechanisms to avoid interference, ensure fair sharing with other unlicensed technologies, and enable flexible spectrum use.

LAA (Licensed Assisted Access) is a standardized version of LTE-U as governed by 3GPP. In certain markets such as the United States a protocol called Listen-Before-Talk (LBT), which was designed to address fair coexistence, is not a necessarily implemented in all solutions since it is not a regulatory requirement.

LWA stands for LTE WLAN Aggregation. LWA configures network to allow use of both Wi-Fi and LTE network simultaneously. Unlike LTE-U and LAA which requires hardware changes to co-exist with WLAN networks, LWA relies on aggregation of the Wi-Fi and LTE traffic in the core network w/o any explicit hardware modifications / requirements on radio nodes.

LAA, LWA and LTE-U can enable the MSO to combine and better utilize the Wi-Fi and CBRS/LTE spectrum and provide their users a much superior Quality of Experience as a managed end to end wireless service.

2.2 Outdoor Mobile Access

MSO have invested in deep fiber (HFC plant) successfully for the last 20 years to support the high speed DOCSIS 3.0 broadband access to their subscribers. MSOs can now leverage the dense HFC plant in general, and the fiber nodes in particular, to strategically install outdoor CBRS/LTE metro cells to further densify MSO mobile access (beyond indoor densification using the *inside-out* strategy).

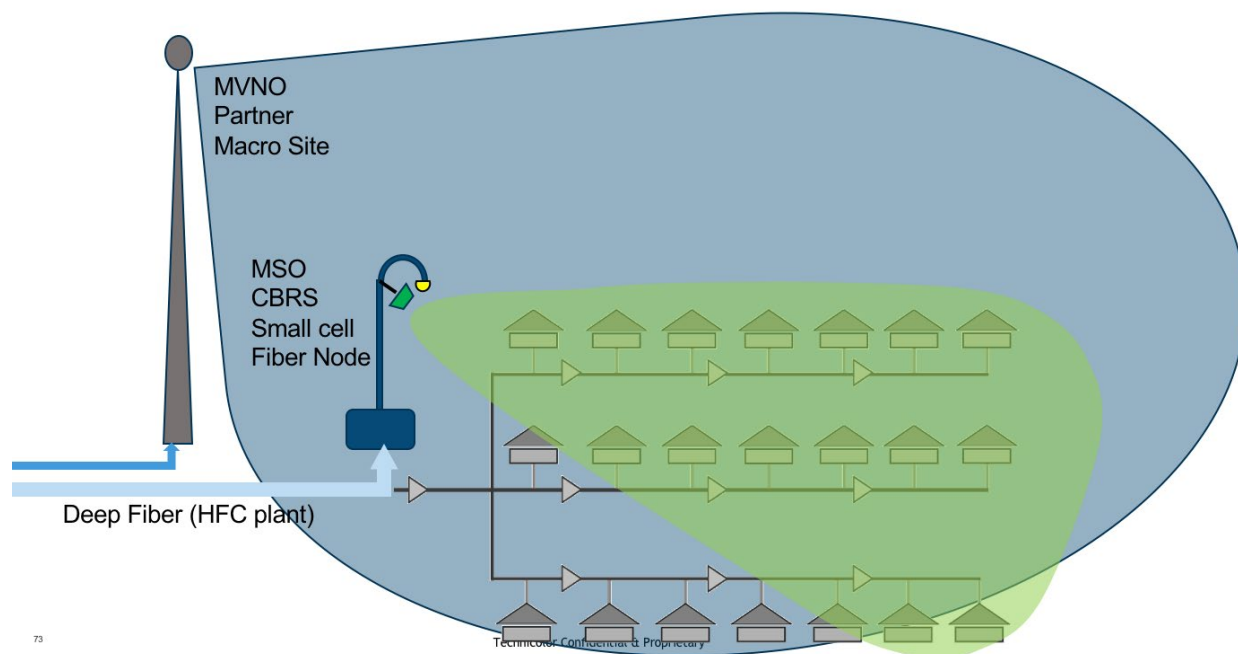


Figure 12 CBRs strategic outdoor mobile access

Leveraging HFC / fiber nodes to provide outdoor CBRs/LTE mobile coverage is a great option for MSOs to build out an LTE network and make the Mobile Virtual Network Operator (MVNO) economics work in their favor. LTE service across both host macro network and owned CBRs small cell network may simplify network integration efforts and will likely result in more predictable user experience than offloading to Wi-Fi. Since US cable operators do not yet own much licensed spectrum, this is a big upgrade from Wi-Fi. MSOs can capture additional subscriber mobile traffic on the 3.5 GHz band with LTE and 2.4/5 GHz bands with Wi-Fi to reduce the amount of charged traffic going over to the host mobile operator network. The profitability of a MVNO business case is heavily dependent on lowering the amount of traffic going over to the host mobile operator network. Since an MVNO pays a mobile operator for traffic going over the host operator's network, higher subscriber usage directly translates to higher network cost. Hence, for the cable operators, this means offloading subscriber traffic over to owned networks as much as possible.

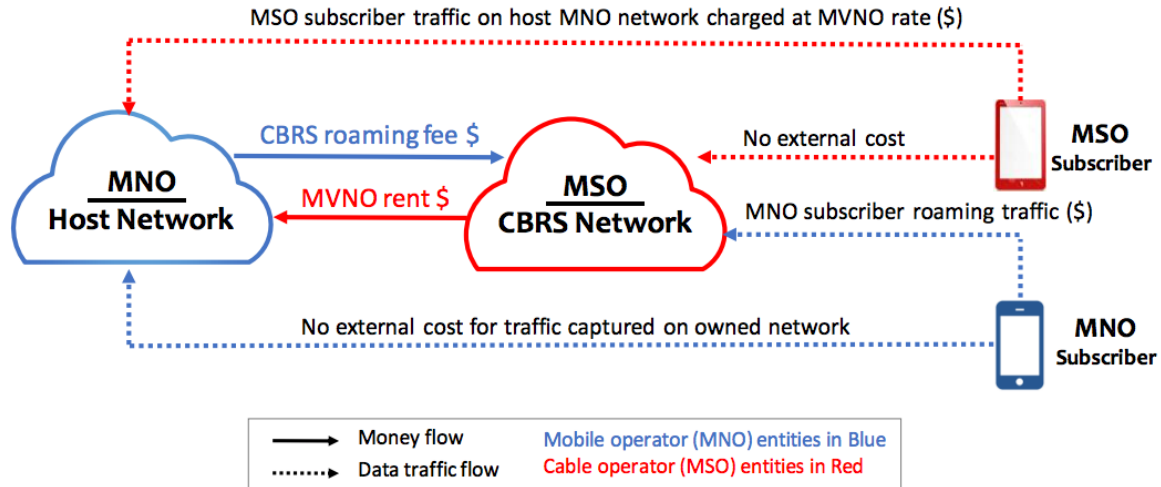


Figure 13 CBRS MVNO economics

Another business model involves a bilateral roaming arrangement with the host mobile operator to allow MNO subscribers to roam onto the MSO’s LTE small cell underlay network in exchange for a lower MVNO terms. Having owned LTE-based network in strategic places where most of subscriber traffic is generated or consumed, affords additional optionality for the cable operators. Besides reducing MVNO expenses through traffic offloading, the cable operators can negotiate for better MVNO terms involving a potential “swap” deal.

2.3 Private LTE Networks

Large enterprises have traditionally deployed Wi-Fi networks to satisfy the growing wireless data demand. However, it has been a poor substitute for critical mobile wireless internet or seamless mobile voice services indoors. The FCC has democratized LTE by making CBRS/LTE a shared spectrum as opposed to exclusive use licensed spectrum that Tier 1 operators use. Like Wi-Fi access points, Enterprises and venues can run seamless LTE services and create a private LTE network, in a similar manner as Wi-Fi, to run enterprise- or venue-specific applications on mobile devices of consumers or workers, enabling tremendous flexibility; it also allows enterprises to tap into broader device and app store ecosystems that already exist. For instance, a large corporation can run secure enterprise CRM and communication tools on workers’ mobile devices through a private LTE network at enterprise campuses. In another example, a heavy industry company can set up a private LTE network at a remote mining site and run industrial IoT applications on LTE devices.

MSOs can create new revenue streams by creating turnkey Private LTE solutions for Enterprises.

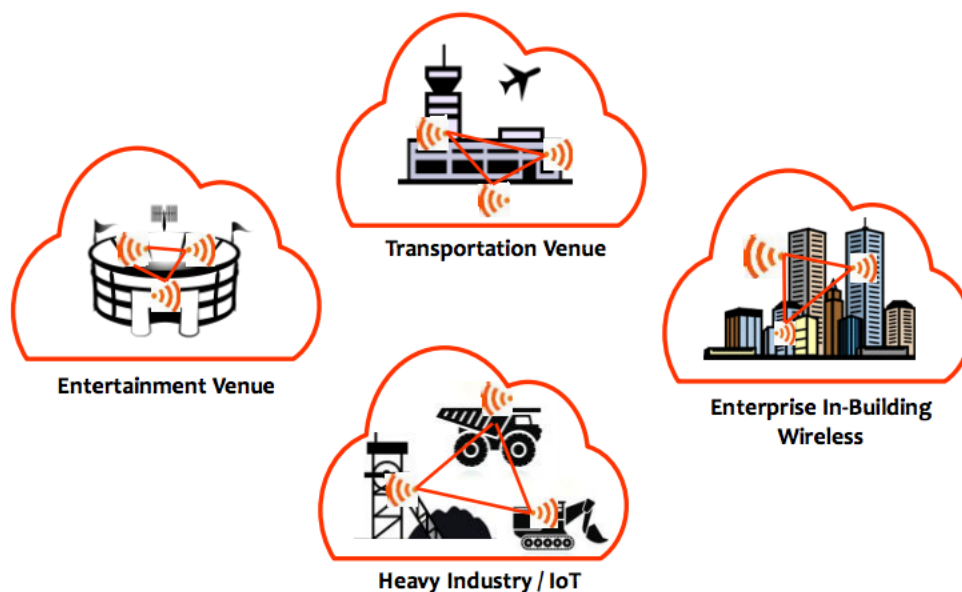


Figure 14 CBRS private LTE networks

2.4 Neutral Host Networks

There is a growing need for neutral host providers to bridge the gap between very large projects with direct mobile operator involvements and large numbers of smaller projects that are too small for mobile operators to consider, but too complex for enterprises to handle on their own. There is an opportunity for MSOs with CBRS/LTE deployments that involve SAS coordination and managing core network integration with mobile operators. Beyond the obvious large public venues such as stadiums and airports, hi-rise buildings, large hospitals, and university campuses are well suited for neutral host providers to address a growing, pent-up demand for in-building wireless coverage and capacity expansion. For enterprises with limited IT/telecom resources, a neutral host provider can take over the technical work and coordination with the operators. About 30 billion square feet of US commercial floor space has poor mobile coverage. With broad support from all four major operators and leading device platform vendors, neutral host providers can create a major new category of mobile coverage, funded by the enterprise or property owner.

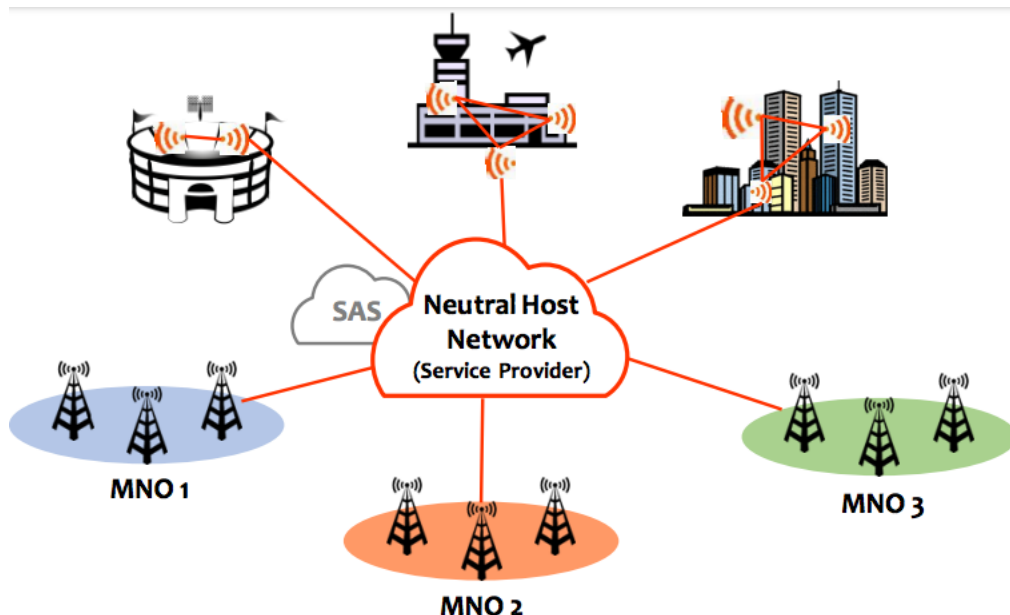


Figure 15 CBRS Neutral Host Networks

2.5 Industrial IOT Networks

As the FCC democratizes LTE in 3.5GHz (CBRS) band, it also paves the way for using private LTE for mission / business critical Industrial IOT applications. Unlike licensed spectrum based IOT that is operated and managed by the cellular operators, CBRS based Industrial IOT networks are owned by the enterprise, managed locally, using dedicated network LTE RAN that can be optimized for the specific Industrial process.




Figure 16 CBRS Industrial IOT networks

2.6 Fixed Wireless Access (FWA)

Within the US, the fixed broadband divide is significant. According to the January 2016 Federal Communication Commission “2016 Broadband Progress Report”, 34 million (10 percent of all Americans) do not have access to broadband service at current FCC minimum standards of 25 Mb/s download and 3Mb/s upload speeds for fixed services; 23 million people (39 percent of rural Americans) lack access to broadband; and 1.6 million people (41 percent of Americans living on tribal lands) do not have access to broadband. Rural communities are any geographic census block with 25 or fewer homes per square mile. These areas have a high cost to deliver broadband service with the FCC defines at 25 Mb/s Downlink and 3 Mb/s Uplink.

Note: Globally the digital divide is a significant issue even in developed countries in EU, Latin America and Asia. According to the WBA July 2017 report (source) approximately 1.75bn citizens in the world’s 8 richest countries remain unconnected. Please see report Global Fixed Wireless (TBD) Technicolor white paper.

The significance of this digital divide in the US has resulted in government incentive programs to bring broadband to rural America. The first program developed was the Connect America Fund (CAF) established in 2014 which provided, and the second and much larger program was the Connect America Fund phase-2 (CAF-II) which allocated \$1.675 Billion per year for 6 years.

With CBRS spectrum availability in rural markets for GAA in Q2-2018, and the eminent auction of PAL in Q3-2018, CBRS leveraging LTE technology becomes a viable option for CAF-II service providers to deliver their broadband coverage commitment.

2.6.1 Connect America Fund

2.6.1.1 Connect America Fund II

In August 2015, price cap carriers either accepted or declined “statewide commitment” to provide voice and broadband in their study areas. Of the \$1.675 billion per year available from 2015 to 2020. Price cap carriers elected to receive \$1.5 billion of the annual fund. The 4 biggest recipients were CenturyLink, AT&T, Frontier (included Verizon original CAF locations), and Windstream.




Figure 17 CAFII Award amount

The award of this funding covered




Figure 18 CAFII target number of households

Over 3.6 Million locations will receive broadband service based on FCC buildout requirements starting at 40% in 2017 and completing 100% of locations by 2020. Multiple service tiers can be offered as long as they offer at least one standalone voice plan and one service plan that meets the performance and latency requirements. The baseline tier is 25 Mb/s, with 10Mb/s being the minimum, higher tiers are and classified as above 100 Mb/s and Gigabit if 1 Gb/s. Latency is either Low < 100 ms, or High <750 ms & Mean Opinion Score for voice of >4.

2.6.1.2 Connect America Fund II Phase 2

As you can see in table figure above all most 1.5 Million locations did not get allocated in the 2015 CAF-II auction. In 2017 the FCC decided to have a new auction for these unallocated locations and create more lenient requirements for bidders. The new action # 903 is schedule for some time in 2018 based on eligible census blocks, where filing for eligibility of short-form application is due on March 30, 2018. The funding for these CAF-II Phase 2 locations will be \$198 Million over 10 years, or approximately \$2 Billion over the course of 10 years.

2.6.2 Broadband Technology Considerations

The technologies available to close this gap are Copper, Fiber, Cable or Wireless. Advancements in VDSL2, G.Fast make copper a low cost alternative but many rural markets do not have good copper plants and investments have limited upside in terms of capacity growth over the longer term. Ideally, Operators look to Fiber to the premise due to the long-term investment benefits, but the initial cost is very high with average distances to each home or business much higher than suburban/urban applications. Cable DOCSIS 3.1 can achieve Gigabit speeds but are high cost with requirements for deep fiber for backhaul and civil engineering costs to bring coax to homes. Alternatively, LTE wireless presents a compelling technology option to address the most challenging and high cost regions of the rural market.

Wireless broadband access offerings in unlicensed (5 GHz), Microwave (or lightly licensed (3650 – 3700 MHz) have been available for some time in the US via Wireless Internet Service Provider (WISP), however these networks remain fragmented today. There are over 300 WISPs serving US customers mostly in high cost rural areas. Most of these fixed wireless solutions utilize closed ecosystem solutions from the likes of Canopy, Ubiquity, Motorola Solutions (others). Product offerings range from point-to-multipoint and point-to-point.

With the introduction of LTE-A and LTE-A pro performance enhancements such as:

- Carrier Aggregation
- MU-MIMO
- Virtualization of LTE core network

And a healthy ecosystem driving the economics of LTE to Wi-Fi and IoT economics in wider band Licensed and Unlicensed spectrum, LTE has become a viable fixed broadband access technology competing with Cable and Fiber.

LTE is now capable of utilizing Unlicensed spectrum in 5Ghz band, in addition to new shared licensed spectrum CBRS 3.5Ghz band in US. Wider spectrum bands such as 3.5 GHz and existing LTE mid bands (e.g. B40, B41) can leverage TDD and LTE-A optimization for FWA scale. Of interest is the technical viability and business case for utilizing LTE for FWA in high-cost rural markets where Fiber is too expensive and copper not viable as a long-term asset.

Mid to longer term 5G technology and spectrum options are emerging quickly, thus increasing the value of Service providers implementing a Fixed Wireless alternative or augmentation to their broadband service offerings. However, to meet the CAF II requirements, 5G mm-Wave technologies might be too expensive to deploy v/s CBRS/LTE wireless as LTE is a mature and well-established technology with economies of scale.

MDU & Enterprise FWB (Fiber Route)




Figure 19 CBRS economics for green field broadband

3 CBRS/LTE FWA Network Design & Deployment

The design considerations for a FWA deployment will need to address the following aspects:

Radio Network Design: Radio network design is the most challenging part of any FWA solution and requires many factors both technical and economical to be considered. While each NSP will have different environments, this paper selects some generic use case and take a wide brush look at this aspect of the solution. Factors such as customer premise equipment, radio node placement, service tier performance and capacity, radio node lease and backhaul costs must all be factored into the decision.

Service Integration in Core Network: Depending on NSP type there can be significant differences. For example, a Mobile Networks Operator (MNO) may have more emphasis on following 3GPP standard integration of the FWA service into their Mobile Packet Core and IP Multimedia Services (IMS) platform where FWA is a network slice leveraging same Mobile assets but with less complexity. Conversely, a Telco NSP may want to further simplify the Packet Core network and integrate into existing Fixed Line core network where FWA looks like just another fixed access.

Integration with SAS and CBRS spectrum coexistence is also a new spectrum sharing mechanism that relies on a close, real-time coordination between network nodes that would need to be proven to work at scale.

3.1 CBRS Radio Dimensioning

Dimensioning a 3.5 GHz Fixed Wireless Access network requires a detailed link budget analysis to answer how far FWA subscribers can be from the CBSD to achieve the performance of the service offering based on how many subscribers the NSP needed to support at different KPI levels. Bandwidth capacity of the radio is based on a combination of LTE-A radio performance, antenna power and technology, radio propagation characteristics of the deployment environment and receiving CPE capabilities.

While a detailed site survey and empirical test data can provide more accurate account of the deployment dimensioning there are several industry modules that utilize empirical data sets that can be used to baseline dimensioning.

The use of CBRS spectrum for a rural fixed broadband service is interesting as it offers better economics and performance than Fiber or Copper investments. CBRS also has the added benefit of time to service and flexibility of much easier network and technology upgrades.

The overall performance of an FWA deployment will depend on the following characteristics:

- Quality of the radio link (e.g. Path Loss, Interference)
- Spectral channel bandwidth
- Broadband service traffic model (UL/DL, QoS etc.)
- Base Station and Antenna characteristics
- Radio optimizations and efficiencies (e.g. MIMO, Carrier Aggregation)

Typical FWA deployments use one of the following options:

3.2 Mounting CBRS to existing Macro Cell structures

In this scenario, the Service Provider utilizes the existing tower infrastructure to add new CBRS macro base stations to provide FWA to households in the area. This approach is beneficial when the primary goal is to provide wide coverage e.g. in very sparsely populated rural areas; such a design is optimized for coverage but lacks capacity.

3.3 Low Power CBSD-B and ODU

In this scenario, the Service Provider adds new low power radio to a smaller tower or to existing street furniture like utility poles to add new CBRS small cells to provide FWA to households in the area. This approach is beneficial when the primary goal is to provide adequate capacity in semi-dense populated areas like suburban towns and cities; such a design is optimized for capacity, but has much smaller coverage radius (typically 25-30 houses per small cell).

The ODU (Outdoor unit) is a standalone external Active LTE antenna that is installed on an outside wall of the customer premise. ODUs are typically powered via a PoE interface. The ODU is installed and mounted directionally to the eNode-B and is connected via Power of Ethernet (PoE) to an indoor unit with PoE injector, or separate PoE injector. The height of the eNode-B and ODU antenna can make a significant difference.

3.4 Quality of the Radio Link

In CBRS deployment various factors like propagation, interference and other aspects of radio quality that can impact the performance of deployment.

3.4.1 Link Budget

Link budget is a measurement of the gains and losses from the transmitter, through the air to the receiver in a wireless communication system. It accounts for the attenuation of the transmitted signal due to propagation, as well as the antenna gains, feedline and miscellaneous losses.

A simple link budget equation is represented as follows:

$$\text{Received Power (dB)} = \text{Transmitted Power (dB)} + \text{Gains (dB)} - \text{Losses (dB)}$$

An example link budget for a sample Base Station and ODU deployment could be shown as:

$$\text{eNode-B to ODU Rx Power} = 33 \text{ dBm} + x \text{ dBi Tx Gain} + 12\text{dB Rx Gain} - \text{Losses (dB)}$$

$$\text{ODU to eNode-B Rx Power} = 23 \text{ dBm} + 12 \text{ dBi Tx Gain} + \text{dBi Rx Gain} - \text{Losses (dB)}$$

3.4.1.1 Losses

Assuming the antennas are in an acceptable Fresnel Zone and line of site ground clearance to the Free Spaces Link Budget, a least squares approximation can be used to factor in loss due to distance without obstacles. Path loss can be defined as the ratio of the transmitted to received power expressed in decibels.

If the estimated received power is sufficiently large (typically relative to the receiver sensitivity), the link budget is said to be sufficient for sending data under perfect conditions. The amount by which the received power exceeds receiver sensitivity is called the link margin.

3.4.1.2 Free-Space Path Loss (FSPL)

In a line-of-sight radio system, losses are mainly due to free-space path loss (FSPL). FSPL is proportional to the square of the distance between the transmitter and receiver as well as the square of the frequency of the radio signal. In other words, free-space path loss increases significantly over distance and frequency.

Other losses in a radio system to consider are due to antenna cabling and connectors. In the case of the high gain ODU there is negligible loss from antenna. In the case of the Base Station which has external antenna a rule of thumb is 0.25dB loss per connector and 0.25dB loss for every 3-ft of antenna cable should be included in the calculation. For a radio system with a 3-ft LMR400 cable and 2 connectors, 0.75dB loss should be included.

FSPL equation is as follows for distance (d) in Km and frequency (f) in MHz:

$$FSPL(dB) = 20\log_{10}(d) + 20\log_{10}(f) + 32.45$$

The table below provides FSPL for 3.5 GHz from 1Km to 12Km

Table 1

Distance	FSPL (dB) @ 3550 MHz
1 Km	103.45
2 Km	109.48
3 Km	112.99
4 Km	115.50
5 Km	117.43
6 Km	119.02
7 Km	120.36
8 Km	121.52
9 Km	122.54
10 Km	123.45
11 Km	124.28
12 Km	125.04

Free Space Link Budget Calculation

Given a Small Cell and ODU assumption deployed in Line of Site the following link budget is possible from 1Km to 12Km..

Table 2

Distance	Tx Power (dBm)	Antenna Gain (dBi)	Rx Gain (dBi)	FSPL (dB) @ 3550 MHz	Link Budget dB
1 Km	33	12	17	-103.45	-41
2 Km	33	12	17	-109.48	-47
3 Km	33	12	17	-112.99	-52
4 Km	33	12	17	-115.50	-55
5 Km	33	12	17	-117.43	-57
6 Km	33	12	17	-119.02	-59
7 Km	33	12	17	-120.36	-60
8 Km	33	12	17	-121.52	-61
9 Km	33	12	17	-122.54	-62
10 Km	33	12	17	-123.45	-63
11 Km	33	12	17	-124.28	-64
12 Km	33	12	17	-125.04	-65

3.4.1.3 Link Margin

Fading due to multipath can result in reduced signal and should be included in the model. A rule of thumb is to maintain a 20 dB to 30 dB of fading margin to compensate.

3.4.1.4 Signal-to- Noise Ratio (SINR)

Higher modulation techniques such as 64-QAM, 256-QAM and MIMO require higher SNR to achieve higher bandwidth capacity on the same carrier. SNR is the ratio of LTE signal to background noise and based on modulation scheme can deliver different data rates.

Conclusion

The shared spectrum model is a first of its kind innovative dynamic, 3-tiered, shared spectrum approach adopted by the FCC for the Citizens Broadband Radio Service. It is a bold and historic shift in spectrum allocation that hopes to combine the best of unlicensed and licensed technologies together. Developing and deploying an effective spectrum sharing mechanism through CBRS would be a significant achievement. It is an exciting opportunity because it makes available a significant amount of spectrum without the need for expensive auctions and is not tied to a particular operator.

It is expected that CBRS will create many new opportunities and revenue sources for MSOs for the new business models and use cases described in this paper. Market analysis shows that MSOs could benefit significantly from the new, shared CBRS bands. 3GPP LTE evolution is the key software enabler while the regulatory framework supports the availability of more spectrum that create economic value for operators. CBRS is an opportunity for the US to demonstrate new technology, business models and inject regulatory innovation.

The proposed opportunities enable MSOs to retain their existing Triple play customers, as well as acquire new customers for a Quad Play service (mobile wireless) and strengthen their overall market position by offering a new and improved personalized mobile broadband data services. An important use case for CBRS is the improvement in building coverage and capacity increase, Local Indoor Mobile Access (LIMA) using LTE based *inside-out* strategy. It has the advantage of being a more Wi-Fi like business model and economics with the Quality of Service of the LTE network. CBRS, thus, opens new business models for in-building wireless solutions. Further, with relaxed the Base Station Transmission power requirements, FCC has helped MSOs to extend the business cases to use CBRS outdoors, leveraging their deep fiber HFC plant.

In terms of relevance to MSOs, this table captures all the specified new business models with their relevance to generate new profit pools for MSOs.

Table 3

Use-Case	Revenue opportunity	Time Horizon	Notes
Local Indoor Mobile Access (LIMA); Inside-Out	++++	2019	By far the most relevant opportunity for MSOs in near term
Outdoor mobile Access using fiber node assets	+++	2019	Leverage the deep fiber and deploy CBRS small cells at strategic optical nodes for outdoor coverage.
Private LTE Networks	++	2019	MSOs offer turnkey private LTE enterprise wireless networks for Medium to Large enterprises.

Use-Case	Revenue opportunity	Time Horizon	Notes
Neutral Host Networks	++	2019	MSOs offer turnkey LTE Radio Access network in the CBRS spectrum for large venues and neutral hosts like hospitality.
Fixed Wireless Access	+	2019	Though MSOs have deep fiber and a cable plant that covers 80% of US homes and businesses, CBRS based FWA can complement the cable plant to provide strategic broadband access in rural areas and the learnings can be extended to future evolution to 5G FWA in urban areas.
Industrial IOT	+	2020	Unlike IOT in unlicensed spectrum, Industrial IOT like in manufacturing plants, refineries, chemical plants etc. requires highly reliable / mission critical wireless communication that can only be provided by a QoS capable technology like LTE. Industrial IOT ecosystems provide MSOs to offer turnkey Industrial IOT networks to automate manufacturing and smart cities.

Abbreviations

CBRS	Citizens Band Radio Service
LIMA	Local Indoor Mobile Access
IOT	Internet of Things
LTE	Long Term Evolution
MVNO	Mobile Virtual Network Operator
FCC	Federal Communications Commission
DoD	Department of Defense
SAS	Spectrum Allocation Server
PAL	Priority Access License
GAA	General Authorized Access
FSS	Fixed Satellite Service
EIRP	Equivalent Isotropically Radiated Power
ESC	Environmental Sensing Capability
CBSD	CBRS Broadband Service Devices
DFS	Dynamic Frequency Selection
EUD	End User Device
PSD	Power Spectrum Density
ITM	Irregular Terrain Model
DPA	Dynamic Protection Area
LAA	LTE Assisted Access
CNO	CBRS Network Operator
ODU	Outdoor Unit

Bibliography & References

- [1] <https://www.ustelecom.org/sites/default/files/files/USTelecom-White-Paper-2.pdf>
- [2] <http://www.analysismason.com/>
- [3] [https://ecfsapi.fcc.gov/file/12280695017091/Final%20Charter%203_5%20GHz%20NPRM%20Comments%20\(12-28-17\).pdf](https://ecfsapi.fcc.gov/file/12280695017091/Final%20Charter%203_5%20GHz%20NPRM%20Comments%20(12-28-17).pdf)
- [4] <https://ecfsapi.fcc.gov/file/104190157420141/IIoT%20Coalition%20ex%20parte%20filed%20041918.pdf>
- [5] <https://www.cnet.com/news/cox-hangs-up-on-cell-phone-service/>
- [6] <https://www.statista.com/statistics/489169/canada-united-states-average-data-usage-user-per-month/>
- [7] <https://www.multichannel.com/news/comcast-charter-benefit-scuttled-sprint-t-mobile-ma-talks-analyst-416755>
- [8] <http://bgr.com/2018/06/04/spectrum-mobile-best-unlimited-plan-2018-vs-verizon/>

[9] <https://www.fiercetelecom.com/telecom/from-comcast-to-hawaiian-telcom-tracking-top-15-residential-broadband-service-providers-q3>

[10] <http://senzafiliconsulting.com/resources-2/the-total-cost-of-ownership-tco-for-fixed-ongo-in-the-3-5-ghz-cbrs-band/>

Characterization of Spectrum Resource Scheduling in FDX DOCSIS

A Technical Paper prepared for SCTE•ISBE by

Tong Liu

Principal Engineer, Office of the CTO
Cisco Systems Inc.

300 Beaver Brook Road, BOXBOROUGH, MASSACHUSETTS 01719, UNITED STATES
978-936-1217
tonliu@cisco.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	4
1. FDX DOCSIS: Potential and Challenges	4
1.1. DS to US Interference	5
1.2. US to DS Interference	6
1.3. FDX Spectrum Resource Scheduling Challenges	7
2. FDX Spectral Efficiency	7
2.1. Interference-Limited Network Model	7
2.2. FDX DS and US SIRs	9
2.3. Interference-Limited Spectrum Resource Characteristics	10
3. Sub-band Directional Assignment.....	11
4. FDX Spectrum Capacity Gain	12
5. FDX Fairness	13
6. An Illustrative Example.....	14
6.1. Physical and RF Topologies	15
6.2. DS and US Spectral Efficiencies.....	16
6.3. Scheduling Decisions.....	17
6.4. Scheduling Performance	18
Conclusion.....	19
Acknowledgement.....	19
Abbreviations	19
Bibliography & References.....	20

List of Figures

Title	Page Number
Figure 1- FDX Resource Block and Allocated Spectrum Options	4
Figure 2 DS to US Interference and Echo Cancellation at the FDX RPD Node	5
Figure 3 Mitigating CM to CM Interference with Simplex-Duplex Rule	6
Figure 4 Interference-Limited Network Model for FDX Operation	8
Figure 5 Correlation Between the DS SIR and US SIR in FDX Spectrum	9
Figure 6 Comparable vs. Incomparable Fully Assigned RBA Sub-band Direction Sets	11
Figure 7 Spectrum Resource Distribution Hierarchy and Load Balancing Options	13
Figure 8 Example - FDX System Physical and RF Topologies	15
Figure 9 Example - Resource Distribution Hierarchy.....	17

List of Tables

Title	Page Number
Table 1 Example - IG Options.....	14
Table 2 Example - DS and US Spectral Efficiencies	15
Table 3 Example - Scheduling Performance	15
Table 1 Example - IG Options.....	15
Table 2 Example - DS and US Spectral Efficiencies	16
Table 3 Example - Scheduling Performance	18

Introduction

Prior to Full Duplex (FDX) DOCSIS, the downstream (DS) and the upstream (US) were scheduled independently, as the DS and the US transmissions were isolated from each other in frequency. In FDX DOCSIS however, the DS and the US operate at the same frequency and at the same time. Coordinated DS and US scheduling is required to avoid interference and to balance the bi-directional traffic need. Fundamentally, FDX spectrum is directionally fluid with a unique set of constraints that confines the DS and US scheduling decisions. How to manage the FDX spectrum resource, maximize the DS and US throughput and maintain fairness is a big design and deployment challenge faced by both Cable Modem Termination System (CMTS) vendors and operators.

In this paper, we tackle this problem by studying the characteristics unique to the FDX spectrum, quantify the optimization objectives, and identify scheduling options to maximize both spectral efficiencies and fairness. The rest of the paper is organized as follows. Section 1 overviews the FDX operation principles and the scheduling constraints. Section 2 then examines the correlations between the DS and the US spectral efficiencies. Section 3 characterizes the FDX spectrum directional assignment and its impact on system throughput. Section 4 quantifies the FDX spectrum capacity gain with respect to the FDD spectrum capacity. Section 5 discusses fairness in the context of the FDX spectrum resource distribution hierarchy. Section 6 provides an illustrative example to demonstrate the FDX spectrum resource scheduling framework. A summary is provided at the end to conclude the paper.

Content

1. FDX DOCSIS: Potential and Challenges

FDX DOCSIS is targeted at significantly increasing the US capacity without sacrificing the DS capacity by enabling simultaneous bidirectional transmissions in a frequency band between 108 MHz and 608 MHz[1][5][6]. Within this band, the RF spectrum allocated for FDX operations is organized in sub-bands, with each sub-band containing one FDX DS channel and one or two FDX US channels, as shown in Figure 1.

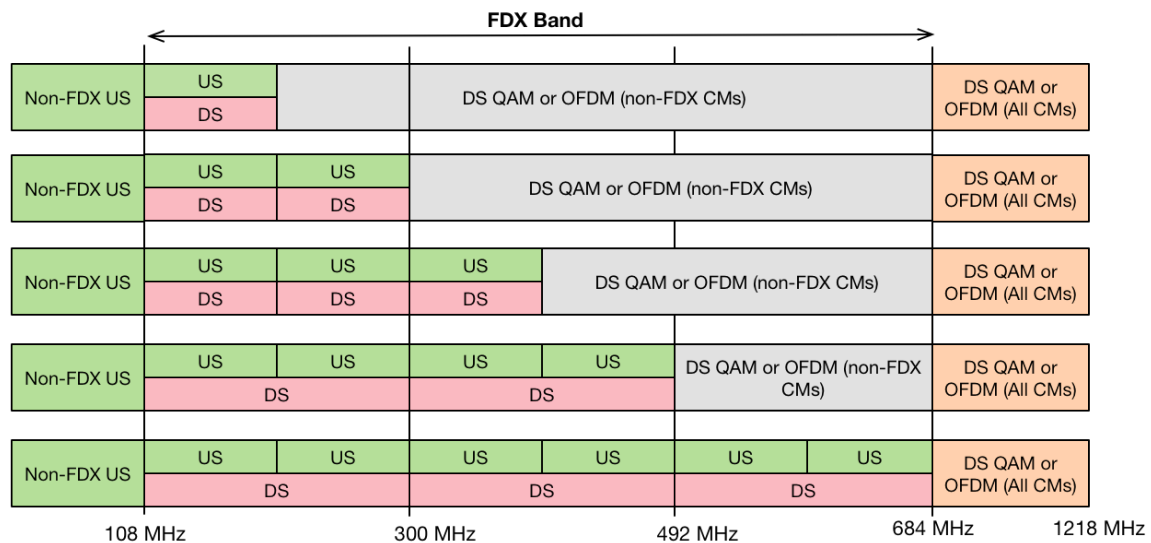


Figure 1- FDX Resource Block and Allocated Spectrum Options

The biggest challenge to enable FDX is the interference from transmitter to receiver at the CMTS (DS to US interference) and among cable modems (CM)s (US to DS interference), as described below.

1.1. DS to US Interference

At an FDX remote PHY device (RPD) node the transmitted DS signal, which has a much higher signal level than the received US signal, can be echoed back to the US receive path through the internal and external coupling. Such interference will completely wipe out the received US signal if there is not sufficient isolation between the transmitter and the receiver. Normally, a diplexer filter is used to keep the DS signal from entering the US receive path. However, the FDX RPD Node will not have a diplexer in order for it to receive the US signal in the same spectrum sent by the CMs. Instead, an echo canceller (EC) is used to suppress the interference sourced from the DS transmissions and thus provide the required isolation[3]. The EC removes the interference by reconstructing the echoes from the transmitted DS signals based on a proper estimation of the echo channel characteristics, as shown in Figure 2.

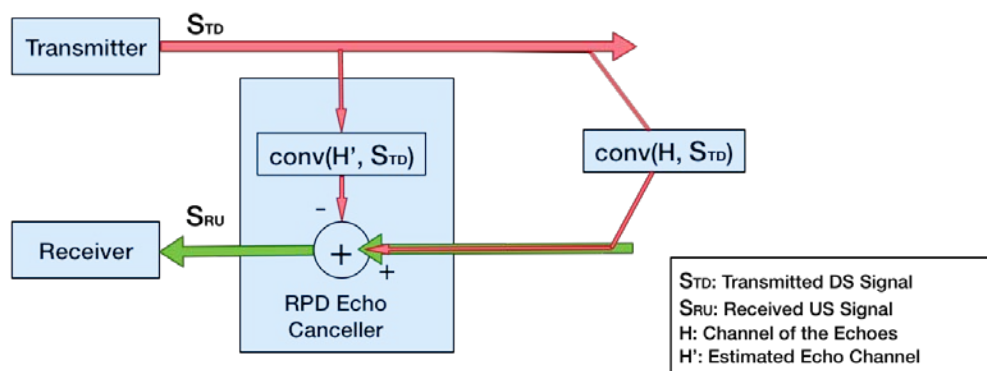


Figure 2 DS to US Interference and Echo Cancellation at the FDX RPD Node

1.2. US to DS Interference

From the CM perspective, the FDX spectrum will still appear to be frequency division duplexed (FDD), in the sense that, a sub-band can only be used by the CM in either DS or US direction at a given time. FDX is achieved by assigning one set of CMs to use the sub-band for US at the same time that a different set of CMs is being assigned to use that sub-band for DS. The main sources of the US to DS interference experienced by a CM are internal, from its own US transmission in an adjacent sub-band, and external, from neighboring CMs US transmissions. While the echo cancellation technique can remove the internal self-interference, it cannot be used to mitigate the external interference, as the source of the interference is unknown to the receiving CM.

The external CM to CM interference is due to the imperfect isolation in the cable plant. If one CM is transmitting in the US in one sub-band while another CM is trying to receive in that same sub-band, energy from the first CM's US transmission can, in some cases, leak into the location of the second CM and prevent it from successfully receiving the DS transmissions. Such interference can be avoided with proper scheduling such that CMs that interfere with each other will not transmit and receive at the same time and at the same frequency. This mechanism is referred to as the “simplex-duplex” rule required for FDX operation. Specifically, FDX DOCSIS separates a CM's neighboring CMs into two categories based on the CM's interference group (IG)[2][4]:

1. CMs within the IG, which need to operate in simplex mode (uni-directional at any particular point of time and frequency) with the respect to the CM to avoid co-channel interference;
2. CMs outside the IG, which have an isolation above the IG's boundary (a threshold that bounds the isolation among CMs inside the IG) and can operate in full duplex mode with respect to the CM.

In this nomenclature, simplex refers to a unidirectional path at a given instance of time and frequency, and duplex refers to a bidirectional path at a given instance of time and frequency.

In FDX DOCSIS, the “simplex-duplex” operation rule is enforced via the IG-TG-RBA mapping hierarchy as shown in Figure 3. IGs are grouped into different transmission group (TG)s for scheduling purposes. Each TG is associated with a resource block assignment (RBA) that enforces FDD within the TG. Different RBAs enable bidirectional use of the FDX spectrum.

Figure 3-b shows the RBA directional assignment options with three FDX sub-bands. RBAs with complement DS and US assignment form the minimum RBA sub-band direction sets that fully assign the spectrum in both directions. When more than two TGs are used, a sub-band will have at least two TGs sharing the sub-band in the same direction.

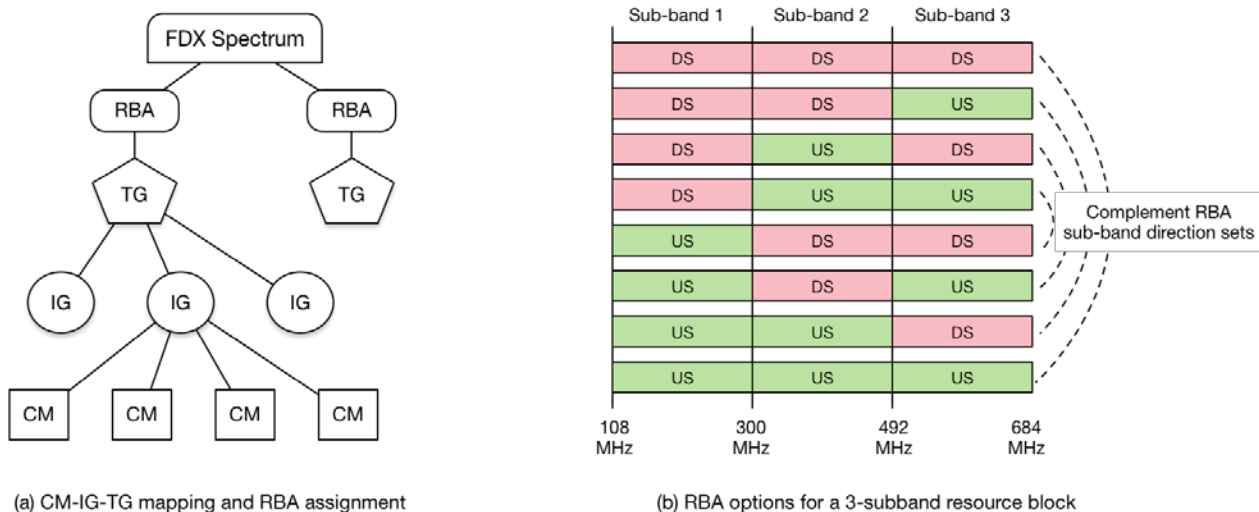


Figure 3 Mitigating CM to CM Interference with Simplex-Duplex Rule

1.3. FDX Spectrum Resource Scheduling Challenges

The FDX operation imposes a unique set of challenges to spectrum resource scheduling. First, as described in more detail in Section 2, the DS and US spectral efficiencies can no longer be separately controlled, as the intended transmit signal of one direction is the source of the interference to the other direction. Second, the DS and US resource blocks accessible by a CM are constrained by the simplex-duplex rule through a multi-level resource distribution chain. Third, the IG based spectrum sharing may conflict with the ability to maintain service level agreement (SLA) goals, as the interference environment is independent to the traffic distribution and service offerings.

Due to the DS to US correlation, and the conflicting optimization goals of efficiency and fairness, FDX spectrum resource scheduling becomes a highly integrated problem set with multiple sub-problems that need to be individually analyzed.

2. FDX Spectral Efficiency

The objective of this section is to analyze the achievable DS and US spectral efficiencies when the interference is a dominating factor affecting system performance. Due to the path loss of the coax plant, the DS and US transmit power ends being much higher than the converse receive power, resulting in a high level of interference at the receiver. When interference is a dominating factor impacting system performance, we can use signal to interference ratio (SIR) to evaluate the spectral efficiency. If the background noise is not negligible, the SIR serves as an upper bound of the achievable spectral efficiency.

2.1. Interference-Limited Network Model

To analyze the impact of the interference, we model the FDX system as a two-port (a kind of four-terminal) network connecting the transmitters and the receivers of both the DS and US directions, as

shown in Figure 4. The directional connections of the DS and US signals and the interference paths in between can be characterized with the following parameters:

- Plant Path Loss

This is the path loss along the intended DS or US transmission path through the coax plant. For a N+0 network, the plant loss is reciprocal between the DS and the US.

- DS to US Isolation

This is the total isolation from the DS transmitter to the US receiver at the FDX RPD node. It combines the passive DS to US coupling loss, X_{DU} , and the DS echo return loss, G , provided by the EC.

- US to DS Isolation

This is the isolation between a pair of CMs attached to the same coax plant. In FDX operation, the IG boundary is the lower bound of the US to DS isolations, X_{UD} for its member CMs with respect to rest of the CM population.

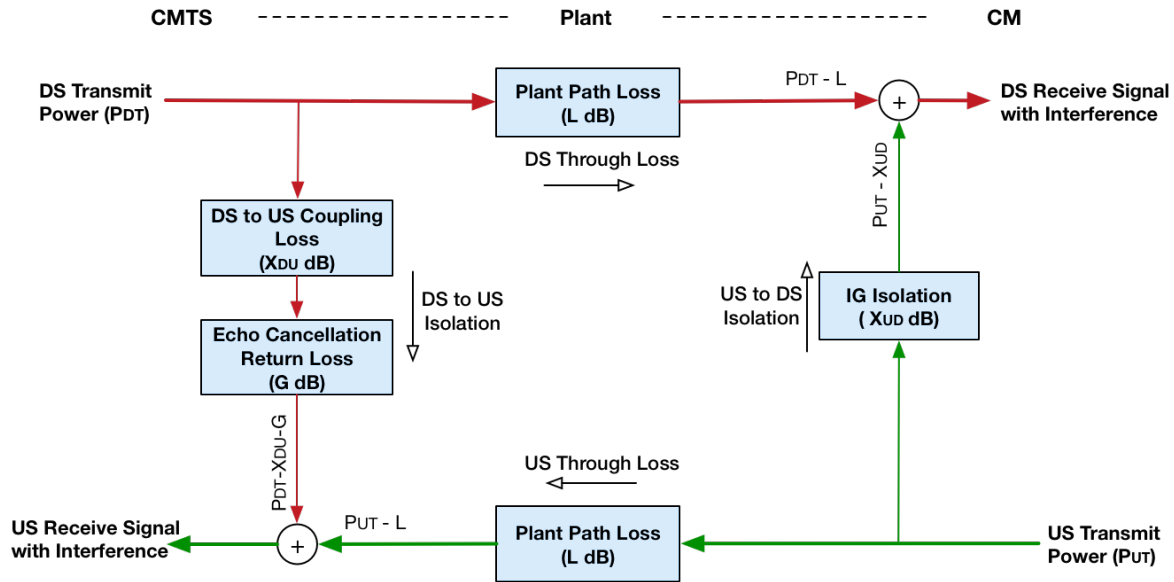


Figure 4 Interference-Limited Network Model for FDX Operation

2.2. FDX DS and US SIRs

Based on the network model, the DS and US SIR can be easily setup using logarithmic power units (dB) as below:

$$SIR_U = P_{UT} - L - P_{DT} + X_{DU} + G \quad Eq\ 1$$

$$SIR_D = P_{DT} - L - P_{UT} + X_{UD} \quad Eq\ 2$$

Where,

P_{UT} : US transmit power per 6 MHz in dBmV

P_{DT} : DS transmit power per 6 MHz in dBmV

X_{DU} : DS to US isolation in dB

X_{UD} : US to DS isolation in dB

G : DS echo return loss in dB

L : DS or US path loss in dB

Eq1 and Eq2 reveal the correlation between the DS and the US spectral efficiencies in terms of SIRs. To visualize this correlation, the DS and the US SIR values are plotted in Figure 5 as the DS transmit and US receive power change. Other parameters assumed are listed below:

X_{DU} : DS to US isolation = 20 dB

X_{UD} : US to DS isolation = 65 dB

G : DS echo return loss = 50 dB

L : DS or US signal propagation loss = 33 dB

Note that the above values are illustrative only.

From Figure 5, we can observe that as the US transmit power increases, the US SIR improves, but the DS SIR declines; and as the DS transmit power increases, dB by dB, the DS SIR increases, the US SIR decreases. The sum of the DS SIR and US SIR however remains constant regardless of the power level.




Figure 5 Correlation Between the DS SIR and US SIR in FDX Spectrum

2.3. Interference-Limited Spectrum Resource Characteristics

When the system operates in the interference limited regime, the FDX spectrum efficiency exhibits following properties:

- Monotonicity

When the transmit power increases in one direction, the spectral efficiency decreases in the opposite direction. This implies that we can trade off the DS spectral efficiency against the US spectral efficiency with proper transmit power adjustment in either direction.

- Exchangeability

Adjustments are exchangeable among elements affecting the spectral efficiency. For example, as shown below, increasing US transmit power is equivalent to decreasing the DS transmit power with respect to the US SIR; increasing the DS transmit power is equivalent to increasing the IG isolation boundary with respect to the DS SIR,

$$SIR_U = P_{UT} \uparrow - P_{DT} \downarrow + X_{DU} + G - L$$

$$SIR_D = P_{DT} \uparrow - P_{UT} + X_{UD} \uparrow - L$$

This property provides the flexibility required to achieve the spectral efficiency target under certain operational constraints. For example if the DS transmit power is limited, a larger IG can be used to get to the higher DS SIR desired.

- Conservation

Conservation is a direct result of the monotonicity. For given signal path loss and isolation loss between the DS and the US, the sum of the DS and US SIR remains constant. This can be seen by adding Eq1 with Eq2 as below:

$$SIR_U + SIR_D = X_{UD} + X_{DU} + G - 2L \quad \text{Eq 3}$$

Eq3 makes sense intuitively. When the interference dominates, less signal loss and better isolation result in better spectral efficiency.

3. Sub-band Directional Assignment

In FDX operation, the DS or US sub-band direction assignment must follow the “simplex-duplex” rule such that a sub-band is only used in a single direction for a particular TG. Between the TGs, however, there are no direction restrictions, so when one TG uses a sub-band for US operation, a different TG can simultaneously use the same sub-band for DS operation, doubling the usage of the spectrum.

In FDX DOCSIS, the sub-band direction assignment is done through RBA to TG mapping and at any given time a TG can only be associated with one RBA. To examine the impact of sub-band direction assignment on system performance, we derive the maximum achievable FDX DS and US throughput as below.

Assuming there are M TGs sharing N sub-bands, we model the per TG per sub-band direction assignment as a $M \times N$ matrix for each direction shown as below, where A^D denotes the DS assignment and A^U denotes the US assignment, a_{ij}^D and a_{ij}^U are the directional assignment coefficients for TG_i in j th sub-band of the corresponding US and DS directions.

$$\begin{array}{l} \text{Assignment of the DS: } A^D = \begin{array}{c} \xrightarrow{N \text{ Sub-bands}} \\ \begin{bmatrix} a_{11}^D & \cdots & a_{1N}^D \\ \vdots & \ddots & \vdots \\ a_{M1}^D & \cdots & a_{MN}^D \end{bmatrix} \\ \downarrow \\ M \text{ TGs} \end{array} \\ \\ \text{Assignment of the US: } A^U = \begin{array}{c} \xrightarrow{N \text{ Sub-bands}} \\ \begin{bmatrix} a_{11}^U & \cdots & a_{1N}^U \\ \vdots & \ddots & \vdots \\ a_{M1}^U & \cdots & a_{MN}^U \end{bmatrix} \\ \downarrow \\ M \text{ TGs} \end{array} \end{array}$$

The direction assignment coefficient varies between 0 and 1, representing different assignment scenarios described below:

- 0 : The given TG is not assigned to use the given sub-band in the given direction
- 1: The given TG is assigned to use the full sub-band in the given direction
- Between 0 and 1: The given TG is assigned to use a fraction of the given sub-band in the given direction. This will be the case when multiple TGs share the same sub-band in the same direction.

The following restrictions apply to the sub-band direction assignment:

- A sub-band cannot be used for both DS and US for a given TG at the same time. This implies the entry wise product (also known as the Hadamard product) of the DS and US assignment matrices is zero:

$$a_{ij}^D .* a_{ij}^U = 0 \quad \text{Eq4-1}$$

- The sum of the assignments of a sub-band in a given direction across all TGs cannot exceed 100% of the sub-band spectrum in the corresponding direction. The sub-band is claimed to be fully assigned in both DS and US directions if the following conditions are met:

$$\sum_{i=1}^M a_{ij}^D = 1; \quad \text{Eq4-2}$$

$$\sum_{i=1}^M a_{ij}^U = 1; \quad \text{Eq4-3}$$

Eq4-1 to Eq4-3 imply that a minimum of two TGs ($M \geq 2$) are required for a sub-band to be fully utilized in the FDX operation. For the special case of two TGs, the directional assignment coefficients a_{ij}^D and a_{ij}^U are either 1s or 0s, and the sub-band direction sets of the two TG need to be complement to each other for each sub-band to be fully assigned. Figure 6 shows the two TG and three TG mapping examples that have the three sub-bands fully assigned.




Figure 6 Comparable vs. Incomparable Fully Assigned RBA Sub-band Direction Sets

An additional observation is that when there is an odd number of sub-bands fully assigned to two TGs, the numbers of the DS and US sub-bands are incomparable between the TGs. For example, in Figure 6-a, TG1 has two DS sub-bands and one US sub-bands, while TG2 has two US sub-bands and one DS sub-bands. For the three sub-band case, comparable directional assignment can be achieved however with three TGs instead, with each sub-band shared by at most two TGs in each direction. An example of the comparable directional assignment is shown in Figure 6-b.

4. FDX Spectrum Capacity Gain

The FDX spectrum capacity is defined here as the aggregate DS and US throughput achievable across all TGs over the allocated FDX spectrum. For an individual TG, denote the achievable DS and US throughput across the N sub-bands as,

DS max throughput set for TG_i : $D_i = [d_{i1}, \dots, d_{iN}]$, for $i = 1, \dots, M$

US max throughput set for TG_i : $U_i = [u_{i1}, \dots, u_{iN}]$, for $i = 1, \dots, M$

where d_{ij} and u_{ij} are determined by the DS and US spectral efficiencies of the member IGs included in TG_i and the sub-band width.

After applying RBAs, the FDX spectrum capacity collectively achievable with M TGs can then be expressed as:

$$\text{FDX Spectrum Capacity} = \sum_{i=1}^M (A_i^D \cdot D_i + A_i^U \cdot U_i) \quad \text{Eq5}$$

where, A_i^U and A_i^D denote the i th row of the RBA direction assignment matrices A^U and A^D .

We then define the FDD spectrum capacity as a reference to evaluate the FDX spectrum capacity gain. Assuming the maximum DS or US throughput in the non-FDX mode across the spectrum span of the N sub-band as:

$$\text{DS max throughput when spectrum is used in DS only: } D_0 = [d_{01}, \dots, d_{0N}]$$

$$\text{US max throughput when spectrum is used in US only: } U_0 = [u_{01}, \dots, u_{0N}]$$

Further assuming an even split of the DS and US spectrum in FDD, the FDD spectrum capacity is:

$$\text{FDD Spectrum Capacity} = \sum_{j=1}^N (d_{0j} + u_{0j})/2 \quad \text{Eq6}$$

Hence, the FDX spectrum capacity gain is:

$$\text{FDX Spectrum Capacity Gain} = \frac{\sum_{i=1}^M (A_i^D \cdot D_i + A_i^U \cdot U_i)}{\sum_{j=1}^N (d_{0j} + u_{0j})/2} \quad \text{Eq7}$$

Eq7 can serve as a benchmark to evaluate an FDX system performance. Maximizing it combines the FDX spectral efficiency optimization with the RBA sub-band directional assignment optimization. However, maximizing Eq7 alone will result in biased spectrum allocations in favor of the less interfered users, fairness also needs to be considered as discussed below.

5. FDX Fairness

A typical dilemma in resource scheduling is the conflicting optimization objectives between throughput and fairness. In FDX DOCSIS, maximizing the throughput alone does not necessarily translate into the maximum value for the operators, as it may favor the least interfered subscribers who may not be subscribed to the highest service levels. Fairness is needed to retain happy customers with an allocation proportional to the SLAs.

Service level fairness is a global goal irrelevant to the localized constraints or scheduling decisions. The fewer constraints or the fewer local scheduling points, the easier it is to achieve fairness. FDX DOCSIS, however, has many localized scheduling points along the resource distribution hierarchy. Achieving fairness requires resource or traffic balancing at each of the branching point. As shown in Figure 7, from bottom up, the following options can be used to optimize fairness:

- A CM can have a different spectral efficiency in a direction by adjusting the transmit power or shrinking / extending its IG boundary. The spectral efficiency change results in different DS and/or US data profile settings.

- By shrinking the IG boundary to a lower spectral efficiency, an IG can split into smaller IGs, which can then be mapped into different TGs for load balancing purposes.
- A TG can be associated with a different RBA in order to achieve a different DS to US capacity ratio.
- The legacy DS and US spectrum resource, accessible to all the CMs, can be used as a fluid resource to improve fairness across the system.

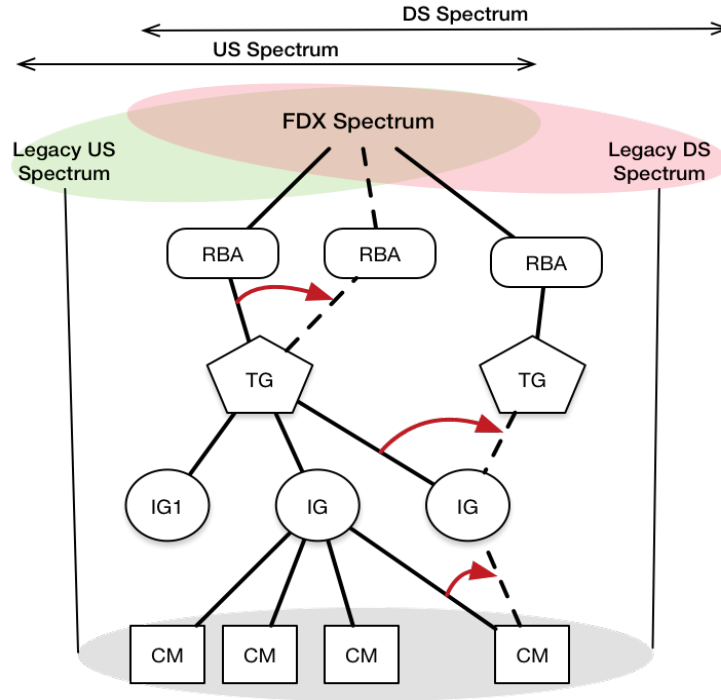


Figure 7 Spectrum Resource Distribution Hierarchy and Load Balancing Options

To evaluate fairness, a proper fairness measure is required. Choices of the fairness measure may include throughput, latency, or a quantifier expressing the subscribers' quality of experiences. More study is required in this area to identify a proper fairness measure for the high-end FDX users, as new applications may appear to take advantage of the FDX spectrum. When best effort service and uniform traffic distribution are assumed, Jain's fairness index [7] can be used to describe the relative equality of the average allocation per user, quoted below:

$$f(x) = \frac{[\sum_{i=1}^n x_i]^2}{n \sum_{i=1}^n x_i^2} \quad x_i \geq 0 \quad Eq8$$

where $f(x)$ is the fairness index, ranging from $1/n$ (worst case) to 1 (best case), and the parameter x is the average throughput per user. The fairness index is maximum when all users receive the same allocation.

6. An Illustrative Example

In this section, we put everything together with an example to illustrate the FDX spectrum resource scheduling framework.

6.1. Physical and RF Topologies

Without loss of generality, we use a single leg N+0 network as shown in Figure 8. The FDX RPD Node has one port connected to 24 FDX CMs evenly across 6 taps. Other legs can of course present, but not considered to be significant for the purpose of this example.

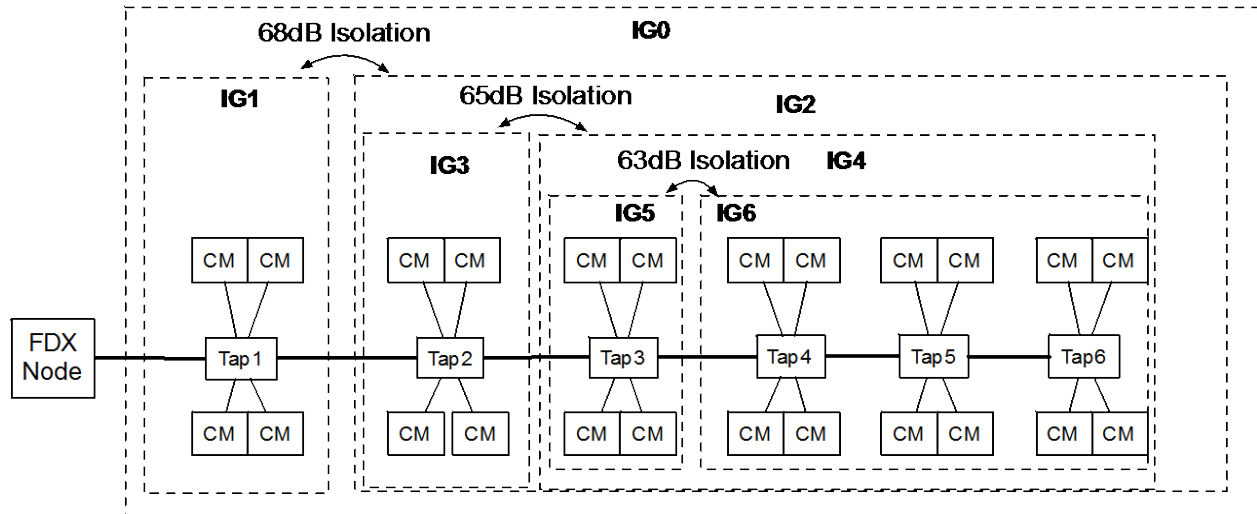


Figure 8 Example - FDX System Physical and RF Topologies

The 24 CMs are separated into 4 IGs, namely IG1, IG3, IG5 and IG6 as marked in Figure8. Note that the IG selection is part of the scheduling decision intended to optimize both throughput and fairness. Table 1 lists all IG options available.

Table 1 Example - IG Options

IG Index	IG Scope	DS Spectral Efficiency (bits/subcarrier)		
		Sub-band 1	Sub-band 2	Sub-band 3
IG0	Tap1 to Tap6	11.5	10.5	9.5
IG1	Tap1	11	10	9
IG2	Tap2 to Tap6			
IG3	Tap2	10	9	8
IG4	Tap3 to Tap6			
IG5	Tap3	9	8	7
IG6	Tap4 to Tap6			

The available DS and US spectrum include:

- Legacy upstream 5-85 MHz, estimated around 400 Mbps with a mix of single carrier quadrature amplitude modulation (SC-QAM) and orthogonal frequency division multiplexing with multiple access (OFDMA) channels, around 400Mbps
- FDX Spectrum 108 – 684 MHz, with 3 sub-bands of 192MHz width. The maximum directional capacity is estimated around 4.8 Gbps if entire band is used for DS operation or US operation.
- Legacy downstream above 684 MHz, estimated around 2Gbps with a mix of SC-QAM and orthogonal frequency division multiplexing (OFDM) channels.

6.2. DS and US Spectral Efficiencies

The DS and the US spectral efficiencies are listed in Table 2, together with other performance impacting parameters. To make the example more interesting, the DS and the US transmit power levels are up-tilted to compensate the increasing path loss across the three sub-bands.. In this example, the background noise is assumed to cause a 3dB degradation to the spectral efficiency with respect to the SIR in either DS or US direction. Note that, in actual FDX operation, the DS and US spectral efficiencies can be determined based on the receive modulation error ratio (RxMER) measurements, during and after IG discovery.

Table 2 Example - DS and US Spectral Efficiencies

	Operation Parameters	Sub-band 1	Sub-band 2	Sub-band 3
Common to all IGs	US Tx Power (dBmV/6MHz)	33	36	39
	DS Tx Power (dBmV/6MHz)	34	37	40
	Path Loss (dB)	30	33	36
	DS to US Coupling Loss (dB)	20	20	20
	DS Echo Processing Loss (dB)	50	50	50
IG1, IG2	IG Boundary (dB)	68	68	68
	DS SIR (dB)	39	36	33
	US SIR (dB)	39	36	33
	DS Spectral Efficiency (Bits/subc)	11	10	9
	US Spectral Efficiency (Bits/subc)	11	10	9
IG3, IG4	IG Boundary (dB)	65	65	65
	DS SIR (dB)	36	33	30
	US SIR (dB)	39	36	33
	DS Spectral Efficiency (Bits/subc)	10	9	8
	US Spectral Efficiency (Bits/subc)	11	10	9
IG5, IG6	IG Boundary (dB)	62	62	62
	DS SIR (dB)	33	30	27
	US SIR (dB)	39	36	33
	DS Spectral Efficiency (Bits/subc)	9	8	7
	US Spectral Efficiency (Bits/subc)	11	10	9

6.3. Scheduling Decisions

In this example, best effort service is offered to all 24 subscribers in both the DS and US directions. All 24 CMs are active with even traffic distributions. Per CM peak rate is limited to 3 Gbps for the DS and 1 Gbps for the US. All CMs are active with even traffic distributions.

Figure 9 shows a snapshot of the resource distribution hierarchy with the set scheduling decisions described below:

- IG1, IG3, IG5 and IG6 are selected with the DS and the US data profiles set to the achievable FDX DS and US spectral efficiencies.
- IG1 and IG3 are included in TG1 with a total of 8 CMs. IG5 is included in TG2 with a total of 4 CMs, and IG5 is included in TG3 with a total of 12 CMs.
- A set of symmetrical RBA sub-band direction sets are selected with respect to the three TGs, with 2 DS sub-bands and 1 US sub-bands allocated to each TG as shown in Figure 9. This arrangement together with legacy spectrum will allow the FDX system to meet the 3 Gbps DS and 1 Gbps US SLA requirements.

The average RBA directional allocation coefficients are listed below. Note that for the US direction, each sub-band is used exclusively by one TG; while for the DS direction, each sub-band is shared by two TGs. The DS assignment coefficients used in this example are based on proportional scheduling and the assumption of uniform traffic distribution across all CMs in all TGs.

$$\text{Average US Assignment Coefficient : } A^U = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$$

$$\text{Average DS Assignment Coefficient : } A^D = \begin{bmatrix} 2/3 & 0 & 2/5 \\ 1/3 & 1/4 & 0 \\ 0 & 3/4 & 3/5 \end{bmatrix}$$

- The legacy DS and US spectrums are used in this example to balance the FDX spectrum capacity allocation differences among the TGs.

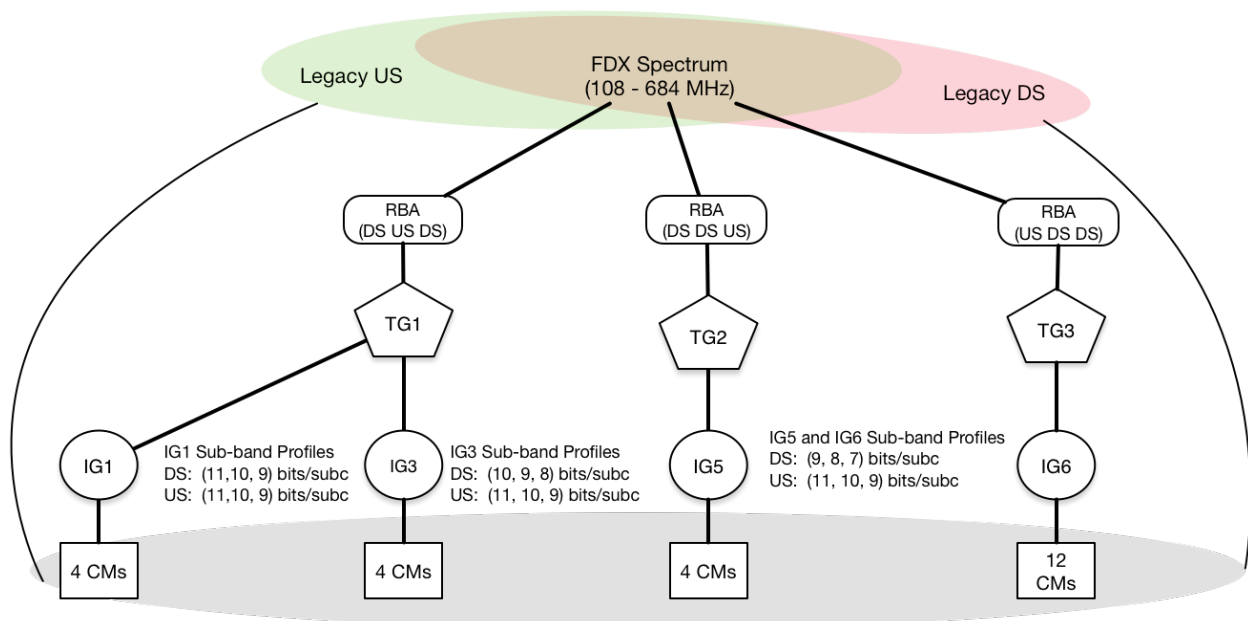


Figure 9 Example - Resource Distribution Hierarchy

6.4. Scheduling Performance

The scheduling performance in terms of throughput and fairness are calculated based on the above scheduling decisions. The FDX spectrum capacity gain (Eq7) is used to evaluate the overall FDX throughput performance, where the FDD spectrum capacity is calculated by assuming the FDD DS or US spectral efficiency is 0.5 dB higher than the FDX DS or US spectral efficiency.

The fairness performance is evaluated with Jain's fairness index (Eq8) with respect to the per CM average DS and US throughput. As indicated in Table 3, the US is less fairly shared due to the unbalanced CM population among the IGs, and the common spectrum resource, in this case legacy US spectrum, is insufficient to compensate the resultant traffic load differences.

Table 3 Example - Scheduling Performance

		TG1	TG2	TG3
Per TG Average FDX Capacity (Gbps)	DS	2.06	0.77	1.57
	US	1.54	1.38	1.69
FDX Spectrum Capacity Gain		1.86		
Legacy Capacity (Gbps)	DS	2		
	US	0.4		
Per TG Max Capacity (Gbps)	DS	4.9	4.60	4.3
	US	1.9	1.8	2
Per TG CM Average (Mbps)	DS	241	234	247
	US	192	345	174
Jain's Fairness Index	DS	0.99		
	US	0.91		

Conclusion

An optimal FDX spectrum scheduler maximizes the FDX spectrum capacity while maintaining fairness among subscribers based on the DS and US SLAs. It balances the DS and the US capacities with traffic loads by adjusting the DS and US spectral efficiencies, the DS and US spectrum widths, and the competing CM population. To do this, it must follow the FDX spectrum resource distribution hierarchy enforced by the FDX DOCSIS operation, which involves RBA, TG, IG and data profile assignments. Essentially, the FDX spectrum scheduling is a process to identify the best paths along the distribution hierarchy connecting the spectrum resource to the CMs that achieve the optimization objectives.

This paper characterizes the FDX spectrum scheduling by analyzing the bidirectional spectrum capacity, the resource sharing constraints, the optimization objectives and the delivery mechanism. We observed that when system is interference limited, the DS and US spectral efficiency is inversely related, and the combined spectral efficiency remains constant relative to the passive plant topology and the active processing gain for interference suppression. We showed how this property can be used to balance the DS and US performance under various operation constraints. We further modeled the FDX spectrum resource sharing mechanism as a single rooted scheduling tree, and revealed the resource sharing fluidity between the DS and the US, and among all the distribution points. We explained how the utilities defined in FDX DOCSIS can be applied for the resource sharing purpose, including RBA allocation, TG assignment, IG discovery and data profile settings. We also defined a set of performance metrics that could be used to evaluate system performance and guide optimization. An example was present at the end to illustrate the FDX spectrum resource scheduling process.

Acknowledgement

The author would like to thank John Chapman and Hang Jin of Cisco for their support and contributions to this paper.

Abbreviations

CM	cable modem
CMTS	Cable Modem Termination System
DS	downstream
FDD	frequency division duplexed
FDX	full duplex
IG	interference group
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiplexing with multiple access
RBA	resource block assignment
RxMER	receive modulation error ratio
SC-QAM	single carrier quadrature amplitude modulation
SIR	signal to interference ratio
SLA	service level agreement
TG	transmission group

Bibliography & References

- [1] John T.Chapman, Hang Jin (2016). *Full Duplex DOCSIS*, INTX 2016, May 18, 2016
- [2] Tong Liu, John T.Chapman, Hang Jin (2016). *Interference-Aware Spectrum Resource Scheduling for FDX DOCSIS*, SCTE 2016 Journal
- [3] Hang Jin, John T. Chapman. *Echo Cancellation Techniques for Supporting Full Duplex DOCSIS*, SCTE 2017
- [4] Tong Liu *IG Discovery for FDX DOCSIS FDX* , SCTE 2017
- [5] DOCSIS MULPIv3.1 *DOCSIS3.1 MAC and Upper Layer Protocols Interface Specification CM-SP-MULPIv3.1* May 9, 2018
- [6] DOCSIS PHYv3.1 *DOCSIS 3.1, Physical Layer Specification, CM-SP-PHYv3.1*, May 9, 2018
- [7] Jain, R “A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer System” DEC Research Report TG-301

Coherent Access Applications for MSOs

A Technical Paper prepared for SCTE•ISBE by

Harj Ghuman

Network Architecture & Technology Strategy
Cox Communications
6305 Peachtree Dunwoody Road, Atlanta, GA, 30328
404-269-8547
Harj.ghuman@cox.com

David Job

Principal Engineer
Cox Communications
6305 Peachtree Dunwoody Road, Atlanta, GA, 30328
404-269-5126
David.Job@cox.com

Table of Contents

Title	Page Number
Table of Contents	2
1. Introduction.....	4
2. What is Coherent Optics?	4
2.1. Coherent Time Line.....	5
2.2. Modulation In Coherent Systems	5
2.3. FEC in Coherent Systems.....	6
2.4. Coherent DWDM Grid and OSNR Requirements	7
2.5. Coherent Equipment Type	8
2.6. Coherent Bi-directional Transmission Capability	8
3. OCML – MDM Network	9
3.1. 10G and Coherent Coexistence.....	10
3.2. Impact of 10G and DCMs on Coherent 100G/200G	11
3.2.1. Cross Phase Modulation (XPM) In Mixed 10G/Coherent Signals	11
3.2.2. Guard Bands In Mixed 10G and Coherent 100G/200G	12
4. Testing of Coherent and 10G Coexistence in OCML.....	12
4.1. Test Set-Up	13
4.2. Tests Performed.....	14
4.3. Test Results	15
4.4. Test Conclusions.....	16
5. Coherent Access Applications	17
5.1. Converged Cable Access Network	17
5.2. Remote-PHY Coherent Optical Trunk.....	17
5.2.1. Muxponder	18
5.2.2. Ethernet Switch	18
5.2.3. Router	18
5.3. RPhy 10G DWDM Optical Trunk Transition to Coherent.....	19
5.4. Coherent Business Service Applications	19
5.5. Coherent Hub Consolidation	20
6. Conclusion.....	22
Abbreviations	22
Bibliography & References.....	23
Acknowledgements	23

List of Figures

Title	Page Number
Figure 1 - Coherent Deployment Timeline	5
Figure 2 - Coherent Modulation and Reach.....	6
Figure 3- Pre-FEC and Post-FEC Operational Schematic	6
Figure 4– Uni and Bi-directional Coherent Transceivers	9
Figure 5- OCML - MDM Network	10
Figure 6 - Remote-PHY High Level Solutions Architecture	10
Figure 7 - Signal Power Density 10G and 100G.....	12
Figure 8 - Guard Band Example	12
Figure 9 - Test Set up for 10G and Coherent 100G/200G Coexistence.....	14

Figure 10 - Test Wavelengths (ITU Ch Pairs) and Optical Signal Types	14
Figure 11 - Link Margin Test Results	16
Figure 12 - Converged Cable Access Network.....	17
Figure 13 - Remote-PHY Physical Distribution.....	18
Figure 14 - 10G DWDM Optical Trunk Transition to Coherent.....	19
Figure 15 - Coherent Business Services	20
Figure 16 - Coherent Hub Consolidation	21

List of Tables

Title	Page Number
Table 1 - Coherent Optical Characteristics For Typical CFP2 – DCOs	7

1. Introduction

While Coherent optics are ubiquitous in the MSO metro core and backbone optical networks, initial planning and use cases are now being investigated for Coherent technology in the Access domain. Cable HFC optical networks are rapidly evolving from traditional analog to digital DWDM infrastructures that enable the adoption of Distributed Access Architectures (DAA) such as Remote-PHY and Remote MAC, which provide higher performance, lower overall optical costs and simplified network designs.

The available digital DWDM infrastructure within the Access network will accommodate the pervasive 10G NRZ DWDM connections associated with DAAs as well as the introduction of Coherent technology to address high bandwidth applications (i.e. MDU service delivery and Enterprise Business services.) In that regard we detail proof of concept testing over 40 and 60 km distances with bi-directional Coherent 100G/200G and 10G NRZ DWDM wavelengths placed adjacent to each other on the same fiber, using our Cox designed Optical Communications Module Link extender (OCML) and a Mux/DeMux. Use cases of a converged Access network consisting of 100G/200G Coherent, 10G DWDM and PON are presented, allowing operators to deploy a solution which is technology agnostic and able to support a multitude of services over the same network, while fully utilizing existing fiber assets.

2. What is Coherent Optics?

Coherent was first unveiled at OFC 2008 and began to be deployed in long-haul optical networks around the world in 2010. Coherent-based technology today is the de-facto standard for high-speed long-haul optical transport networks at 100G and beyond. Most long haul networks have already been deployed at 200G, with 400G emerging. Eighty eight wavelengths at 200G equate to an enormous capacity of 18 terabits per second which can be transported over distances of up to 2000 km.

Coherent optics were developed to overcome two fundamental limitations as you go beyond 10G: chromatic dispersion (CD) and polarization mode dispersion (PMD). Network capacities are increasing by a 25 to 50 percent Compound Annual Growth Rate (CAGR) every year, and systems running at 10 Gb/s cannot keep up with this kind of growth. At its most basic level, Coherent optical transmission is a technique that uses modulation of both the light amplitude and phase, as well as transmission across two polarizations to transport considerably more information through a fiber optic cable. At the receiver, a local oscillator (LO) measures the phase and amplitude, while a digital signal processor (DSP) recovers and demodulates the phase information. Phase modulation also reduces susceptibility to optical impairments, thus Coherent systems have improved OSNR tolerance, and better reach with multiple advanced modulation formats such as DP-QPSK, DP-8QAM and DP-16QAM being utilized.

Advanced Coherent optical technology has a number of key attributes, including:

- High-gain soft-decision Forward Error Correction (FEC), which enables signals to traverse longer distances while requiring fewer regenerator points which also provides more margin, allowing higher bit-rate signals to traverse farther distances.
- Coherent Digital Signal Processors (DSPs) account for dispersion effects after the signal has been transmitted across the fiber, including compensating for both CD and PMD and also improved tolerances for Polarization-Dependent Loss (PDL). The use of Dual Polarity (DP) modulation in the optical domain also doubles the effective bit rate for many Coherent systems.

2.1. Coherent Time Line

As shown in Figure 1, the introduction of Coherent Optics enabled the “10G Speed Limit” to be broken, initially for 40G and soon after for 100G long-haul transmission. Figure 1 represents the commercial deployment of products. By 2010 to 2011, the technologies had reached a point of market maturity at which they could genuinely allow 100G Coherent signals to be sent over the same (and sometimes greater) distances as 10G, with mass market deployments beginning in 2012. Today 200G systems are common in metro and long haul networks, 400G is emerging, 600G is on the horizon, and one Tbps per wavelength has already been demonstrated in the lab. The development of optical Coherent technologies has been an incredible technical achievement allowing for an enormous increase in capacity.




Figure 1 - Coherent Deployment Timeline

2.2. Modulation In Coherent Systems

Various types of modulation techniques are used in Coherent systems, depending on system requirements like reach and also existing fiber deployments which need to be upgraded to Coherent. DP-QPSK is one form of modulation used in Coherent 100G networks as this provides a robust system with long reach. In addition, the transmitter complexity is low and the DSP algorithms can be performed more simply. The sensitivity of QPSK is suitable for long-haul distances such as transoceanic links. However as the capacity increases to 200G and beyond, the common modulation format is DP - 8,16 or 64 QAM with its associated baud rates. Since Access distances are usually less than 100km, the most common modulation format planned for Coherent systems is 100G QPSK while 200G can utilize QPSK, 8-QAM, or 16-QAM formats. Figure 2 shows the various modulation formats as a function of distance.

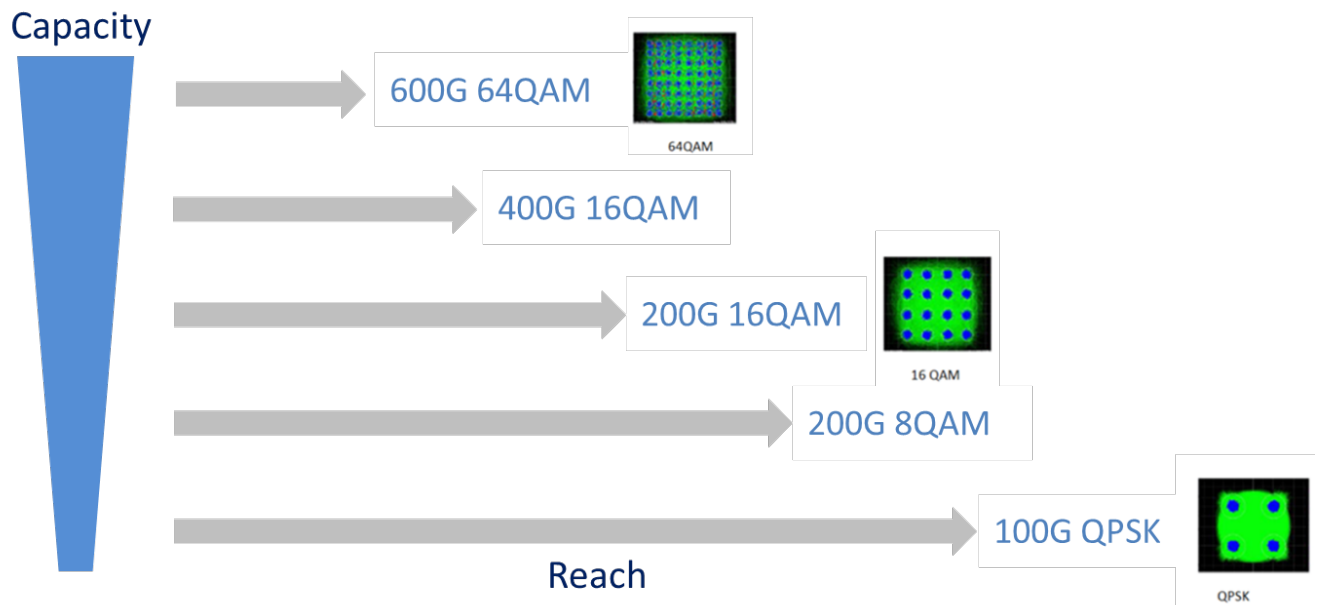


Figure 2 - Coherent Modulation and Reach

2.3. FEC in Coherent Systems

As data rates increase, chromatic dispersion increases which reduces the optical signal-to-noise ratio (OSNR) of a given system. Coherent technologies can digitally compensate for degradation caused by chromatic dispersion and polarization mode dispersion using DSPs. In addition, Coherent systems also employ forward error correction (FEC) codes that allow very significant levels of errored bits to be recovered to deliver an error free digital signal. The latest generation of FEC uses a soft-decision algorithm that significantly improves the net coding gain (NCG) of previous FEC schemes.

Figure 3 shows the basic principle of FEC in optical transport network (OTN) systems. The optical transport unit (OTU) has pre-FEC bit errors (shown by the red lines). The FEC decoder corrects certain level of error, producing a post-FEC, error-free optical data unit (ODU). Powerful electronic processing is required in the receiver to process FEC, and as the complexity of FEC algorithms has increased, the electronics have had to evolve to keep up.

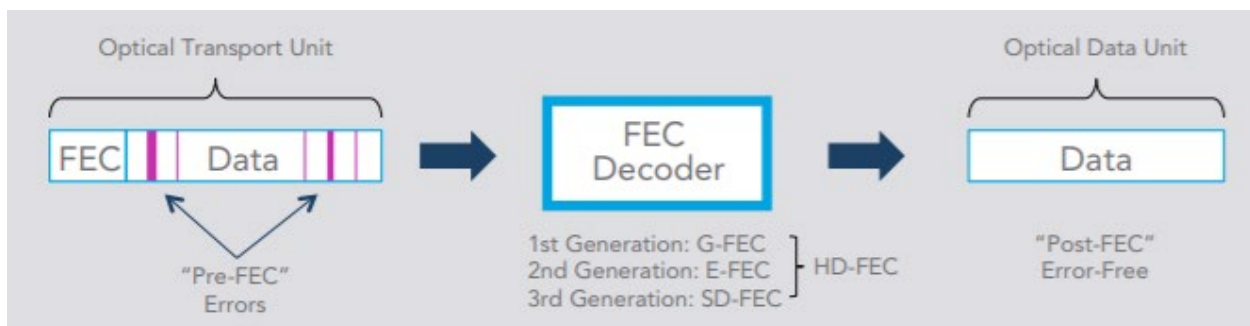


Figure 3- Pre-FEC and Post-FEC Operational Schematic

FEC evolution is generally regarded as having three distinct generations:

First generation: A generic FEC (G-FEC) which uses hard-decision decoding defined in ITU-T G.709 which allows interoperability between vendors but only delivers around 6 dB of NCG with a 6.69 percent overhead.

Second generation: An enhanced FEC (E-FEC), uses hard-decision decoding defined in ITU-T G.975.1., but no interoperability between multi-vendor E-FEC systems. It delivers between 8 dB and 9.5 dB of NCG with overhead between 6.69 and 10 percent. G-FEC and E-FEC can be used with both non-Coherent and Coherent transmission systems. Second-generation E-FEC implementations for submarine networks typically have overhead as high as 25 percent.

Third generation: Uses soft-decision decoding (SD-FEC), enabled by advances in electronic signal processing for Coherent systems at 100G and beyond. It can deliver a NCG of 11 dB or more with an overhead of 15 to 35 percent, depending on the implementation. Newer techniques to further enhance SD-FEC performance continue to be developed.

Three generations of FEC technologies have been spanned in a very short time since the introduction of Coherent systems. FEC is of key importance in achieving the necessary system margin to offer high Quality of Service (QoS) networks.

2.4. Coherent DWDM Grid and OSNR Requirements

Since operators may have existing DWDM filters deployed and wish to run Coherent through them, it is important to understand filter bandwidths and associated channel spacings required for Coherent transmission. The filter bandwidth/channel spacing requirements are determined from the modulation and its accompanying baud rate. Table 1 gives the channel spacing (grid) requirements for different Coherent modulations, as well as the OSNR requirements for typical CFP2-DCOs most likely used in Access networks. If CFP2-ACO or discrete optics are used with external (and higher power DSPs), the FEC NCG or threshold will be different and the OSNR tolerance will be different/better, assuming the same modulation schemes. The actual filter bandwidth is directly proportional to the baud rate utilized and also depends on the FEC used, in particular the % overhead employed.

Table 1 - Coherent Optical Characteristics For Typical CFP2 – DCOs

Data Rate (Gbps)	Modulation	Baud Symbol Rate (Gbaud)	Minimum Grid GHz	OSNR (dB) 10^{-15} Post FEC
100G	DP-QPSK	32	50	12
200G	DP-QPSK	64	100	16
200G	DP-8QAM	42	50	18
200G	DP-16QAM	32	50	20
400G	DP-16QAM	64	100	25
Assumes Soft decision FEC, 10.8 - 11.3 dB NCG, BOL measurements				

2.5. Coherent Equipment Type

There are various types of Coherent equipment manufactured by vendors, depending on cost and system requirements. For Access networks, transmission distances are normally less than 100km, but for long haul > 2000km transmission systems, a heavy focus on reach and performance is required.

Access applications require compact, scalable, cost-efficient, and easy-to-use Coherent transport solutions, which are met by pluggable transceiver modules such as CFP2-ACO (Analog Coherent Optics) or CFP2-DCO (Digital Coherent Optics) that include all optics and the associated digital signal processing.

The key differentiator between the two types of modules is that in the CFP2-ACO, the DSP is located outside the module with the rest of the electronics. This is a more complex solution that requires the user to interface the optical module with the DSP, but users can choose their own DSP, which makes the ACO a good fit for network equipment manufacturers who want to incorporate their own proprietary DSPs. In the CFP2-DCO the DSP is located within the module, making it a plug-and-play, simple to-deploy and operate solution for enterprise connectivity and data center interconnect (DCI).

In order to achieve the capacity and cost-per-bit targets of Access and Metro/Regional systems, it is desirable to be able to increase the data rates from 100Gb/s to 200Gb/s while maintaining sufficient signal robustness in narrower optical filtering (50GHz grid) applications. Thus it is desirable that 200Gb/s CFP2-DCO/CFP2-ACO support higher modulation formats like 8QAM and 16QAM. Higher modulations do however, have increased susceptibility to distortions, requiring higher OSNR values.

Higher performance Coherent equipment designed for regional, long haul and submarine applications, typically have discrete optical front ends with flexibility to add EDFAs, tuneable filters and high end modulators (LiNbO3 Mach Zender for example). The actual transmitted power could be much higher in discrete implementation (up to 5dB higher launch power) than typical CFP2-ACO/DCOs. Similarly, Rx sensitivity can be a lot lower.

As described in the previous sections, Coherent system performance and reach also depends on many factors such as FEC and modulation schemes employed in addition to the optical front end used. To summarize, Coherent system performance depends on:

- DSP FEC and Gain (NCG) with Soft decision Forward Error Correction.
- Optical front end: Transmit output power, fixed or adjustable (built in EDFA), type of modulator, e.g (LiNbO3 MZM), laser linewidth
- Receiver sensitivity

In this paper we will present proof of concept test results of bi-directional, mixed 10G and Coherent 100G/200G using two very different types of equipment. We will designate them as equipment A (used in long haul/submarine deployments) and equipment B (Metro/Access platform with pluggable CFP2-ACOs designed for enterprise and data center interconnect (DCI) connectivity).

2.6. Coherent Bi-directional Transmission Capability

An important requirement for Cox networks is the ability to have bi-directional transmission, where both downstream and upstream signals are transported over the same fiber in opposite directions. This is important for applications where fiber resources are limited or where there is a requirement for fiber

redundancy. Other operators may have “unidirectional” transmission (uni) that utilize two fibers, one for each direction.

Figure 4 below shows a high-level block diagram of the uni and bi-directional transmission Coherent equipment. The receiver uses a reference light signal (local oscillator or LO) as a comparison to measure the phase and amplitude of the incoming light wave. In unidirectional transmission, the same laser is used for both the transmitted wavelength and as an LO for the receive signal, so the downstream and upstream must be the same wavelength. This is the normal case where two separate fibers are utilized for the downstream and upstream directions. In bi-directional Coherent transmission, the optical front end contains two lasers so the transmitted and receive signals can be at different wavelengths.

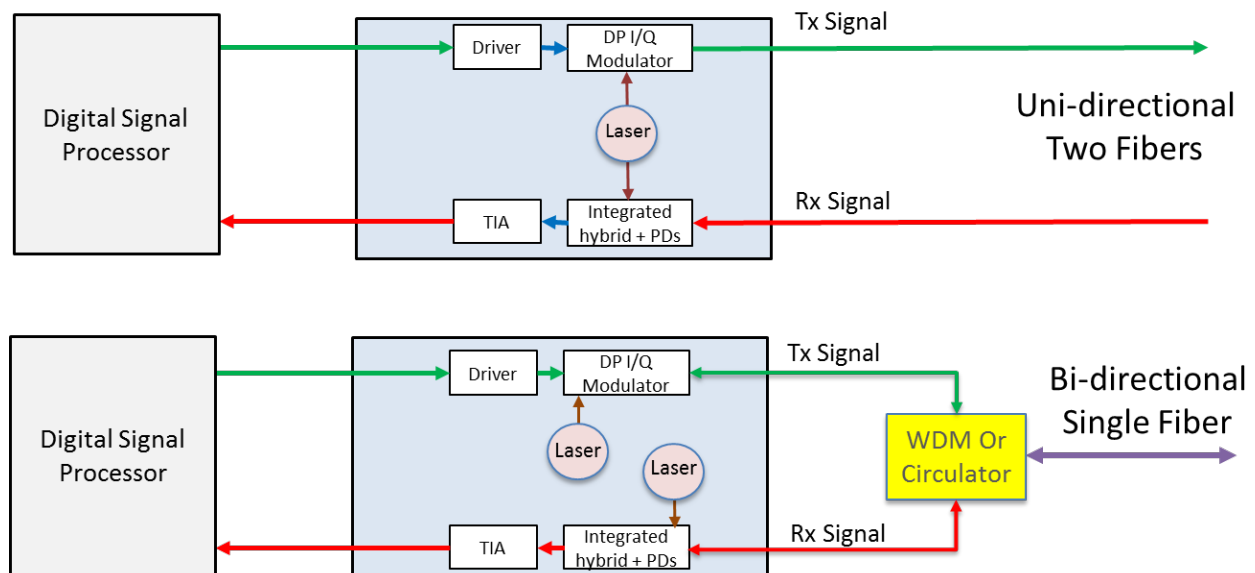


Figure 4– Uni and Bi-directional Coherent Transceivers

3. OCML – MDM Network

Cox Communications Access networks are unique in the sense that most of its primary HFC nodes are protected by a secondary back-up fiber, providing optical path redundancy which significantly increases reliability. Additionally, we have bi-directional traffic flow on the same fiber. To maintain this architecture in the transition to DAA networks, we had to solve the problem of how to effectively transport optical signals over these bi-directional dual fiber rings. We had the additional requirement to carry GPON/10GEPON plus 10G DWDM Remote-PHY signals. It was also desirable to be able to carry Coherent 100G/200G over the same network.

To meet these exacting requirements, we developed the Optical Communications Module Link Extender (OCML)¹ and the MDM. The OCML supports next-generation fiber deep DWDM Access networks and may transport up to 20 X 10G bi-directional wavelengths, plus future Coherent 100G/200G over variable 5 to 60 km path redundant fiber links. PON/10GEPON signals may also be transported in the platform through an innovative WDM filter mechanism which passes through all wavelengths and blocks the 10G C band. As the industry begins to implement DAA rollouts, the OCML allows use of available 10G NRZ

DWDM technology, but the underlying infrastructure is based on ITU standard wavelength plans in anticipation of the requirement to deploy Coherent 100G (and beyond) wavelengths. The MDM is a field-based passive Mux/DeMux filter incorporating a 3 dB splitter for connection to the primary and backup secondary fiber. The basic OCML – MDM network shown in figure 5 below allows active components of a DWDM Access network to be integrated into one module which can then be placed in the central office/headend/hub, while the outside plant aggregation point can be a simple Mux/Demux (MDM).

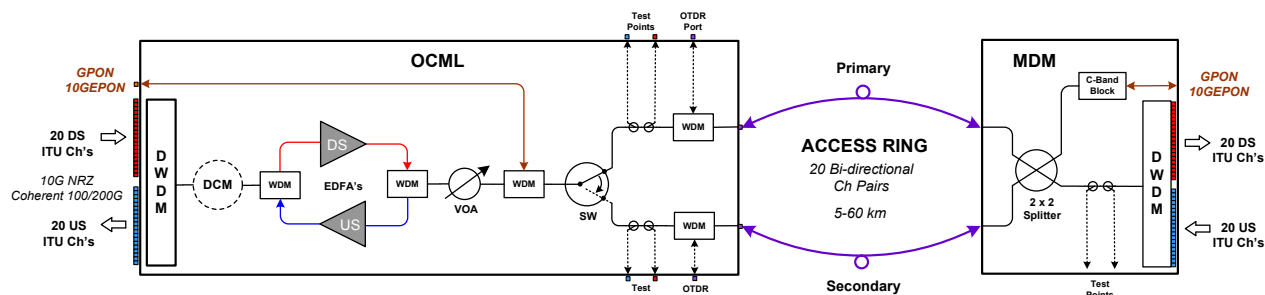


Figure 5- OCML - MDM Network

A high level remote-PHY architecture showing the OCML - MDM is shown in Figure 6 below. This basically shows how a 10G DWDM optical trunk can be used to effectively feed a large number of Remote-PHY nodes. While ten RPD nodes are shown, the OCML – MDM can support up to twenty RPD's.

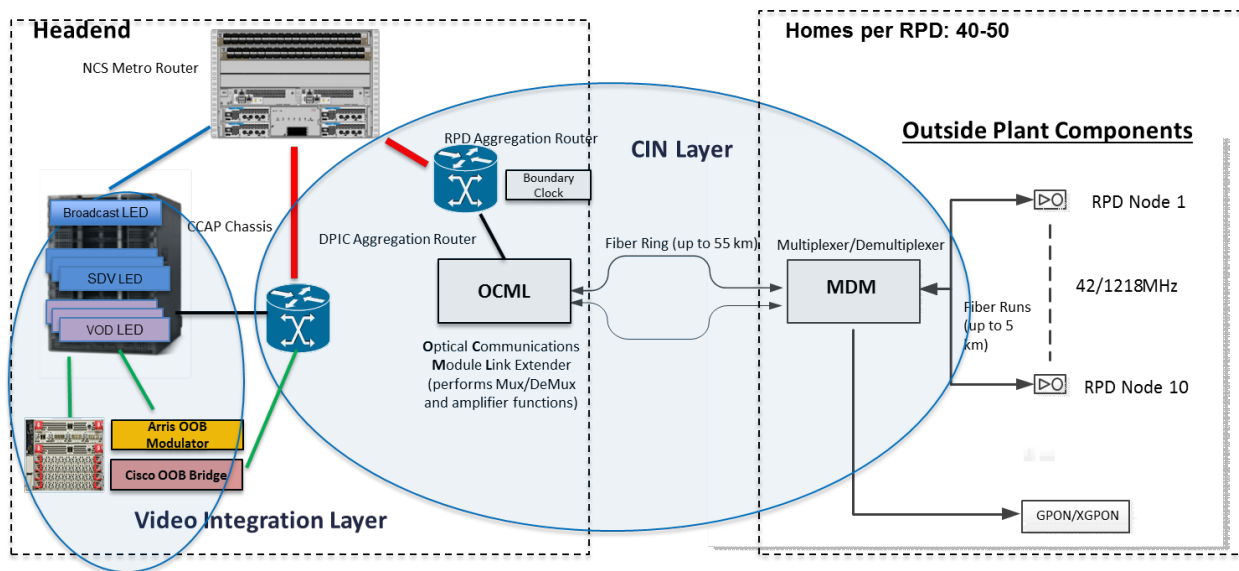


Figure 6 - Remote-PHY High Level Solutions Architecture

3.1. 10G and Coherent Coexistence

One of the key advantages which DAA networks provide is the ability to transport various types of signals across MSO Access networks. Up to now, MSO HFC networks have utilized very restrictive analog signals which normally have high power levels. The analog optical signals typically required

elaborate wavelength plans to mitigate fiber non-linearities such as Four wave Mixing (FWM). Cox is planning to roll out Remote-PHY networks utilizing a 10G DWDM bi-directional network via the OCML-MDM. Our challenge was to utilize this same network to also transport Coherent 100G and 200G signals plus GPON and 10GEAPON. In addition, the network had to carry bi-directional wavelengths over a primary and backup secondary fiber. We will discuss some of the technical considerations to be taken into account in deploying mixed 10G and Coherent optical signals over fiber networks.

3.2. Impact of 10G and DCMs on Coherent 100G/200G

Existing 10G networks pose two performance challenges to Coherent 100G/200G. The first is the 10G DWDM signals themselves, while the second is due to the dispersion compensation modules (DCMs) typically found in 10G networks. Since one of the OCML variants utilizes a DCM, we needed to ensure that Coherent signals can be transported through the OCML, both with and without dispersion compensation.

3.2.1. Cross Phase Modulation (XPM) In Mixed 10G/Coherent Signals

10G systems use amplitude (or power) based On-Off Keying (OOK) modulation, while Coherent 100G/200G transmission most commonly uses DP-QPSK or DP-8/16QAM modulation formats. These modulation schemes alter the phase of the transmitted signal in addition to the amplitude. 10G signals have a much higher power spectral density (large amount of power in a very small spectral range) than Coherent 100G/200G, as shown in figure 7 below. This has a greater impact on the refractive index than Coherent signals. Since Coherent signals make use of phase modulation, they are more severely impacted by effects that alter the signals phase. For these reasons, cross phase modulation (XPM) from 10G wavelengths can have a significant impact on the reach of Coherent 100G/200 wavelengths in a mixed 10G/100G network. Fortunately, since Access networks are typically less than 80km, XPM impacts induced by 10G signals should not be too severe, but still ought to be evaluated in mixed 10G/Coherent 100G or 200G networks.

10G DWDM receivers based on OOK can typically tolerate 80km~100km of chromatic dispersion. For this reason, dispersion compensation modules (DCMs) are typically deployed in 10G networks. Since the performance of 10G DWDM systems with EDFAs depend on various factors such as OSNR, fiber dispersion and optical receive power, DCMs are also used to help offset the effects of low OSNR, as can be the case in Cox OCML - MDM networks. However, low chromatic dispersion is actually a disadvantage for Coherent transmission. In the absence of chromatic dispersion, the symbols of each channel could all alter the refractive index of the fiber at the same time, thus maximizing the impact on the fiber's refractive index and thereby increasing XPM. By causing the channels to travel at slightly different speeds, chromatic dispersion reduces the time correlation between the symbols, thus reducing the buildup of nonlinear penalties including XPM and Self Phase modulation (SPM). Given the above, we investigated the impact of 10G signals and DCMs on Coherent 100G/200G signals through the OCML – MDM network.




Figure 7 - Signal Power Density 10G and 100G

3.2.2. Guard Bands In Mixed 10G and Coherent 100G/200G

Guard bands (not using one or more channels between the 100G and the 10G channels) can be employed to mitigate XPM in mixed 10G and Coherent signal networks as shown in Figure 8. For example, if a 100G wavelength is deployed on channel 25, then channel 24 and channel 26 are left empty to accommodate 10G wavelengths on channels 23 and 27. Alternatively, 10G channels can use one end of the C- band and Coherent signals deployed from the other end. However, guard bands cause reduced spectral efficiency and planning challenges since they do not use certain wavelengths. Guard bands are typically deployed on very long (>1000km) links where 10G and Coherent coexist. We evaluated whether guard bands would be needed in short (< 80km) links for mixed 10G and Coherent signals.




Figure 8 - Guard Band Example

4. Testing of Coherent and 10G Coexistence in OCML

The performance of Coherent systems depends on various factors as explained in section 2.0.

In this section we present proof of concept test results for mixed bidirectional 10G and Coherent 100G/200G through the OCML using two types of equipment.

Equipment A: High performance Coherent equipment designed for regional, long haul and submarine applications. Typically this type of equipment has a discrete optical front-end with added flexibility to add EDFAs, tuneable filters and high end modulators (LiNbO3 Mach Zender for example). This allows for links to be optimized for performance and reach.

Equipment B: Lower cost and lower performance Coherent equipment designed for shorter, Metro/Access applications. Typically this type of equipment utilizes off-the-shelf 200G 16QAM pluggable CFP2-ACO optics. Since this is most likely the type of equipment which cable operators will use for Access, we wanted to evaluate its performance against the high performance long haul equipment.

In addition both the Coherent platforms we tested had bi-directional capability in that they could transmit and receive at different wavelengths to allow transmission over a common fiber. This was a major requirement for the Cox Access network.

4.1. Test Set-Up

Figure 9 below shows the test set-up used to evaluate the performance of a network supporting bi-directional multi-channel 10G and Coherent 100G/200G through the OCML – MDM network over 40 and 60 kilometers of fiber. The 60km setup used a version of OCML that has an integrated Dispersion Compensation Module (DCM). The 40km setup used a lower cost version of OCML without DCM. We also used two types of Coherent equipment (A and B) as described above. The OCML – MDM equipment supports twenty bi-directional DWDM channel pairs. Seventeen were used to carry 10G NRZ signals and the remaining three transported Coherent 100G/200G.

The high-capacity Ethernet traffic generators shown in the diagram were used to compare transmitted packets to received packets to determine whether any uncorrected data errors occurred on any of the bi-directional channel pairs during the testing. In the case of the Coherent optical signals, any errors detected by the traffic generator/packet analyzer indicated the presence of uncorrectable (post-FEC) errors. There was no FEC in the 10G NRZ SFP+ optics, nor the host switches they were installed in.

Two different channel loadings were used in the 60km testing (1 and 2). Only channel loading 1 was used in the 40km testing. Figure 10 shows the ITU channels, associated equipment types, modulation types, and data ratings used for both channel loadings. Note that no guard bands were used to separate the 10G and the 100/200G Coherent signals.

Coherent & 10G Bidirectional Coexistence COX LAB TEST SET UP

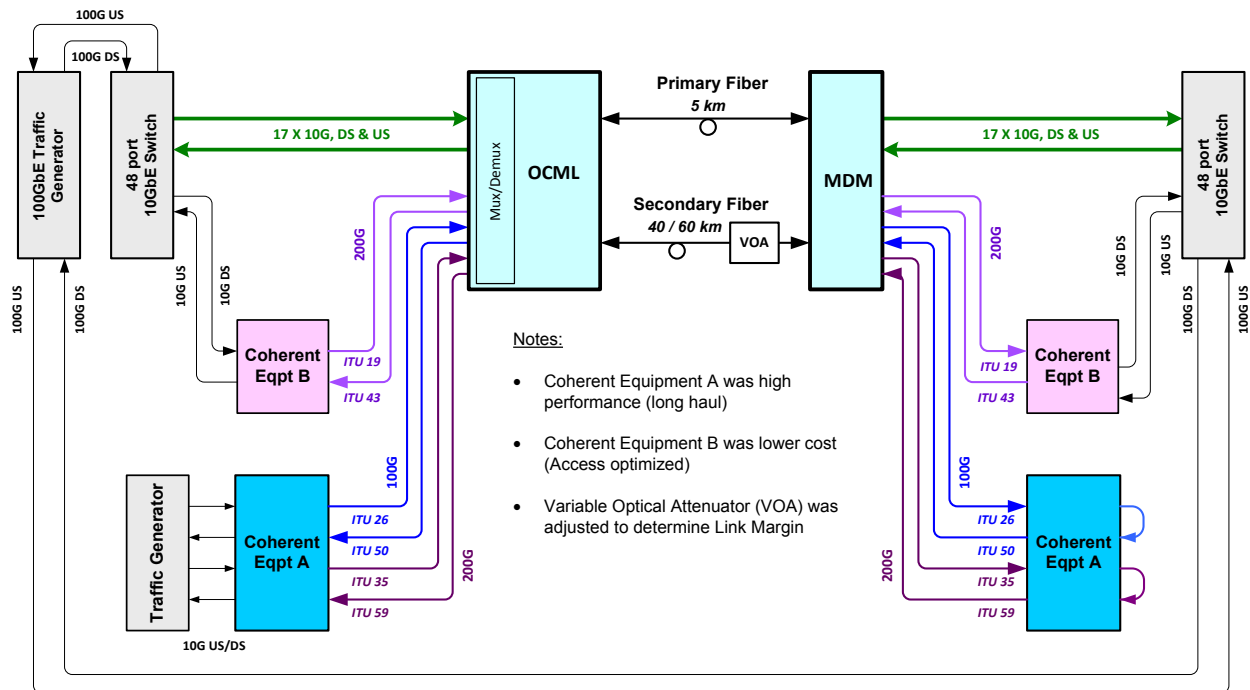


Figure 9 - Test Set up for 10G and Coherent 100G/200G Coexistence

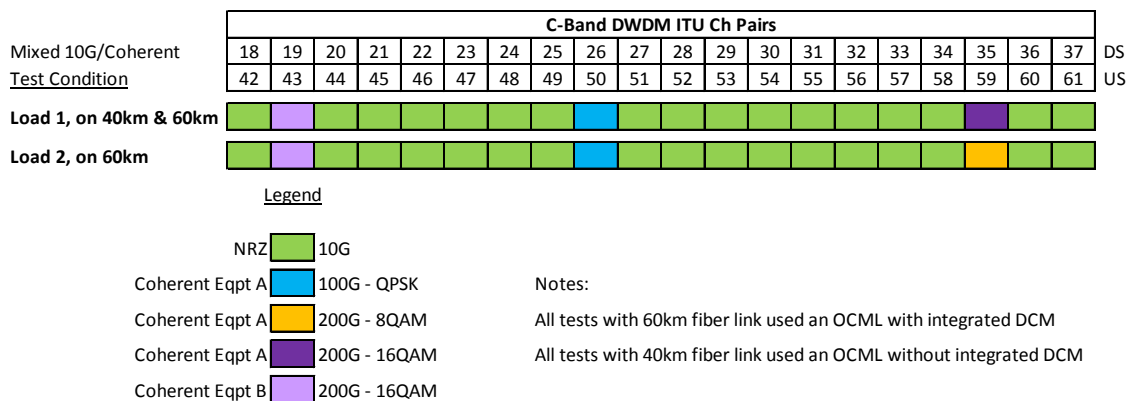


Figure 10 - Test Wavelengths (ITU Ch Pairs) and Optical Signal Types

4.2. Tests Performed

We first performed baseline tests with only the seventeen 10G signals present on the fiber link, and then with only the three 100G/200G Coherent signals present. With the 60km and the 40km fiber link and an additional 2.7 dB of optical attenuation inserted in the link, no uncorrectable errors occurred on any of the

signal paths. Next, we combined the 10G and 100/200G Coherent signals on the fiber with the same link condition and again had no uncorrectable errors on any of the signal paths with either of the channel loadings.


After we established error free performance in the initial coexistence tests, we performed what we called Link Margin tests. In these tests we adjusted the variable optical attenuator (VOA) in the 60km and the 40km fiber link to determine the maximum amount of attenuation that could be inserted prior to the onset of uncorrectable errors for each equipment and modulation type. To do so, we first increased the optical attenuation in the link until we found the onset of uncorrectable errors. We then reduced the optical attenuation slowly in 0.5 dB increments to determine the maximum amount of attenuation that could be added without inducing errors. The Link Margin in dB was recorded in the test results for the various types of optical signals used in the test.

Lastly, we turned the 10G optical signals adjacent to the 100G/200G Coherent optical signals off to determine if it made any difference to Link Margin - which might occur if the Coherent signals were negatively impacted by Cross Phase Modulation (XPM).

4.3. Test Results

The Link Margin Test Results are presented in figure 11 below. Note that there were seventeen 10G NRZ channel pairs used in the test. We monitored all seventeen and recorded the worst case Link Margin in the results. The OCML's used in the testing make use of both downstream and upstream optical amplification, but the optical input power to the upstream EDFA gets lower on longer optical links, which in turn causes the upstream OSNR to be lower than the downstream OSNR. As expected, when we performed the Link Margin tests the upstream was the first to start taking errors, due to the combination of lower OSNR and low optical receive power.

Additionally, when we turned the 10G signals adjacent to the 100G/200G Coherent signals off and on, we found no difference to the measured Link Margin for the Coherent signals, indicating that there was no measurable XPM impact on the Coherent signals.



*Link Margin = Maximum optical attenuation that could be added to the 40/60km link before onset of uncorrectable errors

Figure 11 - Link Margin Test Results

4.4. Test Conclusions

We proved that 10G NRZ and 100G/200G Coherent signals can coexist and perform well across 40 and 60km bi-directional optical links using the OCML/MDM. As expected, the amount of headroom (link margin) available depended on the type of Coherent equipment used and the modulation type. While the margin for the lower cost Equipment B Coherent type that we expect operators to use in Access applications was somewhat lower than with the traditional 10G NRZ pluggable optics, it was the 200G equipment with higher order 16QAM modulation which had the highest OSNR requirements of the types tested. While we did not have any 200G 8QAM or 100G QPSK equipment available for testing in the lower cost Equipment B type, we anticipate such equipment will provide additional margin of roughly 2-7 dB, depending on type.

Link Margin (headroom) can be thought of in two ways. One is that each dB of margin might allow an additional dB of optical reach. The other is that each dB of margin affords additional protection from service interruption in the event of inadvertent conditions in the plant that create additional optical loss. When determining optical design guidelines it is best to consider both aspects and first apply an agreed upon minimum amount of margin to absorb undesired optical losses and account for equipment performance variation, and then consider the rest of the margin available for additional optical reach if

needed. We recommend working with your selected equipment vendor on any designs incorporating Coherent equipment.

5. Coherent Access Applications

Long-haul and metro network systems have been utilizing the tremendous capabilities of optical Coherent transport for many years. Coherent transport is now being investigated for Access optical networks. In this section we will present some of the applications which can be supported by Coherent technologies.

5.1. Converged Cable Access Network

Figure 12 shows how several optical technologies may coexist on the same fiber using the Optical Communications Module Link Extender (OCML) and MDM network. This network supports several optical technologies and applications, resulting in a powerful, scalable and technology agnostic solution. It can provide 10G for cable RPDs (Remote-PHY Devices) and RMDs (Remote-MAC Devices), as well as 100G/200G Coherent for high capacity applications such as businesses, hotels, hospitals and MDUs. A Coherent optical trunk can also support Remote-PHY nodes but would need an outside plant aggregation node device to provide dedicated 10G fiber links to the RPD nodes.

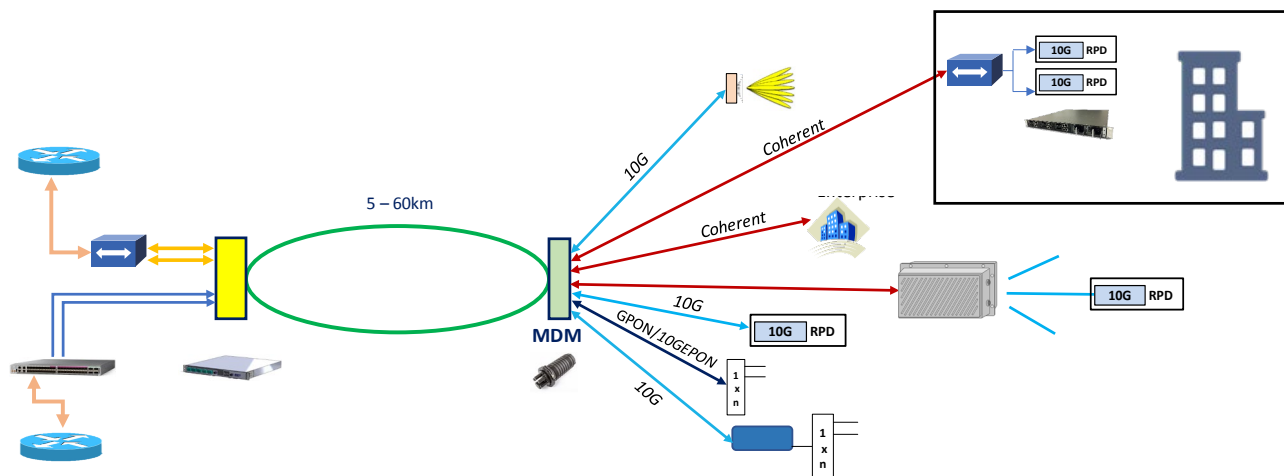


Figure 12 - Converged Cable Access Network

5.2. Remote-PHY Coherent Optical Trunk

Figure 13 shows how a Remote-PHY architecture utilizing a Coherent link can be realized with a high capacity Coherent optical trunk. Essentially a 100G/200G Coherent link would feed an outside plant, hardened aggregating device to provide up to 10G data links to individual Remote-PHY Nodes. These 10G links can be lower cost grey (1310nm) optics, with each RPD supported by a dedicated fiber pair for downstream and upstream.

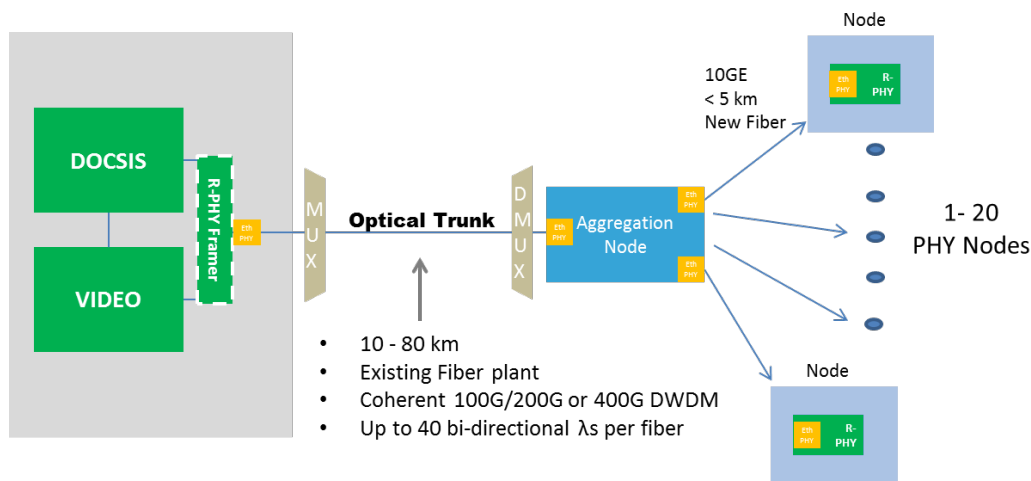


Figure 13 - Remote-PHY Physical Distribution

This hardened (I – Temp) aggregation node can be a Muxponder, Ethernet Switch or Router, depending on operator architectures and preferences. The main features of the various types of aggregation devices are given below.

5.2.1. Muxponder

A Muxponder operates at Layer 1. It is a simple, low-cost device which provides dedicated bandwidth or capacity with 100% mapping, essentially the same inputs as outputs. A Muxponder does not support statistical multiplexing nor over-subscription. The total backhaul capacity is shared equally between the output ports, so a 200Gbps optical link can provide 20 x 10G ports, each with dedicated 10G to RPDs. The muxponder is the easiest and lowest cost option to support DAA architectures.

5.2.2. Ethernet Switch

An Ethernet switch operates at layer 2 and allows both statistical multiplexing and oversubscription. The inputs and outputs need not match, so a 100G optical link can provide 20 x 5G ports to provide 5G links to RPDs. An Ethernet switch can technically operate with less backhaul capacity than a muxponder. An Ethernet switch is a relatively simple device, but more complex than a muxponder. It will also cost more and have a higher power consumption than a Muxponder. It requires a backplane with switching capacity to handle all inputs and outputs. An Ethernet switch is perhaps a good compromise between a muxponder and a full-fledged router as described in section 5.2.3 below.

5.2.3. Router

A Router operates at Layer 3 and can do all the same things as an Ethernet switch. It natively supports multicast replication and for Remote-PHY permits lower backhaul capacity requirements until the service group to RPD ratio becomes 1:1. It is moderately more complex than an Ethernet switch and likely costs more and consumes more power. From the various aggregation node types described in this section, the router affords the most flexibility, but comes with a higher degree of complexity, power consumption and cost.

5.3. RPhy 10G DWDM Optical Trunk Transition to Coherent

Figure 14 shows a typical spine and leaf network where the CCAP aggregation switch is connected to the OCML with multiple 10G signals. The OCML is connected to an MDM through a primary and backup secondary fiber. The MDM provides 10G dedicated links to Remote-PHY devices. In transitioning to a Coherent optical trunk, the Spine Aggregation router could be directly connected to a Coherent transport system and input to the OCML. The “leaf” CCAP Agg switch could then be moved to an outside plant location, but would need to be housed in a temperature hardened device (aggregation node). This type of aggregation node would also need to contain the necessary hardened Coherent transceivers plus the switching fabric to provide lower capacity 5G to 10G connections to Remote-PHY Devices.

We could replace a 10G DWDM ring with a Coherent 100G or 200G optical trunk. This would reduce the number of 10G ports at the headend but we would then need a more complex, active hardened device in the outside plant. The advantage of a Coherent optical trunk is that we would only need a single Coherent wavelength to feed all the RPD devices in a typical node serving area, thus freeing other wavelengths to support higher bandwidth applications like business services, hospitals, schools, etc.

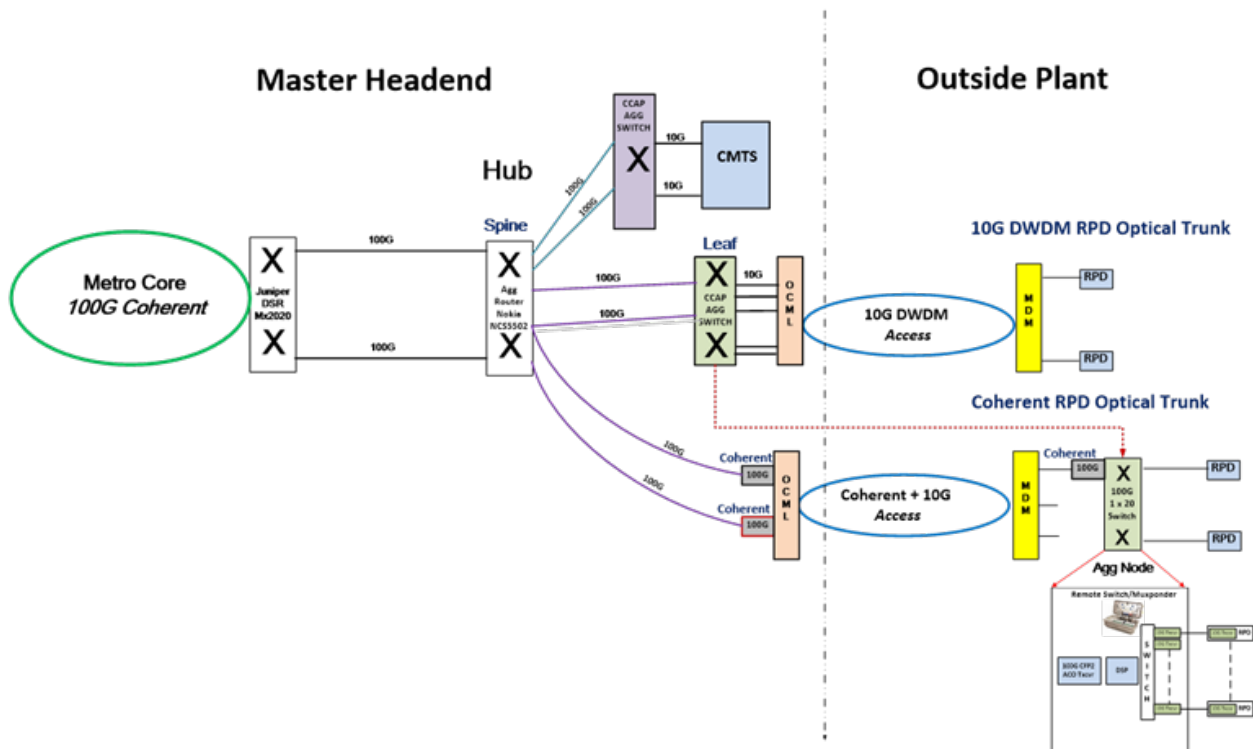


Figure 14 - 10G DWDM Optical Trunk Transition to Coherent

5.4. Coherent Business Service Applications

Figure 15 shows a Coherent 100G/200G DWDM Access network which can support business services and other high capacity applications such as Airports, Conference centers, hospitals, etc. The OCML would be located at the headend while the MDM would be physically located near an existing HFC node

and provide multiple Coherent 100G or 200G outputs. Applications requiring high capacity bandwidths can be provided with a dedicated Coherent 100G/200G link.

Figure 14 also shows how a Coherent 100G link could be provided to existing cabinets where muxponders, switches or routers could be used to provide lower capacity such as 1G and 10G services. This architecture fully utilizes the capacity of Coherent technology to provide ever increasing bandwidth to customers.

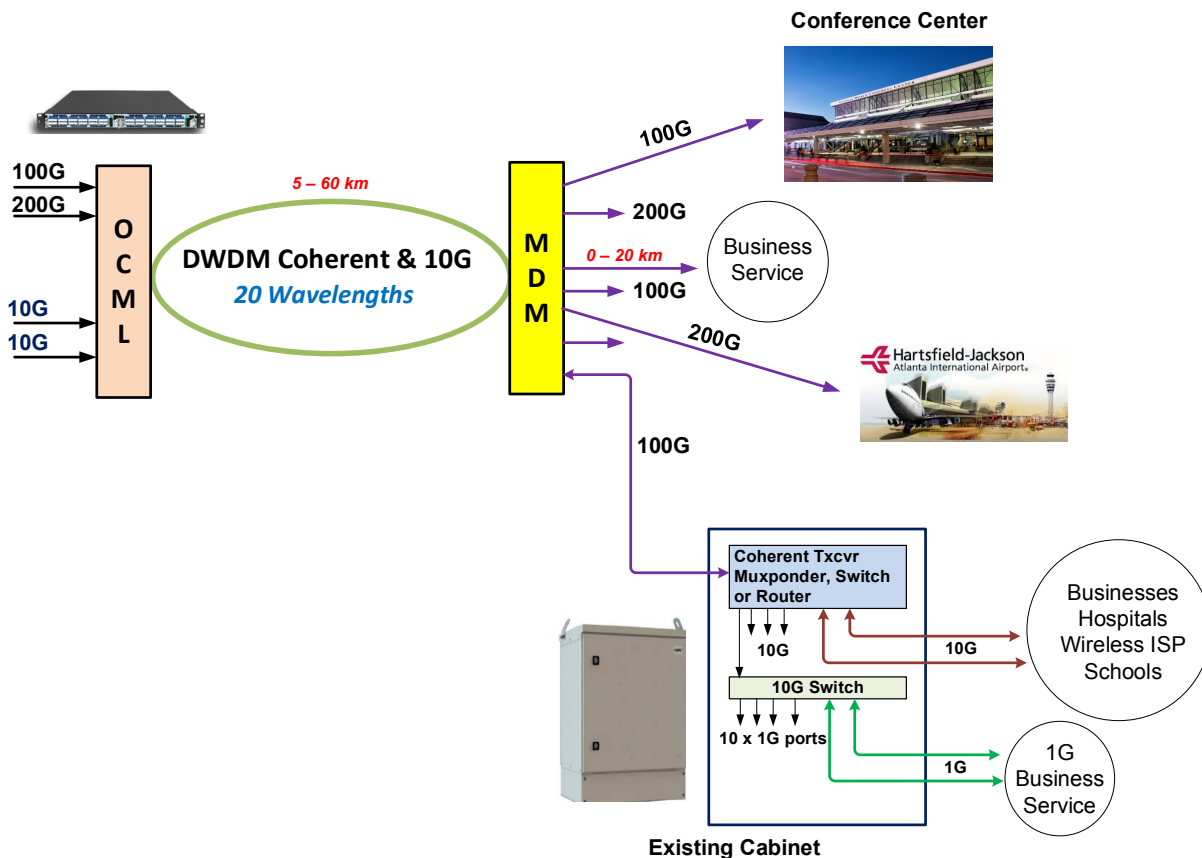


Figure 15 - Coherent Business Services

5.5. Coherent Hub Consolidation

With its high capacity DWDM capability, Coherent transport could also be used for Hub consolidation or collapsing the hubs to smaller sizes. Figure 16 shows a typical metro MSO fiber optical ring where several hub sites are connected. If we have multiple 10G DWDM links feeding Remote-PHY devices at hub A, these could be aggregated via muxponders onto a 200G DWDM ring and transported back to the master headend. This would considerably reduce equipment required at Hub A which could be collapsed into a much smaller footprint, potentially in a hut or hardened cabinet.

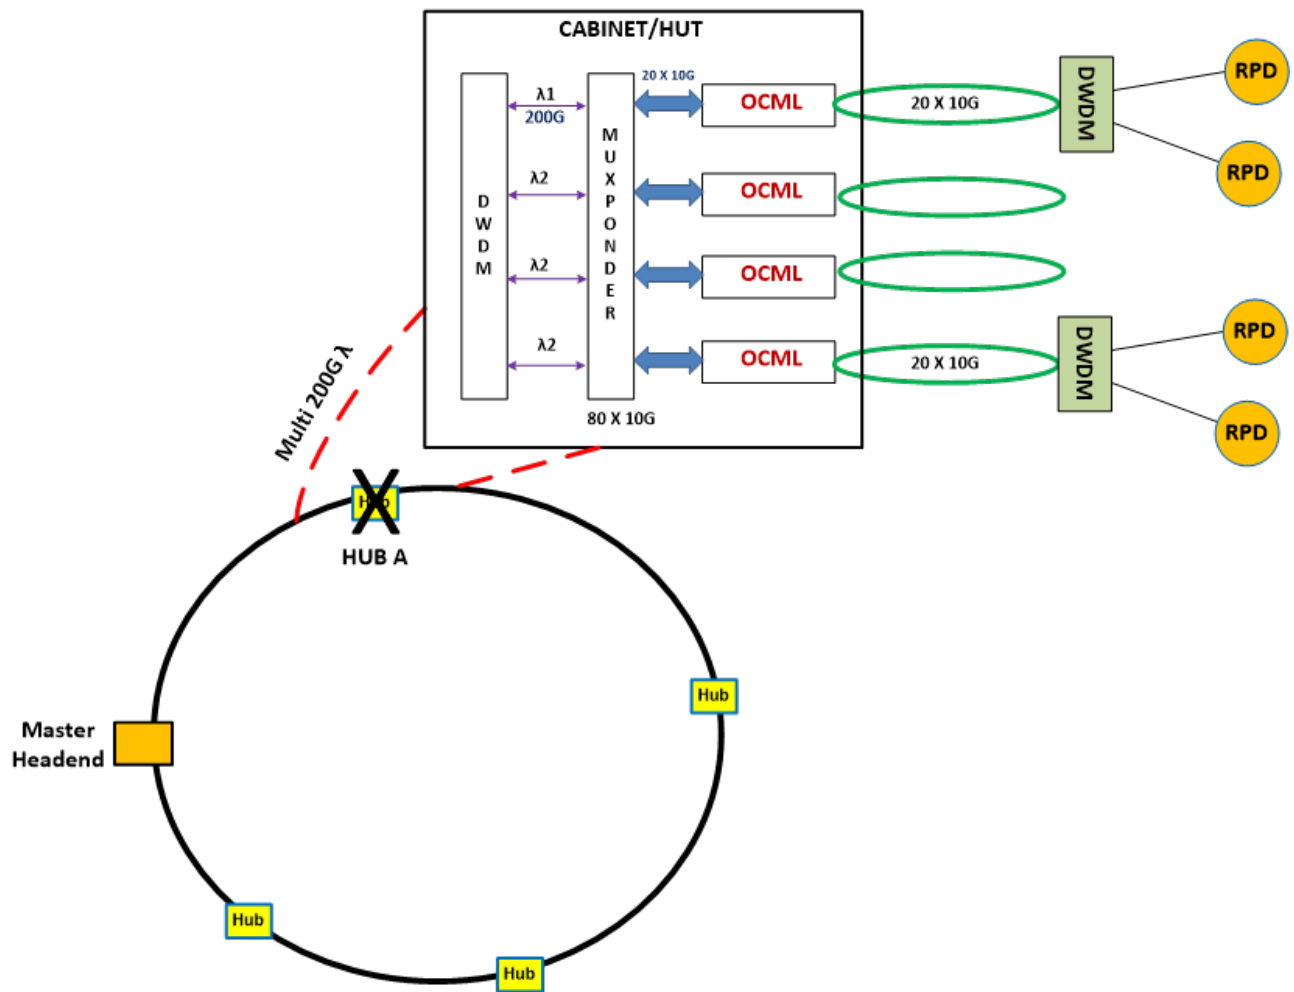


Figure 16 - Coherent Hub Consolidation

6. Conclusion

A high-performance Access network is key to the commercial success of delivering broadband services to the customer. In addition, Access networks should be scalable and technology agnostic. 10G NRZ and 100G/200 Coherent coexistence on a single fiber through an integrated platform (OCML) has been successfully demonstrated over 60km with dispersion compensation and 40km without dispersion compensation. We tested two types of Coherent equipment, a high performance, long-haul (A) and a lower cost, lower performance Access/Edge Coherent device (B). We established that both Coherent 100G QPSK and 200G 8 QAM and 16 QAM could easily be transported through the OCML – MDM infrastructure, and as expected, the Coherent 100G had greater reach (link margin) than the Coherent 200G. Several Access applications using Coherent 100G/200G have also been presented. In conclusion, with data rate requirements increasing every year, the huge capacity capability of Coherent can be used effectively to supplement or transplant 10G OOK signals in the Access region of MSO networks.

Abbreviations

OCML	Optical Communications Module Link Extender
MDM	Mux DeMux
MTC	Master Terminal Center
STC	Secondary Terminal Center
10G	10Gbps
DSP	Digital Signal Processing
LO	Local Oscillator
Bps	bits per second
FEC	Forward error correction
HFC	Hybrid fiber-coax
SCTE	Society of Cable Telecommunications Engineers
OIF	The Optical Internetworking Forum
DCM	Dispersion Compensation Module
NRZ	Non-Return-to-Zero
DWDM	Dense Wavelength Division Multiplexing
RPD	Remote-PHY Device
FTTH	Fiber to the Home
PON	Passive Optical Network
GPON	Gigabit-capable Passive Optical Network
PIN	PIN diode has a wide, undoped intrinsic semiconductor region between a p-type semiconductor and an n-type semiconductor region.
APD	Avalanche photo diode
OSNR	Optical to Signal Noise Ratio
OOK	On-Off keying
NCG	Net Coding Gain
DCI	Data Center Interconnect (DCI)
BOL	Beginning of Life

Bibliography & References

1. DWDM Access For Remote-PHY Networks Integrated Optical Communications Module (OCML), Harj Ghuman SCTE 2017.
2. Optimizing the Performance of Coherent 100G in 10G Dispersion-Managed Networks,
3. Coriant White Paper.
4. Coherent WDM technologies, Infinera White Paper.
5. Soft-decision Forward Error Correction for Coherent Super-channels, Infinera White Paper.
6. Backhaul Capacity for Optics in Remote-PHY, IEEE 802.3 Interim March presentation, Fernando Villarruel, Harj Ghuman, Michael Eggert, Marek Hajduczenia.
7. FEC in Optical Communicatios, A.Tychopoulos, O. Koufopavlou, I. tomkos, IEEE Circuits & Devices, 2006.
8. Evaluating 200G/400G solutions for practical deployments in long-haul network, Electronic Letters, Sep. 2016, Y. Ma, S. Makovejs, Q.Wang, W.Wood, N. Kaliteevskiy, J. Li, C. Zhang
9. Real-time transmission of 16 Tb/s over 1020km using 200Gb/s CFP2-DCO, H. zhang, B. Zhu, S. Park, C. Doerr, M. Aydinlik, J. Geyer, T. Pfau, G. Pendock, R. Aroca, F. Liu, C. Rasmussen, B. Mikkelsen, P. I. Borel, T. Geisler, R. Jensen, D. W. Peckham, R. Lingle , D. Vaidya,. F. Yan, P. Wisk, D. Digiovanni, Optics Express, 2018

Acknowledgements

Daniel Cleere, Cox Communications; Mark Campbell, Maurice Howard, Nazar Neayem, Russell Pretty: Nokia

Comparison Of LPWA Technologies And Realizable Use Cases

A Technical Paper Prepared for SCTE•ISBE by

Satish Chalapati

Telecom Billing and International Roaming Consultant
Tata Consultancy Services
2314 Larchmont Pl, Mt Laurel, NJ, USA, 08054
+1 9729832013
satish.chalapati@tcs.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	4
1. Business Forecast.....	4
2. Architecture	4
3. Standards/Technologies	5
3.1. LoRaWAN	5
3.2. LTE-Cat0/LTE-Cat1/LTE-Cat M1 (eMTC).....	5
3.3. NB-IoT	5
3.4. EC-GSM-IoT.....	6
3.5. Sigfox	6
4. Technology Comparision	6
5. Three dimensional view of Parameters vs Technology	6
6. Radio Features.....	7
6.1. Bandwidth.....	7
6.2. Spectrum	7
6.3. Frequency Bands	7
6.4. Standardization	7
6.5. Uplink & Downlink	7
6.6. Latency.....	8
6.7. Coupling Loss.....	8
7. Non Radio Features	8
7.1. Power Saving Mode	8
7.2. Extended Idle Discontinuous reception Mode.....	8
7.3. Cloud Compatibility	8
7.4. Artificial Intelligence & Machine Learning	8
7.5. Analytics and Big Data	8
7.6. Low Cost	9
7.7. Security Considerations	9
8. Application Use Cases	9
8.1. Agriculture	9
8.2. Smart Wearables	9
8.3. City Management and Metering	9
8.4. Vending Machines.....	9
8.5. Industries.....	10
8.6. Environmental Monitoring	10
8.7. Car Auto-Piloting – Original Idea.....	10
9. NB-IoT Roaming.....	10
10. NB-IoT 5G	10
Conclusion.....	10
Abbreviations	11
Bibliography & References.....	12

List of Figures

Title	Page Number
Figure 1 – Architecture Diagram showing multiple layers.....	5
Figure 2 – 3D view of key parameters vs technology	7

List of Tables

Title	Page Number
Table 1 – Comparision of Technologies	6

Introduction

The evolution of the Internet has opened a wide variety of opportunities for connecting devices remotely. This has resulted in growing demand of remote maintenance and connectivity of devices and has made the Internet of Things evolve at a faster pace than anticipated. The advancement of wireless technologies and data transfer capacities over the same timeframe has resulted in wider usage of internet over wireless. The combination of all of the above factors has led to the development of IoT over wireless communications.

This paper will cover Low Power Wide Area Network (LPWA) technologies which have overcome the challenges of IoT over wireless by reducing power consumption, increasing coverage, tailoring the bandwidth according to the needs, and many other advantages.

There are various LPWA standards developed by GSMA, LoRa(SemTech), Sigfox, NB-Fi(WAVIoT), RPMA(Ingenu) leveraging 2G/3G/4G/5G bands and Industrial-Scientific-Medical (ISM) bands. Currently existing standards are LTE Cat 1, LTE Cat 0, LTE Cat M1 aka LTE-M(eMTC), NB-IoT aka LTE Cat NB1(Narrow Band), EC-GSM-IOT(EDGE), LoRaWAN, Sigfox, NB-Fi. Each of these technologies/standards are discussed and compared in detail in the sections of this paper.

Below are some of the features that are gained with the advancement of above technologies.

- Low power consumption resulting in battery life >10 years.
- Leveraging existing bands and future compatibility.
- Lower device costs of around \$5 and lower maintenance costs.
- Improved connectivity – indoor and long range.
- Tailored bandwidth for lower data rates, latency, and mode.
- Network scalability with ease of capacity upgrade.
- Roaming connectivity.
- Security and authorization.

Content

1. Business Forecast

As per market research issued in July 2018, forecast for the Low Power Wide Area (LPWA) global business market is expected to grow at a compound annual growth rate (CAGR) of approximately 72% during the forecast period from 2017-2023.

According to an LPWA forecast from ABI Research, network connections will grow at a 53% CAGR through 2023.

2. Architecture

The below architecture diagram shows different layers of LPWA technology. Each of the features per respective layer is explained in sections 6 and 7

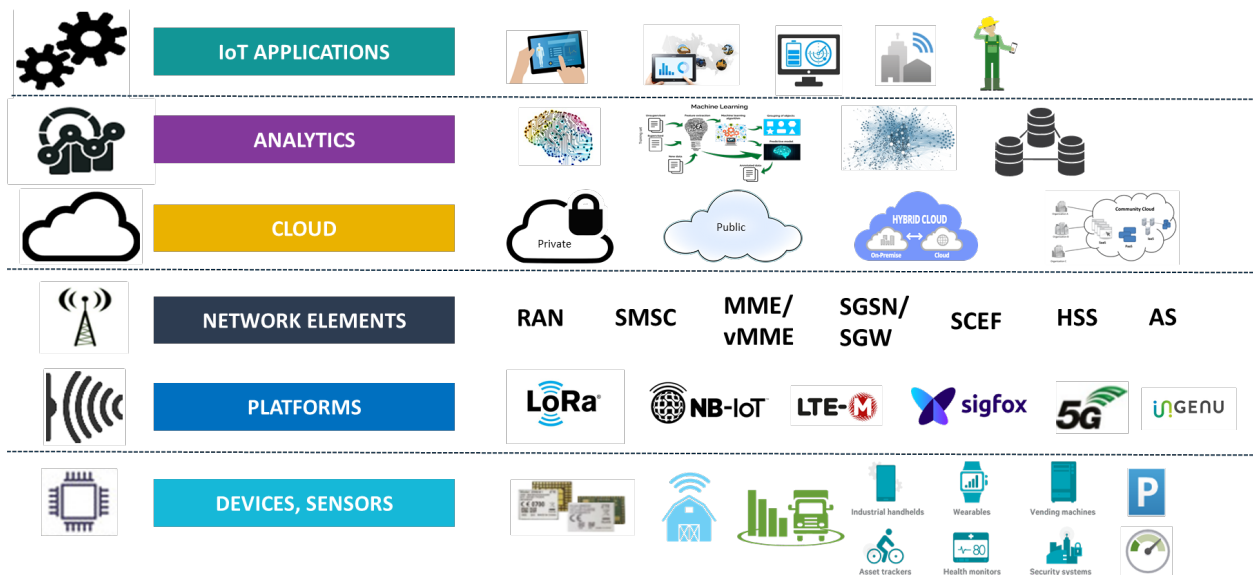


Figure 1 – Architecture Diagram showing multiple layers

3. Standards/Technologies

3.1. LoRaWAN

LoRaWAN is a technology developed by the LoRa alliance. LoRaWAN is a media access layer to LoRa, which is the physical layer. LoRa is a patented IoT wireless data communication acquired by SemTech. LoRa uses ISM frequencies and has very long range of transmissions up to 30 miles. LoRa uses less bandwidth, which helps in saving power mode and there by extending battery life of the devices >10 years.

3.2. LTE-Cat0/LTE-Cat1/LTE-Cat M1 (eMTC)

LTE-Cat0/LTE-Cat1/LTE-Cat M1 (eMTC) technologies are based on the mobile networks developed using 3GPP Release 8 to Release 13 specifications. These technologies will co-exist with the existing 4G networks. Each of the technologies are deferred based on the bandwidth. For the applications which require higher data usage utilizes Cat0 & Cat1 which has the highest bandwidth. For less data consumption or sporadic data consumption Cat M1 is being used to extend the battery life with proper sleep modes.

3.3. NB-IoT

NB-IoT (Narrow Band) technology is also developed based on the 3GPP Release 13 specifications using a subset of LTE standard but with a much narrower band. This can be deployed over the existing 2G/3G/4G spectrums as well. Due to its narrow band, the data uplink/downlink is less than the other LTE technologies thus increasing the battery life. Currently development of NB-IoT Roaming and NB-IoT over 5G network is in progress.

3.4. EC-GSM-IoT

EC-GSM-IoT (Extended Coverage GSM) technology is also developed based on the 3GPP Release 13 specifications based on eGPRS designed for extended coverage with high capacity, long range, low power and lower complexity by leveraging the existing 2G/3G technologies.

3.5. Sigfox

Sigfox technology is a proprietary technology developed by the company with the same name. This technology utilizes an ultra-narrowband within the ISM radio band, thus enabling the low power requirement. This technology is being used for devices which need a limited amount of data. Due to these factors the cost of Sigfox devices is less than \$3 and their battery life is 10+ years.

4. Technology Comparison

Table 1 shows the differences between the key features for the technologies discussed above.

Table 1 – Comparison of Technologies

Technology/ Feature	LTE Cat1	LTE Cat 0	LTE Cat M1 (eMTC)	LTE Cat NB1 (NB-IoT)	EC-GSM-IoT	LoRaWAN	Sigfox
Bandwidth	1.4 – 20 MHz	1.4 – 20 MHz	1.4 MHz	180 kHz	200 kHz	125kHz	200Hz
Spectrum	Licensed	Licensed	Licensed	Licensed	Licensed	ISM	ISM
Frequency Bands	700-2100MHz	700-2100MHz	700-2100MHz	700-2100MHz	700-2100MHz	915MHz	915MHz
Standardization	3gpp Release 8	3gpp Release 12	3gpp Release 13	3gpp Release 13	3gpp Release 13	LoRa Alliance	ETSI
Uplink	5 Mbit/s	1 Mbit/s	1 Mbit/s	250 kbit/s, 20 kbit/s	474 kbit/s, 2 Mbit/s	50kbit/s	100bps
Downlink	10 Mbit/s	1 Mbit/s	1 Mbit/s	250 kbit/s	474 kbit/s, 2 Mbit/s	50kbit/s	600bps
Latency	50–100ms	na	10ms–15ms	1.6s–10s	700ms–2s	1-10s	1-30s
Duplex Mode	Full Duplex	Full or Half Duplex	Full or Half Duplex	Half Duplex	Half Duplex	Full or Half Duplex	Half Duplex
Coupling Loss	144dB	144dB	156dB	164dB	164dB	157dB	153dB
Cost (in \$)	> 10\$	> 10\$	< 10\$	< 5\$	< 7\$	< 7\$	< 3\$
Batter Life (in years)	5	5	10	10+	10	10+	10+

5. Three dimensional view of Parameters vs Technology

Bandwidth, Battery Life, and Range are the three key main features which play a major role in deciding the technology type.

The below picture shows the three dimensional view of WiFi, LTE, LPWA, considering the above three factors.




Figure 2 – 3D view of key parameters vs technology

6. Radio Features

6.1. Bandwidth

Based on the usage, bandwidth plays a crucial role. Bandwidth also plays a major role in power, penetration, range, and latency. Bandwidths vary from 20MHz to 200Hz. Machines which need a limited amount of data utilize less bandwidth.

6.2. Spectrum

Spectrum requirements vary based on the industry and technology. Technologies using Mobile/Cellular communication are based on license based spectrum. LoRaWAN, Sigfox, RMPA, NB-fi are based on the ISM spectrum.

6.3. Frequency Bands

Licensed spectrum bands vary from 700Mhz to 2100Mhz across the globe. 2G/3G/4G technologies fall in these frequency bands. LoRaWAN, SigFox, RMPA, NB-Fi are based on the ISM spectrum, which is centered around 915Mhz.

6.4. Standardization

Standardization provides guidelines for features like security, payload, transmission, inter-operability. LTE-Cat0/LTE-Cat1/LTE-Cat M1 (eMTC), NB-IoT and EC-GSM-IoT are developed based on the 3GPP standards from Release 12 to Release 13. LoRaWAN uses the standards developed by LoRa alliance. Sigfox is developing based on the standards/regulations set by ETSI 300-220 and FCC part 15.

6.5. Uplink & Downlink

Uplink & Downlink play a major role in determining bandwidth, low power utilization, data formats, protocols and security features. Due to power restrictions, devices will not be continuously transmitting and receiving. Time for on-the-air and continuous burst is restricted. The higher the bandwidth, the higher the data rate and higher the power utilization. LTE-Cat0/LTE-Cat1/LTE-Cat M1 (eMTC) technologies have higher bandwidth, which support higher data transmission rates, and with less battery life. NB-IoT,

EC-GSM-IoT and LoRaWAN have intermediate bandwidth, supporting kbps transmission rates. Sigfox utilizes ultra-narrow bandwidth, which supports only hundreds of bps rates.

6.6. Latency

Latency is effected by many other features like bandwidth, range, and power saving mode. It supports extended buffering of the downlink data packets when the user equipment is in “sleep” or power-saving mode, and will start re-transmitting when the UE becomes reachable again.

6.7. Coupling Loss

Coupling loss is calculated based on device transmit power, occupied channel bandwidth, receiver noise figure and signal-to-noise ratio.

7. Non Radio Features

7.1. Power Saving Mode

This feature is supported by some user equipment, enabling the device to reduce power consumption by sending the device into deep sleep mode. This mode is intended for devices with infrequent data transmission and which can accept latency at the termination end/user equipment. Devices use timers to listen to the paging channel and only “wake up” when they hear that network traffic is intended for that device.

7.2. Extended Idle Discontinuous reception Mode

eDRX stands for extended idle discontinuous reception. Unlike power saving mode, this device mode doesn’t listen to paging and downlink channels, and also turns off part of the circuitry to save power. In this case, the network should support the same frequency as the device when it turns back on.

7.3. Cloud Compatibility

To save processing and reduce power consumption, all non-critical data from multiple device types will be sent to the cloud over the network to perform computing, calculations and processing of the data. This will also help in accessing data remotely for any further data analytics.

7.4. Artificial Intelligence & Machine Learning

Artificial intelligence and machine learning are applied to the data that is cloud processed for further analysis. By applying smart algorithms, issues are identified and machine learning is used to identify patterns and anomalies.

7.5. Analytics and Big Data

Data from multiple sensor types is collected and used to identify the patterns. This can also help in identifying possible information gaps and white space areas.

7.6. Low Cost

Considering the above factors would ultimately help in reducing device cost, network utilization, power savings and increasing processing and performance. They also help in reducing the maintenance costs over the device lifespan.

7.7. Security Considerations

Proper security considerations need to be considered as device-network mutual network authentication and encryption/ciphering of device-network data consume processing power and bandwidth. It is assumed that 10% of the power and bandwidth is being consumed for security. LTE-Cat0/LTE-Cat1/LTE-Cat M1 (eMTC), NB-IoT and EC-GSM-IoT use almost the same security features as mentioned for 3GPP.

8. Application Use Cases

8.1. Agriculture

Agricultural monitoring sensors can help farmers in measuring soil moisture, growth of crops, humidity, temperature, livestock tracking, remote harvesting, automatic and remote watering, efficient water usage. These features can help farmers in the automation of agriculture.

The above cases don't need the data to be sent continuously. Data can be sent periodically or when a particular condition is met. These devices can send data in intervals and in smaller bandwidths back to the farmers.

8.2. Smart Wearables

This application is mainly for users who require health monitoring, including the elderly and healthcare patients. Smart wearables are able to capture all the vitals of the person wearing them. This can help in remotely tracking a person's health and triggering alarms automatically, as soon as the vital metric drops as opposed to pressing a button which may not be possible in some cases. Hypothermia is a use case where the heart beat drops instantly and causes the person to become confused.

The above applications need devices to send data periodically or with an event trigger with less latency. Also these devices need ultra-narrow bandwidth thereby resulting in less data consumption and less power consumption, as well as excellent indoor coverage.

8.3. City Management and Metering

Smart metering, smart remote monitoring of electricity grid, smart waste management, and smart parking can enable cities, municipalities, and private corporations to collect data remotely. These events can be triggered periodically. Two event trigger examples are a sudden uptick in grid consumption, or when waste bins are full and require pickup.

All of the above applications require smaller and non-continuous bandwidth.

8.4. Vending Machines

Vending machines need to be monitored for these reasons: credit card payments (including proper authentication, privacy and verification of data), vending stock tracking, device diagnostic reports, and raising burglary alarm.

For the above conditions a channel with less latency and more security needs to be chosen.

8.5. Industries

Remotely monitoring of temperatures, humidity, safety monitoring, machinery control and propane tank monitoring are some examples which can be monitored remotely.

8.6. Environmental Monitoring

Illegal logging of trees or illegal poaching of rare and endangered species can be monitored by using sensors to trigger whenever a particular noise is generated or by tracking the species.

8.7. Car Auto-Piloting – Original Idea

Most Car Auto-Piloting failures are happening whenever a static object is placed in the roadway, due to the auto-pilot being unable to instantly reduce speed. A new idea is to place sensors whenever there is a road block or maintenance activity underway. These sensors can send the updates to the navigation systems like Google maps or Waze, thereby there-by alerting the auto-pilot vehicles of possible hazards.

9. NB-IoT Roaming

On June 4th, 2018 Deutsche Telekom and Vodafone group successfully completed the first international roaming trial for NB-IoT. These service will be helpful in cases when devices need to exceed geographic boundaries and move from one place to another. For example, the tracking of shipping containers which cross international boundaries from one country to another.

10. NB-IoT 5G

3GPP already has standards for 5G LPWA use cases by evolving NB-IoT and LTE-M as a part of the 5G specifications and co-existing with other 5G components.

Conclusion

LPWA has evolved into many technologies based on many key parameters such as bandwidth, performance, power consumption, latency, spectrum, security. IoT devices/sensors connectivity often require less data consumption, with more battery life, and with co-existing/leveraging of existing networks. Based on these needs, different technologies can be chosen. Evolving and emerging technologies are constantly being researched.

Abbreviations

LoRa	Long Range
GSMA	Global System Mobile association
NB-Fi	Narrow Band Fidelity
Wi-Fi	Wireless Fidelity
RPMA	Random Phase Multiple access.
ISM	Industrial Scientific Medical
NB-IoT	Narrow Band Internet of Things
eMTC	enhanced Machine Type Communications
Cat	Category
LoRaWAN	Long Range Wide Area Network
EC-GSM-IoT	Extended Coverage Global System Mobile Internet of Things
EDGE	Enhanced Data Rates for GSM evolution.
CAGR	compound annual growth rate
RAN	Radio access Network
SMSC	Short Message Servicing Center
MME	Mobility Management Entity
SGSN	Serving GPRS support Node
SGW	Serving Gateway
SCEF	Service Capability Exposure Function
HSS	Home Subscriber Server
AS	Application Server.
3GPP	Third Generation Partnership Project
AP	access point
bps	bits per second
Hz	Hertz
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

3GPP Low Power Wide Area Technologies – GSMA White Paper
Mobile IoT in the 5G future – NB-IoT and LTE-M in the context of 5G - GSMA White Paper
LPWA Technology Security Comparison – Franklin Health White Paper
NarrowBand IoT, The game changer for the Internet of Things – Tmobile White Paper
NB-IoT Deployment guide to Basic Feature Set requirements Version 2.0 - GSMA White Paper
Enabling Wide Area IoT Solutions with machineQ, A Comcast Service – Comcast White Paper

Computing At The Edge Still Has An Edge

A Technical Paper prepared for SCTE•ISBE by

Arun Ravisankar

Senior Engineer, Comcast Labs

Comcast Corporation

1701 JFK BLVD, Philadelphia, PA 19103

Phone:2152867558

Arun_Ravisankar@comcast.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Machine Learning Overview.....	4
Computing at the Edge	6
1. Use Cases.....	8
Conclusion.....	13
Abbreviations	14

List of Figures

Title	Page Number
Figure 1 - Evolution of technology and its influence in the society	3
Figure 2 - Machine Learning Process and Events involved.....	5
Figure 3 - Machine Learning Systems	6
Figure 4 - Cloud-based Inferencing Engine	7
Figure 5 - Driver-assist features in a car.....	8
Figure 6 - Example analyses of video from a security camera	9
Figure 7 - Activity Determination using Machine Learning	10
Figure 8 - Edge compute process example	11
Figure 9 - Object recognition using Machine Learning	12
Figure 10 - Sample flow in a Voice command system.....	13

List of Tables

Title	Page Number
Table 1 - Machine learning algorithms and examples	4

Introduction

History is witness to the evolution of civilizations and how humans continue to discover and innovate things that would propel everyone to a newer level 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. The first industrial revolution was about mechanization, which involved the development of machine tools and the rise of huge factories and factory systems. The second revolution, also known as the Technological Revolution, brought about a rapid rise in industrialization, which involved increases in automation. Digitization can be seen as the third industrial revolution, where digital systems of all types saw an increase in adoption.

The fourth industrial revolution could be envisioned as a function of AI (Artificial Intelligence) and ML (Machine Learning), which are vital in building “Intelligent Machines.” It follows that those “Intelligent Machines” could be referenced as “the compute edge,” as opposed to “the network edge” -- in our case, usually defined as the node, where optical-to-RF conversion occurs. AI/ML technologies influence a large part of the devices and services we use on a daily basis, be it a voice assistant or vehicular parking assist, or be it an entertainment platform that understands our preferences and predicts shows and titles we may like. Apart from these examples, many AI/ML-based applications can help improve lifestyles and bring peace of mind to customers.




Figure 1 - Evolution of technology and its influence in the society

AI/ML plays a vital role in almost any products and services that are offered to customers now. Any application or service rendered in a customer’s home, be it via a set-top box (STB), DOCSIS-based gateway, home automation gateway, or IoT device, involves multiple components working in tandem. The “compute edge” discussed in this paper is comprised of those in-home devices. Because “edges” in general vary widely, for the fourth industrial revolution -- AI and ML -- we define the “compute edge” as

the premise. That necessarily includes devices in the premises, linked to applications running on cloud servers that are racked up in a data center.

IoT applications process data from devices at the edge and are subject to a decision tree usually deployed on a cloud server. The decision tree or rules engine determines the course of action for data sent from a device. With the advent of machine learning and artificial intelligence, and given their natural fit with IoT applications, the demand for higher computing power has increased significantly. Now, with the increase in silicon capabilities that accelerate AI/ML algorithms, devices on the edge can process some information locally, which move some parts of the decision tree to the edge. This paper will discuss how edge compute could improve the delivery of IoT applications.

Machine Learning Overview

Machine Learning techniques involve statistical algorithms that give computers the ability to learn. This helps machines to progressively improve their performance on a specific task. The learning process is automatic relative to the data being gathered, and does not involve explicit programming. Most Machine Learning algorithms could be grouped into the following classes (see Table 1). These algorithmic classes add value to service providers:

1. Classifiers
2. Clustering Algorithms
3. Recommender Systems
4. Anomaly Detection Algorithms
5. Linear Regression

The table below shows a basic description and examples of each of the above algorithms.

Table 1 - Machine learning algorithms and examples

Class of Algorithms	Description	Technology Examples	Applications
Classifiers	Assigns new inputs to one or more classes, based on similarity to other data	Neural Networks	Image Classification, Spam filtering
Clustering Algorithms	Groups similar data into clusters	K-Means	User Profiles and anomaly detection
Recommender Systems	Makes recommendations based on historical data	Filtering	Product recommendations
Anomaly Detection	Detects rare events, usually not normal	Joint Probabilistic modelling	Fraud detection, Home Security and healthcare use cases
Linear Regression	Predicts values for continuous variables	Linear Regression	Churn Rate Prediction

Machine Learning application development usually involves two parallel, yet connected, processes. One is a modelling workstream, and the other workstream involves deployment of the models. The first workstream, as its name indicates, is more centered on the modeling effort. The second workstream focuses on ensuring that there is a path to deployment for the models being developed. The two efforts

are viewed as happening concurrently, because of the complex nature of deploying a machine learning solution in a cable system operator's production environment.

Figure 2 shows the process in a typical machine learning-based application. The aspects shown in Figure 2 can be categorized into two tracks. One track is of model development and other track would be of deployment and integration.



Figure 2 - Machine Learning Process and Events Involved

The events and steps shown in Figure 2 could be split into model development and deployment. Model *development* includes the following characteristics:

- Business/Data Understanding
- Data Preparation
- Modelling

Model *deployment* includes the following characteristics:

- Evaluation
- Deployment and Integration

Model development involves the “learning” process, where training data is used to build models. Learning processes are usually done on high performance systems and are resource intensive, as learning process involves processing a large amount of data to prepare models.

The model is deployed and executed based on the data that is received by the system. The execution depends on the use case; the models are built based on those use cases.

Computing at the Edge

Once the models are developed, they are deployed on devices that execute these models, on test data, and arrive at conclusions that depend on the particular use case. Figure 3 shows a simple depiction of the learning process, where systems work on the data to create models. These systems are compute-intensive, and require necessary infrastructure to be set up.




Figure 3 - Machine Learning Systems

Once the models are ready for deployment, they are deployed on systems that can apply them to the live data coming in from the various sources. For example, a video analytics-based ML application would use models that were trained using images and videos. Once the model is trained, images and video from a camera are analyzed by applying these models. In this specific example, the inferencing engine needs to process the video signals and then apply the model as deployed. The use cases could vary between, say, monitoring an area to monitoring facial expressions. Hence the inferencing engine would also require high performance computing in order to provide results accurately, with minimum latency. It would be a stretch for the customer premises equipment presently deployed to meet these compute requirements. Because of the need for high levels of processing, inferencing engines are often deployed in a cloud infrastructure, where units could be racked to meet the compute and power requirements.

Figure 4 shows how a cloud-based inferencing engine would operate on the data being ingested to provide services to the consumer by executing the rules defined by the models.




Figure 4 - Cloud-based Inferencing Engine

Newer, advanced hardware platforms offer higher compute performance, while requiring less power and memory. These hardware acceleration platforms provide the basis to run AI/ML-based algorithms. Compute power throughput is measured in Tera Operations Per Second (TOPS) or Tera Floating-point Operations Per Second ([TFLOPS](#)). Most AI/ML-based applications run on platforms that offer about 0.5 to 1 TFLOPS. Another important metric is the efficiency of the processor architecture, and is measured in GFLOPS/W, which translates to Giga Floating-point Operations Per Second per Watt of energy consumed.

The efficiency factor determines if the system is best deployed in a rack at a location and services are accessed through cloud, or if the system could be deployed at the edge, again meaning the premise.

Major and sustained advancements in silicon manufacturing have led to the development of high efficiency processors that can support a throughput that is comparable to most of the high-performance CPUs that are deployed. Next, we will look at the use cases best suited for these processors, in terms of improving the overall experience with ML/AI.

One of the major advantages of computing at the edge is the improvement in latency of the system, because the data is processed at the premises, rather than being sent over a network to a server-based processing engine. Hence, applications that need fast response times tend to require edge compute resources.

1. Use Cases

Automotive Applications: Most cars now offer several driver-assist features that use a variety of sensors. The data coming in from these sensors needs to be processed in real-time, so that alerts or actions can be executed. This involves processing a lot of data, and the processing needs minimal latency. Apart from driver-assist features, as shown in Figure 5, there is an increased level of interest in the automobile industry to build [self-driving](#) cars. These cars function similarly to airplane auto-pilot mechanisms, where human intervention is required in specific circumstances. Imagine the amount of computing involved, if we need to match the sophistication that is equivalent to an airplane! This needs a prohibitively large amount of computing -- and the computing has to happen in real-time. In such cases, most and in fact all of the computing needs to happen at the edge (in this case, the on-board computer of the automobile). The system can then process signals from various sensors and initiate appropriate actions.




Figure 5 - Driver-assist features in a car

IoT Applications: Beyond how edge computing technology drives the development of next-gen automobiles, these systems could also be leveraged in IoT (Internet of Things) applications, like home monitoring, security, and healthcare, to name a few. Many IoT applications involve anomaly detection, in which the application processes data coming in from various sensors. These applications deploy machine learning models, that process data from various sensors -- for example, an application that detect events based on video feeds from security cameras, which use computer-vision based models to analyze the current situation. This also involves a significant amount of processing, for both video and AI.

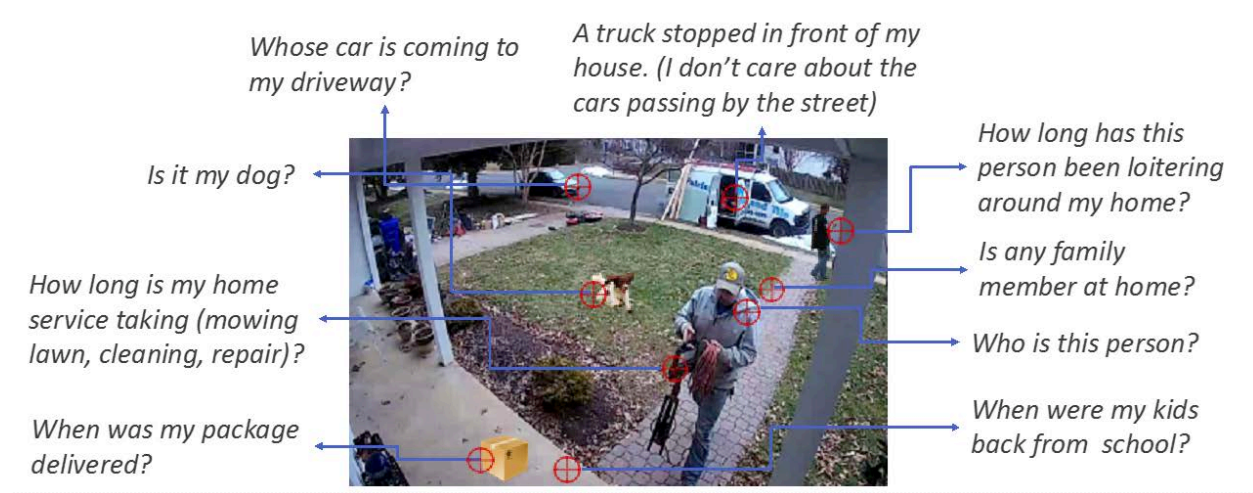


Figure 6 - Example analyses of video from a security camera

Figure 6 depicts examples of events that homeowners tend to be interested in knowing: Who's at the door, what's that truck parked out front, and so on. In a camera-based application, AI/ML based models are deployed on high performance inferencing engines to analyze data derived from the camera feed. These models can be deployed on processors that provide acceleration to AI/ML models to perform tasks that are time-critical at the premise, while further processing and learning could be carried out in cloud-based servers.

If the camera (or any premises equipment) is built with silicon that provides hardware-based acceleration to run AI/ML models and algorithms, there could be significant improvement in latency of the system. Apart from latency, there would also be an improved sense of privacy (in a camera-based application), because the images are being analyzed at the premises and might not ever leave the premises.

Healthcare: IoT technologies contribute immensely to connected healthcare applications. With the use of IoT and AI/ML technologies, monitoring health and wellness could be significantly expanded to provide peace of mind to people who care for family members and patients. Eldercare is a classic example: A combination of IoT and AI/ML technologies can be used to monitor daily activities of the elderly, and notify either the care provider or the family member in the event of perceived abnormalities. Solutions can be built that detect falls, or analyze gaits and alert the appropriate caregiver. Using computer vision, itself a subsystem of AI/ML, such systems can identify both objects and people, and can determine activities, detect falls and otherwise inform a healthcare application. Figure 7 shows an application that can determine the activity of individuals using video analytics.

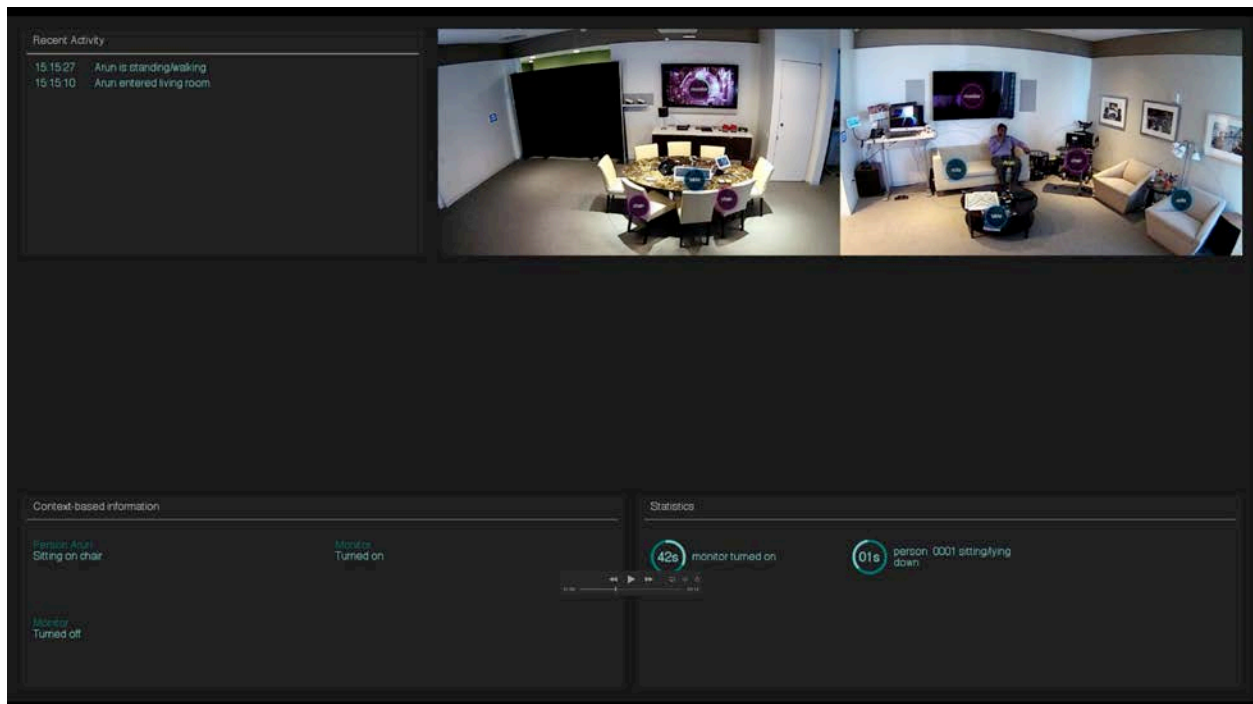


Figure 7 - Activity Determination using Machine Learning

The application depicted here can recognize common household objects like tables, chairs, couches, and TVs. The application can also determine if the TV is on/off, and can identify a person, in a way that is differentiated from a visitor coming into the home. This can be seen in Figure 9, which shows how the objects in the room are identified. Figure 8 explains the process in which an edge compute system could operate. The example chosen is a computer-vision based system, where the models are developed in a cloud-based system by analyzing a vast amount of training data. These models are deployed on the edge system (premises) and the software on the edge uses the on-board AI/ML acceleration features to perform inferencing and display results to the user.

In this example, the camera is used as a sensor. Similar applications could be built using other sensors deployed in home, for example, motion sensors, or door/window sensors. We can also look at the potential of RF sensing for these use cases (including both WiFi and RADAR). The type of sensors used determines the data format and hence the models that are created. Appropriate models have to be developed to work on the sensor, such that it meets the application's requirements. A computer vision-based application was an easier choice for a proof of concept, because of the vast sets of training data available, and because the training data can be continuously generated using a camera.

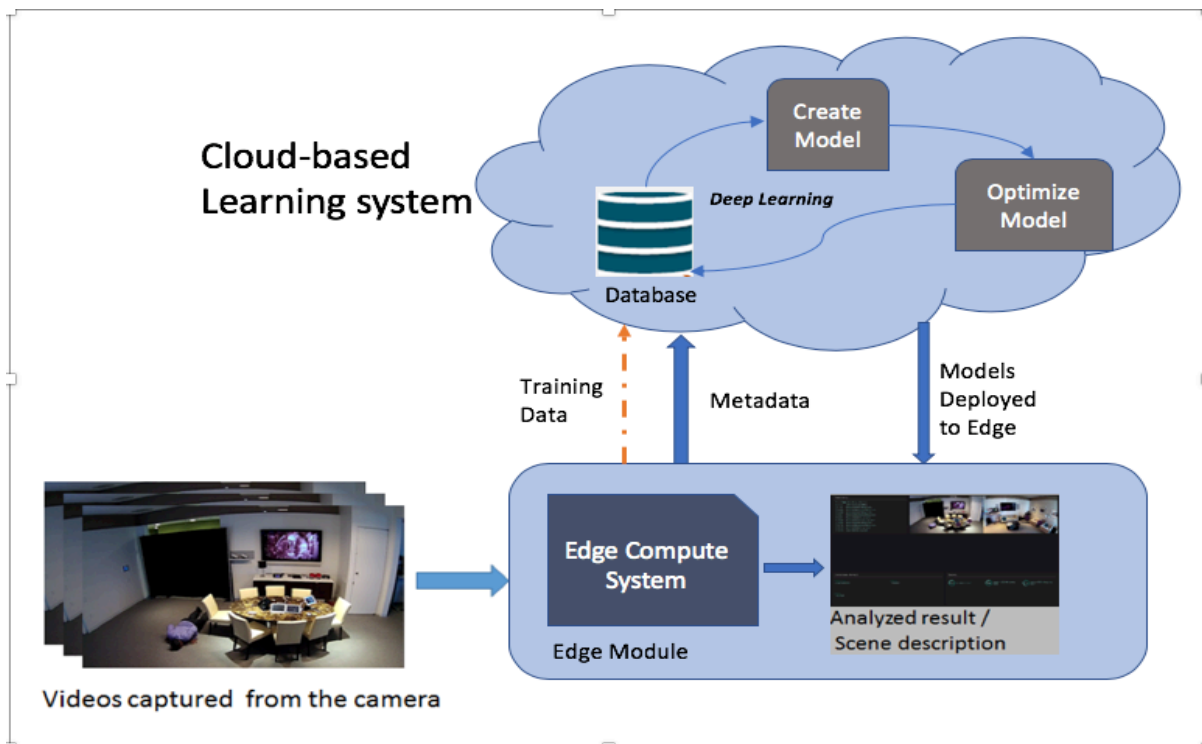


Figure 8 - Edge compute process example

Room before recognizing objects



Room after recognizing objects



Figure 9 - Object recognition using Machine Learning

Home Automation/Smart Assistants: Another relevant aspect of everyday life is the role smart-assistants and home automation can play. Interaction with these devices is gradually increasing. Devices like Alexa and Google Home have become an everyday lifestyle tool for many people. Be it “Alexa, where are my keys?” or “Hey Google play my favorite radio station”, we use these smart assistants for a variety of purposes. These systems, including the voice remote, which is seeing steadily increased usage, are based on machine learning applications. They have to process speech and the language being spoken by the user. For this they use Natural Language Processing (NLP) algorithms. In most cases, a microphone lists the words spoken by the user (upon trigger/wake word) and then sends the corresponding audio packets to a system that converts speech into text, so that a computer can decipher the contents. Once converted to text, the data is processed using NLP to understand the requests from the user, and advanced AI methods are applied to understand the context and intent factor, so that the response to a query is as accurate as possible.

Most of the AI/ML processing is done on high performance systems, and for that reason, the compute edge can play an important role to augment the processing by performing some analysis on the edge, while more complex analysis and learning is done in cloud. This improves latency and is additionally useful in scenarios where network connectivity is poor or lost.

As shown in Figure 10, the processing of speech and analysis mostly happens in a cloud-based system, yet sometimes, the speech to text could be done on the edge, with assistance on NLP and AI in the cloud before the response is sent back to the user.

With the continuing advances in silicon and processor architectures, part of the processing, including the NLP, could move to the compute edge. This would both improve response times and help in situations where the network connectivity is poor or suffering an outage.

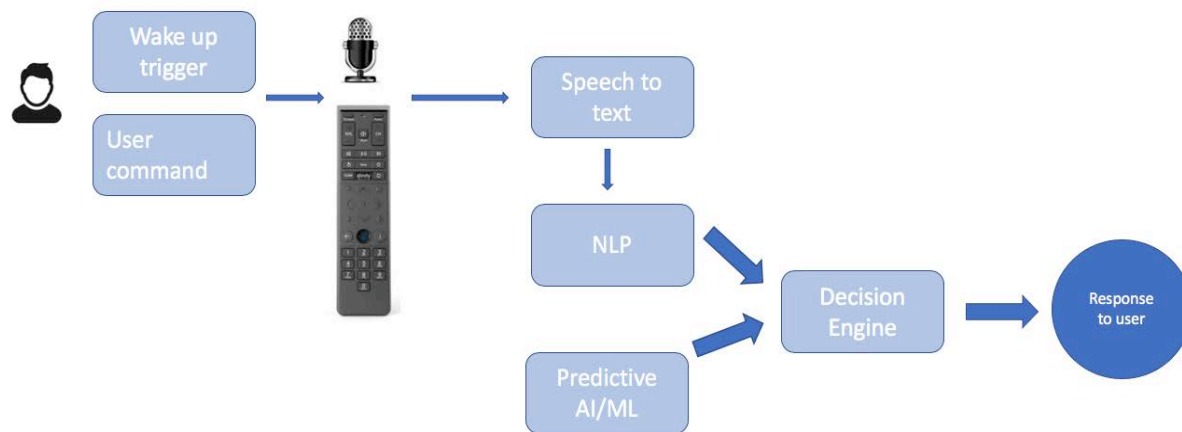


Figure 10 - Sample flow in a Voice command system

Customer Experience: AI/ML applications are playing an important role in improving the customer experience. Predictive AI is an example. With it, the system predicts an issue based on the data points, or can be applied to customer interfaces for quick issue resolution, using chatbots or self-healing techniques. The CPE (Customer Premises Equipment) is home to a vast set of (anonymized) telemetry and troubleshooting data that can be used as an indicator of network health and system status. These data points are used to build models, and when these models are applied to live data, the system can proactively predict issues like outages. These models are simpler than the models discussed earlier, as the data sets are usually available within the premises. These models are a good fit for compute edge use cases, because the system can analyze data and provide recommendations to the user or technician visiting the premises.

Conclusion

If the fourth industrial revolution does indeed turn out to be spawned by the swift and productive rise of AI (Artificial Intelligence) and ML (Machine Learning) technologies, both of which are vital in building “Intelligent Machines,” then those same intelligent machines introduce a new “edge” to the network: The *compute edge*.

In this paper, we looked at the metrics of such compute edge platforms, the efficiency of the platforms and how newer hardware is emerging with higher performance and efficiency. We examined relevant use cases where the compute edge can improve response times and improve the sense of privacy. Also, a compute edge system can augment and complement existing cloud-based systems with more of a near-field analysis.

The edge platforms that are discussed here are not envisioned as replacing the cloud-based systems, but rather to enhance the efficiency and to better distribute the processing responsibilities. One major advantage is to be able to make minimal use of a system, when there is an outage or an intentional

sabotage. The system could provide the first level of AI capabilities and could leverage the cloud systems for further detailed analysis.

While we are excited about the compute edge platforms, we have to note that cloud-based systems are comparatively easy to maintain, because they enjoy a one-to-many relationship. It would add complexity in the system to maintain various compute edge platforms. Such challenges could be mitigated, to an extent, by using the same model structure. Suffice it to say there is still a lot of ground to cover, and such systems would need to be vetted.

Abbreviations

AI/ML	Artificial Intelligence/Machine Learning
bps	bits per second
CPE	Customer Premises Equipment
CPU	Central Processing Unit
DOCSIS	Data Over Cable Service Interface Specification
Hz	hertz
IoT	Internet of Things
ISBE	International Society of Broadband Experts
NLP	Natural Language Processing
RADAR	Radio Detectino and Ranging
SCTE	Society of Cable Telecommunications Engineers
STB	Set Top Box
TFLOPS	Tera floating point operations per second

Converged Multi Access Networks

A Technical Paper prepared for SCTE•ISBE by

Amit Singh

Principal Engineer
Cisco Systems

375 West Tasman Drive, San Jose 95154
408-527-1856
amsingh@cisco.com

Eric D. Heaton

Network Solutions Architect
Intel Systems, Inc.

2200 Mission College Blvd., Santa Clara, CA 95054
Architect, Platform Solutions Group (DPG/NSG)
eric.d.heaton@intel.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Key Technology Evolution.....	3
1. Access Networks Today.....	3
2. Access Technology Evolution	6
2.1. Virtualization & Disaggregation	6
2.2. Deep Fiber Migration.....	7
Advantages of Converged Access Networks.....	8
Multi Access Network Topologies	11
Conclusion.....	14
Abbreviations	14

List of Figures

Title	Page Number
Figure 1 - Access Technology Deployments Today	4
Figure 2 - Today's access network	5
Figure 3 - Virtualization and Disaggregation of Access Technologies	7
Figure 4 - Converged CIN	7
Figure 5 - Common IP Services Chain	9
Figure 6 - Operational Efficiency & Automation.....	10
Figure 7 - Fronthaul Multi Access Network.....	11
Figure 8 - Converged Multi-Access Layer.....	12
Figure 9 - The Converged Multi-Access Network	13

Introduction

This paper discusses the benefits of converged multi access networks. The paper considers three main access technologies. They are Cable, Mobility & PON. As these three access technologies evolve, opportunities to evolve the aggregation and edge networks emerge. Newer network architectures offer unprecedented opportunities for Multi Service Operators (MSOs) and Service Providers (SP) in architecting networks, streamlining and simplifying operations, and opportunities to offer new services at reduced cost. Converged multi access networks also offer benefits for subscribers, as their internet experience will be uniform regardless of the access media/type.

The paper discusses key technology evolutions first. These are enablers for the converged multi access network. It presents the components of the converged multi-access network and discusses the converged architecture while contrasting them to today's network build outs. The benefits of converging the network are presented and discussed. Finally, the converged network topology is presented. The paper concludes with the benefits of such an architecture to MSOs/SPs and subscribers.

Key Technology Evolution

Access network technologies are undergoing two key technology evolutions. These are catalysts to evolve the access network to a converged access network. The two technology evolutions are migration from big iron hardware to virtualization and disaggregation, and deep fiber migration of access networks.

The primary goal of the access network evolution is to provide 10 Gbps to the home at reduced latency and cost. There is uniformity in the way different access technologies are approaching solutions. All three, Cable, Mobility & PON are transforming to deep fiber closer to the subscriber neighborhoods.

Before looking into how these key technology evolutions will help transform the access network, let us look at how Cable, PON & Mobile access networks are deployed today and the consequences of those buildouts on overall operations and associated costs.

1. Access Networks Today

Figure 1 shows a typical picture of big iron deployments of access technologies.

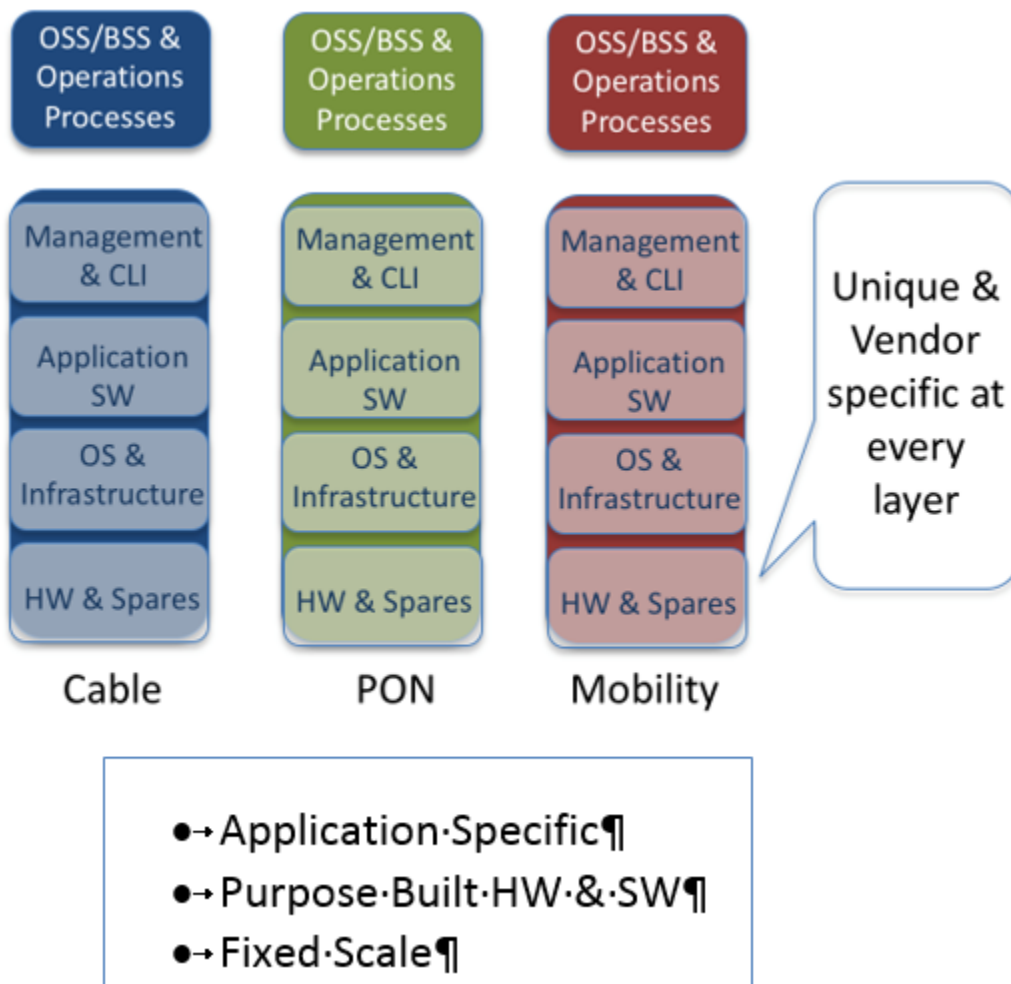


Figure 1 - Access Technology Deployments Today

Access technologies are custom built today. They consist of big iron boxes from network equipment vendors with proprietary hardware, operating systems, application software, management software and command line interface (CLI). Big iron boxes are purpose built and optimized for the access service it provides. As a consequence, operations processes to configure, monitor, upgrade and create new services are unique as they are also tailored to the capabilities and scale of the big iron box. Most operators prefer multiple vendors for the same application. Each vendor has their own unique big iron solution. As the operator manages solution variations from multiple vendors for the same application, it makes every single access technology even more bespoke and further silos the application and its operations processes. As a consequence, the edge and aggregation networks necessary to service the access application also become bespoke.

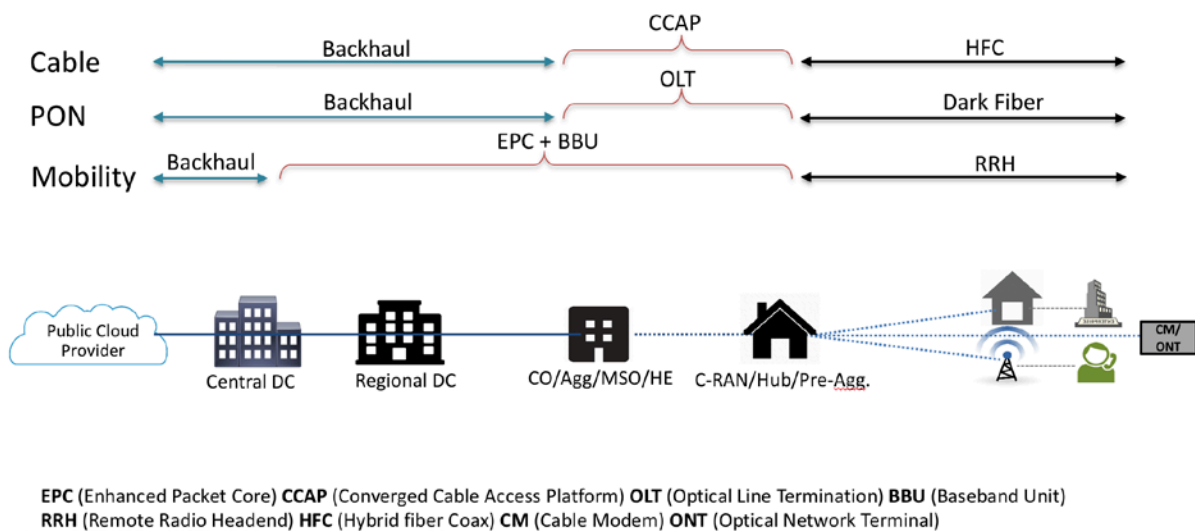


Figure 2 - Today's access network

Figure 2 shows the layout of the CIN today. The CIN is the portion of the network from the Hub/C-RAN to the RPD.

There are several differences in way the fronthaul network is deployed across access technologies.

Firstly, the CIN, is unique per access technology. In Cable access it is a Hybrid Fiber Coax network, while it is a dark or analog fiber network for PON and a 1 or 10 Gbps Ethernet network to the cell tower for Mobility. These networks are largely different in the way they are deployed, managed and operated.

Secondly, the way services are delivered across access technologies varies. For example, Cable operators deploy Converged Cable Access Platform (CCAP) devices. These devices are self-contained units with several network functions bundled together in the sheet metal. They terminate layer 1 coax cables, provide layer 2 DOCSIS capabilities, layer 3 IP routing, subscriber management and an entire stack for video distribution over the Cable network.

PON, on the other hand, disaggregates functions across at least two boxes. The layer 1 of the PON network is terminated in an Online Termination (OLT) and traffic is handed off to an edge router and subscriber management box (not shown).

Mobility also splits functionality into multiple boxes or layers. It terminates the layer 1 over air protocol at the base of the tower and hauls traffic to the Base Band Unit (BBU) and Enhanced Packet Core (EPC). The EPC performs subscriber termination and routing.

Service provisioning on Access technologies varies. Service provisioning in Cable is performed on the CCAP, whereas on PON must be accomplished on two or more devices: the OLT, the subscriber manager and routing devices. For Mobility, service provisioning is done on the EPC and the router behind it.

Thirdly, the location of access technologies in the provider network also varies. Cable operators deploy CCAP devices in hubs or head ends, PON OLTs are also deployed in hubs or head ends but the aggregation router for PON may be deployed deeper in the network. The EPC for mobility is typically much deeper in the network.

The consequence of disparate deployments of access technologies is that operations processes, network growth, subscriber and service additions become very specific to the access type.

It is very difficult to converge operations across multiple access technologies today. Most operators that offer services across multiple access technologies run these in a “ships in the night” model, where one access technology shares nothing with the other in the field network, hub or head end.

Ultimately, traffic does get aggregated at the regional or central data center but that is beyond the scope of the access network, and more in the edge portion of the network.

2. Access Technology Evolution

2.1. Virtualization & Disaggregation

Virtualizing and disaggregating access technologies is well underway. This is illustrated in Figure 3. Large portions of the overall stack, hardware, and software will be common and uniform and will therefore support fungible servicing multiple access applications. Specifically, the underlying hardware (servers and data center switch fabrics), operating system software, application infrastructure and portions of the monitoring, telemetry and automation software will be uniform across access technologies. This creates unprecedented uniformity across these technologies.

Virtualization also enables disaggregation and sharing of network functions. For example, CCAP need not be a single unique bundled instance anymore, as subscriber management and routing can be disaggregated and made common and shared across multiple access technologies.

As large portions of the overall application stack become common, access technologies from multiple vendors will not be disparate anymore. Further, common infrastructure will streamline and simplify operations. An operational component that is specific to the application technology will still exist, however, however this will be a smaller portion of the entire access technology network operation.

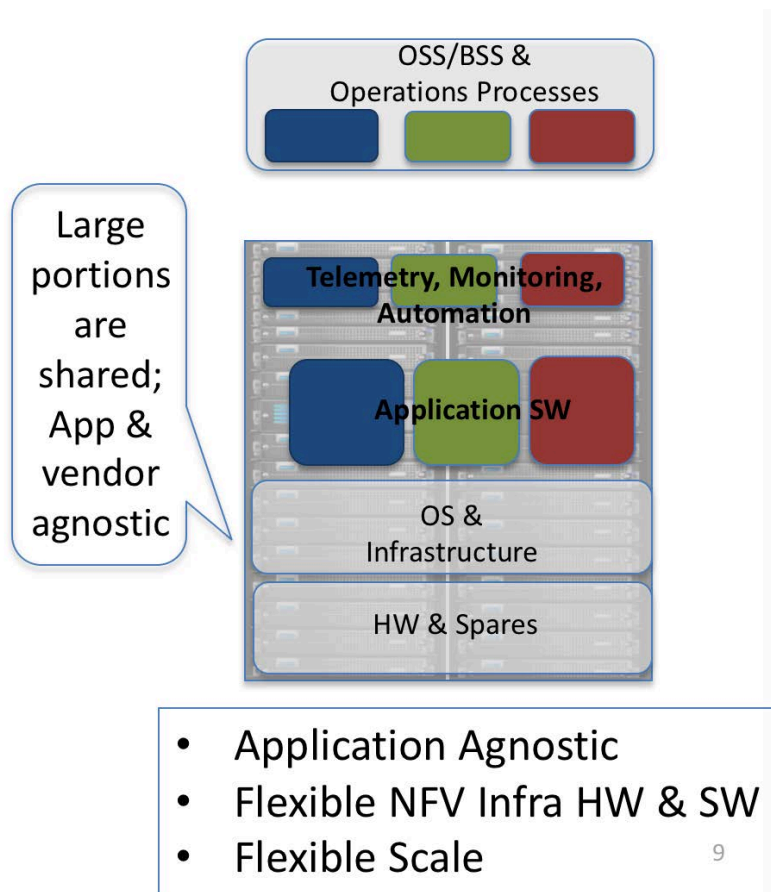


Figure 3 - Virtualization and Disaggregation of Access Technologies

2.2. Deep Fiber Migration

The CIN from the hub/C-RAN to the neighborhood is being converted to digital fiber. This is underway for all three access types, Cable, Mobility & PON. Figure 4 shows a CIN converged across access technologies.

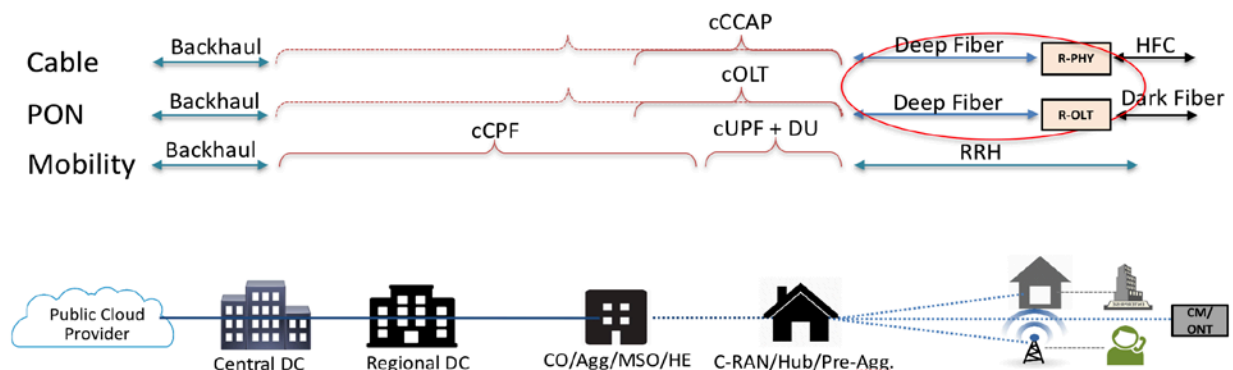


Figure 4 - Converged CIN

Convergence of the CIN for Cable, PON and mobility to 10/100 Gbps deep fiber has big implications on the ability to converge the access networks in the back end.

The deployment and management of the deep fiber portion of the network itself is agnostic of the service that runs on it. Also, it allows access specific service functions such as cCCAP (Cloud CCAP), cOLT (Cloud OLT) and cCPF (Cloud Control Plane Functions for Mobility) and cUPF (Cloud User Plane functions for Mobility) to reside anywhere from the hub to central data center.

Virtualization & disaggregation and the migration to deep fiber are key to architecting the converged multi access network as they eliminate access specific silos.

Advantages of Converged Access Networks

There are six main advantages of converged access networks. These are:

1. Common Multi Access Fronthaul Network (aka. Converged Interconnect Network)
2. Common Access Termination Infrastructure
3. Common IP Services Infrastructure
4. Operational Efficiency & Automation
5. Facilities Optimization
6. Uniform Subscriber Services and Experience

Common Multi Access Fronthaul Network: As shown in Figure 4, a common multi access fronthaul network is a 10/100 Gbps network between the hub and the neighborhood. It displaces a large portion of the access specific network. It is inherently access agnostic and can carry traffic for multiple access technologies simultaneously. This is a big advantage over today's access specific network as this portion of the fronthaul can be operated as a common portion across access technologies.

A deep fiber fronthaul pushes access specific networks to the neighborhood. Access specific distance limitations are eliminated or minimized. Also, as deep fiber covers the major portion of the distance between the subscriber access equipment and the provider access termination equipment, the latency of the network does not increase.

Consider the example of a cable network. With legacy integrated CMTS (Cable Modem Termination System) deployments, the biggest contributor to latency in the network is between the cable modem and the CMTS in the hub. The legacy HFC network is made up of analog fiber and a coax cable network at about 1000 homes passed. The coax portion of the network is the largest contributor to latency. With deep fiber penetration, the fiber portion of the network goes deeper to about 120 homes passed, reducing the length of the coax network, so latency is not impacted adversely.

With deployment of deep fiber, access specific distance limitations are eliminated. This allows access termination equipment such as a cCCAP device, cOLT and cCPF/cUPF to be located deeper in the network and more importantly co-located. Co-location of access technologies allows deployment upon a common infrastructure.

As multiple access technologies can be co-located due to deep IP fiber migration, operators can offer uniform multi-access services using the same infrastructure and fronthaul network.

Common Access Termination Infrastructure: As access technologies migrate to virtualization and disaggregation, as shown in Figure 3, they can all leverage common data center compute and network infrastructure. Coupled with the fact that multiple access technologies can be co-located due to deep fiber migration, an unprecedented new opportunity arises to offer multi-access services using the same infrastructure and fronthaul network.

Common IP Services: As access technologies decouple from the backend IP services framework due to disaggregation, a common set of IP services can be applied via service chains shared across access technologies.

As an example, consider the CCAP device that incorporates the entire routing stack. A disaggregated CCAP could offload routing, DHCP and any other security services to an IP services chain.

Figure 5 shows a sample Common IP services chains.




Figure 5 - Common IP Services Chain

A common IP services chain across access technologies creates a uniform experience for subscribers. For the provider, common IP services provide an opportunity for uniform access policy management and application. Common IP services facilitate application of uniform security policies rather than on an access-specific basis. This will improve security of the entire network.

Operational Efficiency & Automation: Several factors contribute to access network operational efficiencies and automation. As illustrated in Figure 6, the converged deep fiber fronthaul network can be operated as a common entity serving multiple access technologies. The common access services infrastructure can also be operated as a common single entity that is access agnostic. Further, the common IP services chain applied to access technologies can also be operated as a common entity.

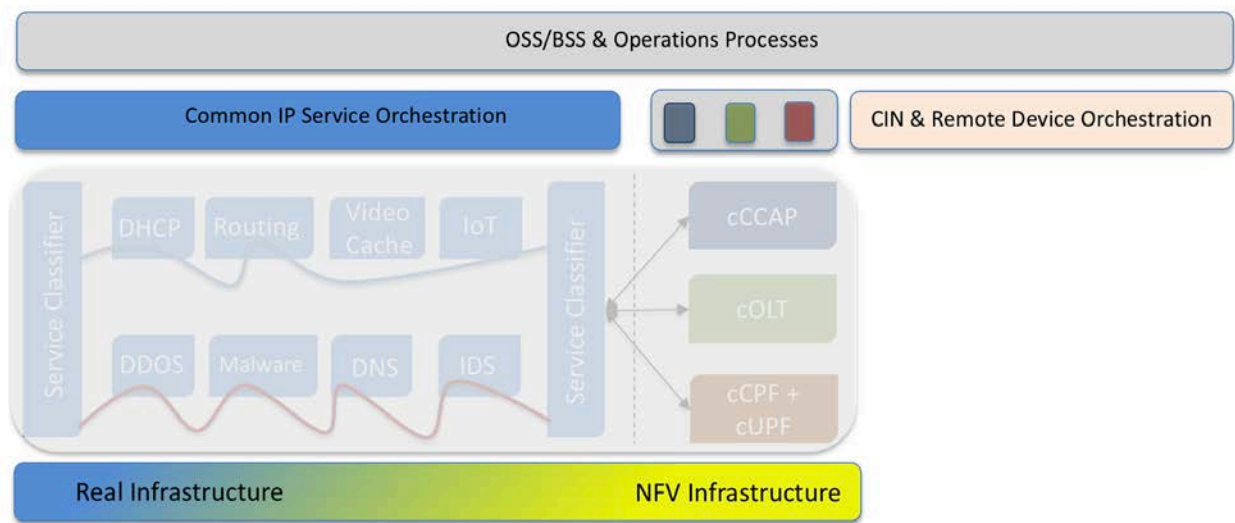


Figure 6 - Operational Efficiency & Automation

The components of the network that are still access specific such as the access technology VNFs and the last few miles of the access network are the only remnants that need access specific operation.

Taken together, most of the access network can be operated as a single common entity. This is an unprecedented opportunity to simplify and streamline operations via automation. Automation can serve to detect and repair faults via fault tolerance built into every layer of the network. Automation can also help with upgrades and application of service policy and in provisioning new services.

Facilities Optimization: A common fronthaul network, disaggregation and virtualization of access technologies and a common IP services network create many new opportunities for facilities optimization.

- Operators can offer multiple access services using the same facilities, leading to better utilization of a facility and better metrics for subscribers/facility or revenue/facility.
- Operators could lease out portions of the facility or co-locate multi-access equipment for better utilization.
- The access network could be leased out to other operators for extending their reach.
- Facility consolidation is yet another advantage, as multiple facilities are not required to host each access network.
- Finally, along with multi-access networks, providers could provide value-added services such as video cache hosting, security and peering deeper in the network.

Uniform Subscriber Services and Experience: As a common IP service chain is applied across access technologies, subscribers will attain a uniform access experience, agnostic of the access used to connect to the internet. Further, if subscriber identity is also unified across access technologies, subscribers will get a uniform experience agnostic of the device they connect to the internet with. This creates new use cases for device to device handoff as subscribers switch from one access network to another via switching access devices.

Multi Access Network Topologies

This section puts all the ideas together to propose topologies of converged multi- access networks. Figure 7 shows the fronthaul multi access network.

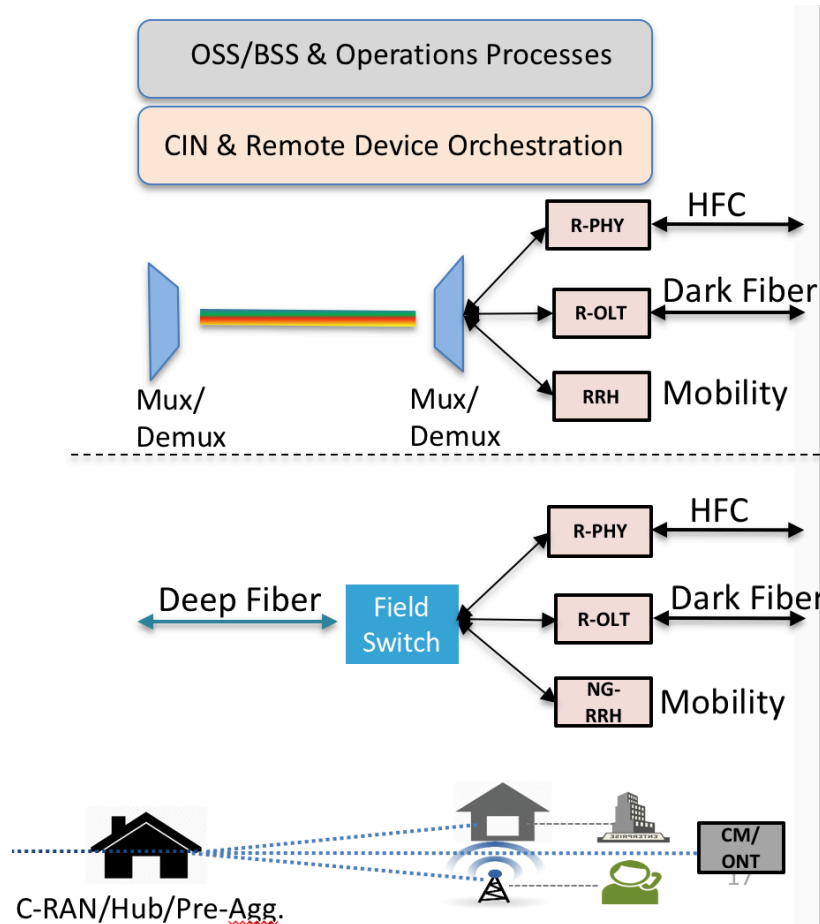


Figure 7 - Fronthaul Multi Access Network

The CIN usually extends from the provider hub to neighborhoods it services. Two topologies are illustrated. The one on the top is a DWDM optical transport topology. In this topology, one pair of mux/demuxes reside in the hub while the other resides in the field. The network is capable of hauling multiple terabits of traffic. Traffic to each of the field access technology devices (such as RPDs and R-OLTs) could be multiplexed on a wavelength or traffic for multiple devices could be multiplexed on the same wavelength. In the latter case, a field switch will be required to de-multiplex traffic for multiple access devices on the same wavelength.

The lower portion of Figure 7 shows a 100 Gbps fiber deep distribution layer from the hub to the field, terminated by a field aggregation switch. Traffic is de-multiplexed by a field aggregation switch to field access devices.

The operator could choose either of these topologies depending on the density of field devices and ultimately the number of subscribers serviced.

As the CIN is access agnostic in both cases, it can be managed via common management software. Thus, a multi access CIN can be managed using the same software layers thereby simplifying and unifying CIN operations.

Figure 8 adds on the converged multi access layer to the outside CIN.

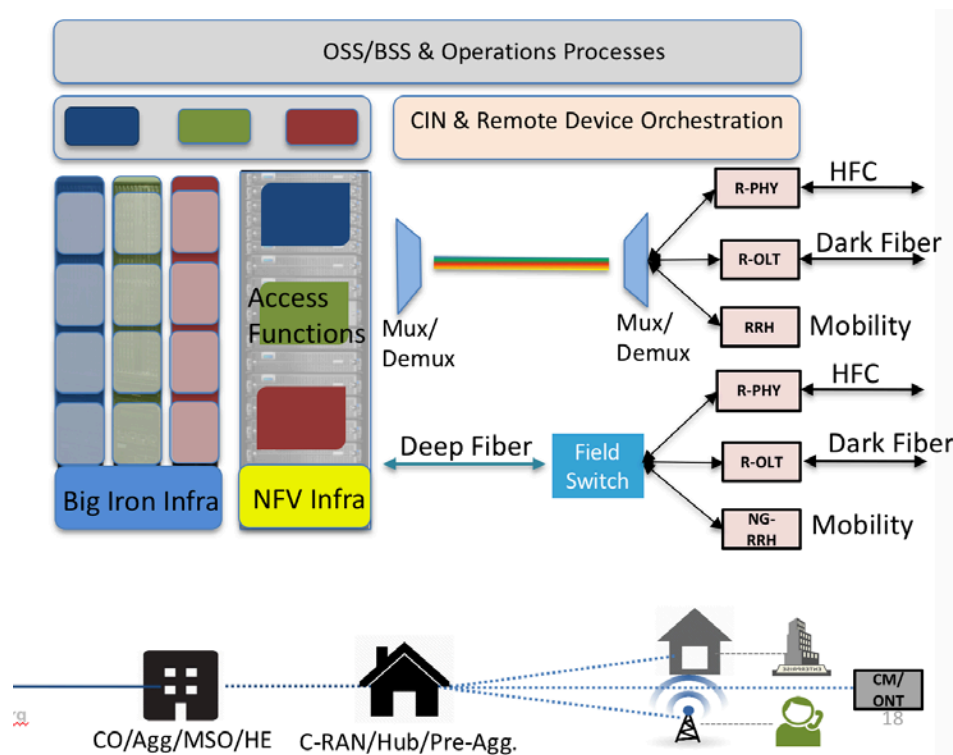


Figure 8 - Converged Multi-Access Layer

Multiple access technologies, such as Cable, Mobility and PON reside on a common virtualized infrastructure attached to the CIN.

It is not required that any operator rip and replace the current revenue generating big iron infrastructure with replace it virtual instances to avail the benefits of a converged CIN. Big iron infrastructure can continue to be deployed and will share the CIN with virtual instances. Although, virtualization provides unique flexibility in scaling an access technology or scaling across access technologies.

There are three variants of the topology with NFV infrastructure. In the first variant, the NFV infrastructure for converged multi access technologies resides in the hub along with the front haul. In the second variant, only the CIN equipment resides in the hub while the NFV infrastructure and multi access CNFs reside in the headend and in the third, the CIN and the data plane CNFs reside in the hub while the control and operations reside more centrally.

As the headend is usually a larger facility than the hub, it can aggregate more subscribers. Aggregation of more subscribers leads to better utilization of the common NFV infrastructure. The choice of variants depends on many factors such as facility power, space and access technology migration to virtualization. Consolidation criteria are unique to every site in the provider network.

Combined real and virtual instances could be managed together with an overall single layer of management software that componentizes application specific portions.

As the NFV infrastructure is common across access technologies, the management layer for the infrastructure is also common across access technologies. There is an access specific component in the management software to manage the corresponding access CNFs and the remote field devices associated with the specific access technology. Provider operations, and OSS & BSS processes extend to cover the NFV infrastructure, unifying management of the CIN, field devices and NFV infrastructure.

Figure 9 shows the complete picture of the converged multi access network by including the IP services network layer. A final convergence on NFV infrastructure is shown. This yields the highest level of convergence and best utilization of infrastructure, though even if portions of the network still have big iron in place a large portion of convergence can be met.

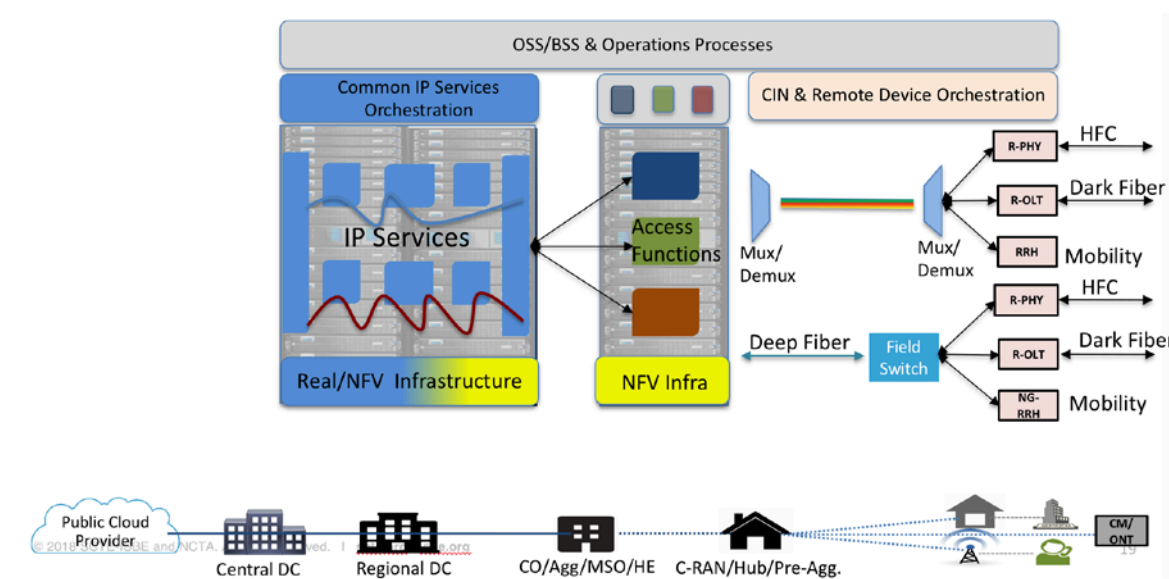


Figure 9 - The Converged Multi-Access Network

As the access layer of the network is now common for all access types and access technologies co-residing at the same location, a single common shared IP services layer can service all access types. Therefore, the entire access network becomes converged and multiple accessed. This simplifies deployments, management and reduces cost to install and operate the network.

The IP services layer could be built using real (purpose built) hardware or virtual instances or a combination of both real and virtual. The IP services layer could be co-located with access workloads or could be deeper in the network anywhere between the central data center to the head end. The deeper in the network the IP services reside, the more subscriber traffic needs to be aggregated and routed. Peering points, video caches and security all move deeper. Note that the IP services layer can be converged regardless of whether the access layer comprises big iron or virtual instances or both.

As the IP services network is common to access technologies, it can be managed as an entity agnostic of access technologies. Combined with the remaining portions of the converged access network, a single

OSS/BSS plane and a common set of operations processes can be applied to manage the entire converged multi-access network.

Conclusion

Access networks Cable, PON & Mobility are evolving because of two underlying technology trends. These are migration of the access technology to virtualization and disaggregation, and rebuilding of the CIN using deep fiber. These technology drivers allow access networks to be built in a common and converged fashion and be co-located.

Once the access technology network and the CIN are converged, the IP services layer can also be converged and can be deployed as an access agnostic network serving multiple access technologies.

As large portions of the access network are common converged and access agnostic, they can be managed and operated using common access agnostic processes enabled by a common converged layer of management software. This is a new opportunity to stream line and simplify operations while providing newer access services.

A converged multi access network creates new unprecedented opportunities for operators. It becomes easier to offer multi-access services, to collaborate and create new opportunities for co-hosting services, as well as to provide new opportunities for facility optimization.

Operators can better monetize the access network by running multiple access types using the same infrastructure.

A converged multi access network also benefits subscribers. The subscriber experience will be uniform across access types, while creating new handoff and access network redundancy opportunities. Subscribers can attain new bundled access services if operators offer multi-access bundled services.

Abbreviations

BBU	Base Band Unit
BNG	Broadband Network Gateway
EPC	Enhanced packet Core
HFC	Hybrid Fiber Coax
CCAP	Converged Cable Access Platform
CIN	Converged Interconnect Network
CM	Cable Modem
CMTS	Cable Modem Termination System
CPF	Control Plane Function
DOCSIS	Data Over Cable Service Interface Specification
DU	Distribution Unit
DWDM	Dense Wavelength Division Multiplexing
FAR	Field Aggregated Router
Gbps	Gigabit per second

HHP	Households Passed
IP	Internet Protocol
OLT	Optical Line Terminal
ONT	Optical Network Terminal
ONU	Optical Network Unit
MSO	Multiple System Operator
PON	Passive Optical Network
UPF	User Plane Function
RRH	Remote Radio Headend
RPHY	Remote - PHY

Converging Edge Caching and Computing Power for Simultaneous Mobile and MSO Networks to Handle Latency Sensitive Services Using Co-Operative Caching

A Technical Paper prepared for SCTE•ISBE by

Sandeep Katiyar
Senior Consultant
Nokia Bell Labs Consulting
Bldg. 9A, 7th Floor, DLF Cybercity
Gurugram, India-122002
sandeep.katiyar@bell-labs-consulting.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Background	4
Caching Strategies	5
Deployment Considerations	7
High Level Business Considerations	10
Conclusion.....	12
Abbreviations	12
Acknowledgments	12
Bibliography & References.....	13

List of Figures

Title	Page Number
Figure 1 – 5G Requirements.....	4
Figure 2 – Co-operative Cache	5
Figure 3 – 5G Application Latency demand.....	6
Figure 4 – Cache Tradeoffs	7
Figure 5 – Framework for Co-operative Caching.....	9
Figure 6 – Service Chaining in Co-operative Caching.....	10
Figure 7– Mobile Traffic by Content Type.....	10
Figure 8 – Stakeholder Benefits.....	11

Introduction

Broadcast and demand-based content networks have been pushed to their limits to reduce latency and to provide faster buffering to seamlessly deliver content. From massive data centers to edge based cache servers, caching has followed Multiple System Operators (MSOs) to the cellular edges to fulfill the demand of its subscribers in delivering emerging latency-sensitive services. Mobile and wireline operators have regularly increased bandwidth to meet growing data and new interactive service demands. But bandwidth itself does not address latency challenges. Caching has been used in services such as YouTube and Netflix to reduce video content delivery and web service latency. MSO deep fiber penetration and the future migration of cable hubs to edge clouds to enable virtualized services can be mutually beneficial to wireline and mobile services by bringing better content to mobile subscribers, providing higher quality reduced latency services, and increasing revenue.

On the other hand, with changing user habits and the resulting reprioritization of mobile data over voice services, along with smart device adoption and usage of personalized and enterprise-level mobility applications, mobile network operators face significant challenges related to redesigning the backhaul to support capacity and latency requirements for 5G deployments. If we closely look to the 5G requirements as depicted in Figure 1, densification of mobile networks is required to bring 5G to full use, leading to a dependency and need for high bandwidth access networks and content caching closer to the edge.

Besides raw data management, the low latency signaling required to coordinate and manage application data flow strains the network in terms of its performance. Greater capacity, unencumbered transmission and continuous coverage are needed. To make this happen specially for growing mobile data traffic i.e., video, the mobile network deployment method needs to be changed.

This paper provides an overview of how MSOs can provide caching to reduce latency for mobile networks. We look at cache tradeoffs, define a high-level architecture and finally discuss a new business service/opportunity for MSO edge content aggregation to meet the needs of mobile services,

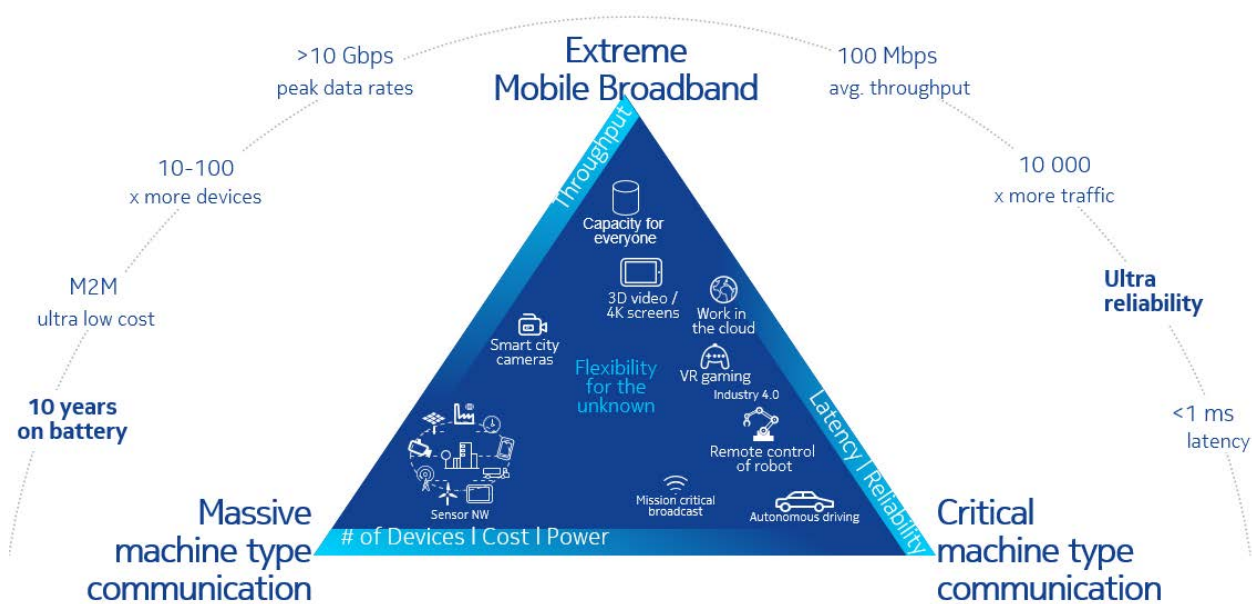


Figure 1 – 5G Requirements

Background

With explosive growth in multimedia traffic, the scalability of Over the Top (OTT) & other video services has become increasingly important. By exploiting the potential cache ability at the edge layer of Mobility and Fixed networks, the performance of multicast delivery can be improved through co-operative caching and realizing the deep reach of MSO networks. This caching technique can help minimize the average bandwidth consumption on the backhaul sides of mobile networks.

A lot of work by researchers and engineers has focused on finding effective ways to reduce duplicate content transmissions. This includes adopting intelligent caching strategies inside mobile networks and enabling edge based caches in MSO networks to access popular content from caches of nearby gateways, using selective Internet Protocol (IP) traffic offload methods. From the MSO's perspective, this also helps reduce traffic exchanged with Internet Service Providers (ISPs) and helps reduce response time required to fetch content. Both MSO and mobile networks face similar problems with respect to the content placement and its delivery, determining the size and location of each cache, and downloading to cache nodes. Co-operative caching addresses the placement issue without compromising Quality of Experience (QoE).

Video is approximately 70-80%¹ of total mobile and fixed network traffic. Given this volume and the need for caching, Co-operative caching can enable MSOs to leverage their network to position caching as a service. Such a service can address QoE and coverage aspects for capacity limited areas, helping reduce video service end-to-end latency, while reducing traffic in core and edge networks.

¹ Bell Labs Consulting traffic analysis




Figure 2 – Co-operative Cache

Caching Strategies

While the potential of co-operative caching within Mobile networks has been evaluated by several recent works [see Reference section], this paper focuses on co-operative caching applied across mobile and fixed networks by providing an overview of the challenges and possible solutions using the edge caches between mobile and fixed network users. For that we need to look at the similarities and the differences between local caching in fixed and mobile edge networks.

From a caching strategy point of view, as depicted in Figure 2, mobile caches generally are either placed at edge of mobile networks providing better QoE, low latency and high complexity. or caches are placed in a core data center with lower QoE, high latency and lower complexity as depicted in Figure 4. Local MSO caches are typically available at a distance of 10-15 km from the last mile. These caches provide low latency, higher capacity and better QoE and are suitable to help achieve real densification in terms of content availability with lower transport latency and processing. That is where the relevance of co-operative edge helps to achieve the low latency and better QoE for streaming content by handling such requests directly at the co-operative edge as described further.

For both fixed and mobile networks, the challenges for caching are similar - where should the content be placed and how it should be delivered. This challenge becomes more relevant for the mobile operator, as the user is mobile and the number of users served by a given mobile network operator access node (5G gNb) may vary with time and hence becomes difficult to find efficient caching placement and optimized cost. . Another aspect impacting cache placement is content delivery latency. As the latency gets tighter in 5G networks and as transport bottlenecks appear, the performance degradation of applications that are sensitive to it, such as video calls, voice, or gaming, result in a worse QoE. Figure 3 illustrates examples of applications and latency requirements in the network. BW-efficient 360° video and 4K video streaming are two examples where co-operative caching can be used.

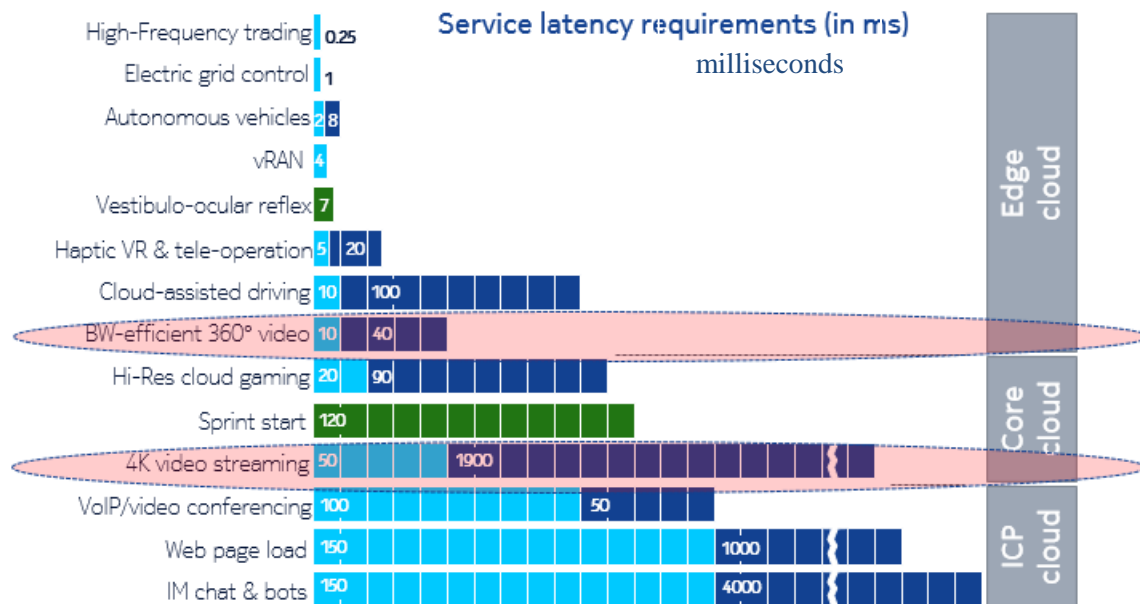


Figure 3 – 5G Application Latency demand

With less storage capacity, the amount of content that can be stored at the local cache is limited.

Processing power can also be limited, and hence it may not be efficient to run some applications from the local cache. End-to-end network complexity increases as network operators deploy, integrate and manage local caches in many locations. Resources may be needlessly duplicated if applications could be efficiently run from an alternate location.

Figure 4 illustrates three caching strategies: mobile Local Cache, Co-operative Cache, and Cloud Core Cache, based on the available latency/storage and processing capacity and distance from the base station, that can be considered by mobile operators.

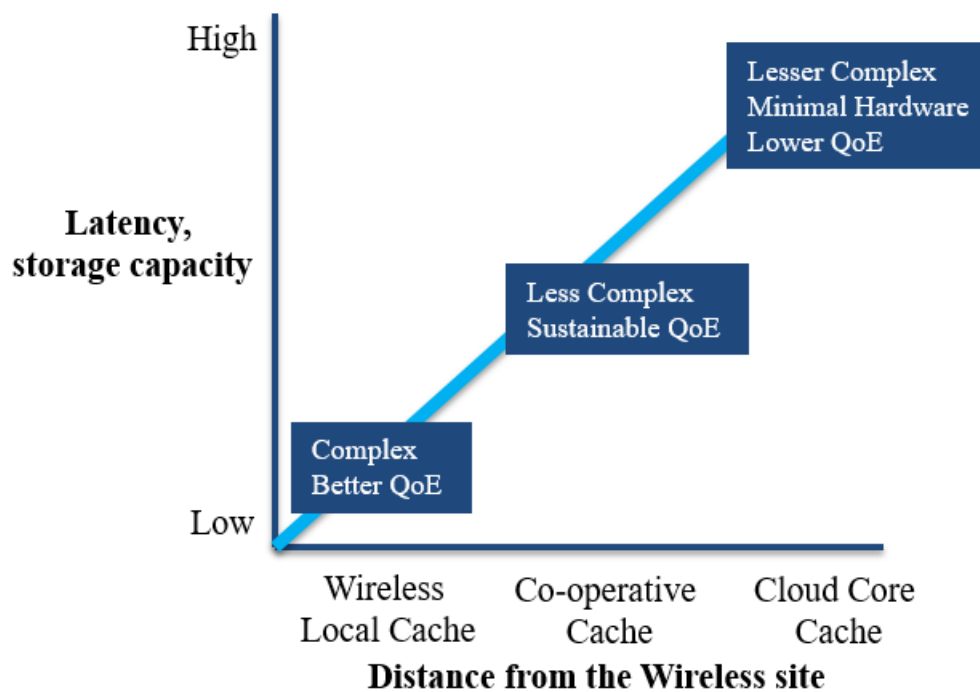


Figure 4 – Cache Tradeoffs

As described earlier, edge caching involves both content placement and content delivery. With mobile operators this becomes more challenging due to densification of the mobile cell sites in a 5G use case. Earlier studies ²show that two factors which affect a mobile cache in comparison to a wired network are:

1. **Low cache-hit probability:** When proactive and reactive caching policies designed for the Internet are not effective for caching at the node and result in insufficient utilization of caches, sharing the cache among the nodes, or redirecting the streaming requests to the co-operative cache can be used as suggested in this paper.
2. **Topology uncertainty:** Fixed networks generally have well known node topologies for subscriber connections, while in a mobility case a user request (i.e., the user connectivity to the base station for processing a request) will always be undetermined due to its mobile nature. This further complicates the determination of expected content and bandwidth. It can be overcome with the provisioning of co-operative edge at different points of networks to cater a certain % of mobile subscribers, together with a deterministic approach of caching the popular content at the MSO co-operative edge cache.

Deployment Considerations

Before moving ahead with deployment considerations, we need to understand some attributes of emerging 5G networks. First, 5G radio access is dependent upon the fiber-based fronthaul due to low latency service requirements. Second, due to massive densification of mobile sites specially in urban areas, there is a need for more capacity in terms of throughput and cache. In Figure 5, the Next Generation NodeB

²Caching at the Wireless Edge: Design Aspects, Challenges, and Future Directions. Dong Liu, Binqiang Chen, Chenyang Yang, and Andreas F. Molisch. IEEE Communications Magazine. 2016

gNBs (the 5G mobile base station) are equipped with Mobile Edge Computing (MEC) servers which can be used for local cache and deliver frequently requested content.

To understand how Co-operative caching works, let's take an example where applications like 4K streaming and virtual reality streaming initiate data requests via the gNB in the 5G network. If a request for un-cached content is initiated, the co-ordination server works with MEC to first check available content sources at the MSO cache and/or the mobile operator cache server in the core network, and the latency over the paths to these content sources. This is achieved by providing feedback to MEC about the results of microburst latency results at regular time intervals. Once the requested content is cached in the co-ordination server, the content can be delivered to the user from the local cache. If the requested content is not cached in the local gNB, but is available in a nearby gNB, the content can be delivered from there.

This also opens a new way to offload all the video based content towards the MSO edge cache, thus relieving the backhaul for other bidirectional latency sensitive services.

The coordinator server plays a key role in co-operative caching, by maintaining the state of the edge nodes, path latency information, and decides the content delivery paths. Together with the MSO edge cache, which has powerful computation and large storage, they deliver co-operative caching for MSO and mobile operators.

The proposed co-operative caching framework is built on top of the gNB's MEC framework, and the MSO Edge Cache and co-ordination server potentially running on a virtualized platform in a MSO edge node as illustrated in Figure 5. The MSO edge node also consists of a headend and an edge router. Together the MSO edge node provides rich computing resources, storage capacity, connectivity, and access to cached contents. The MSO edge node interacts with the mobile operator's gNBs to handle streaming requests via the co-ordination server. To support the coordinated approach, the mobile edge and the MSO coordination server interconnects the set of service functions dynamically so that the request packets traverse the system and get processed by each service function.

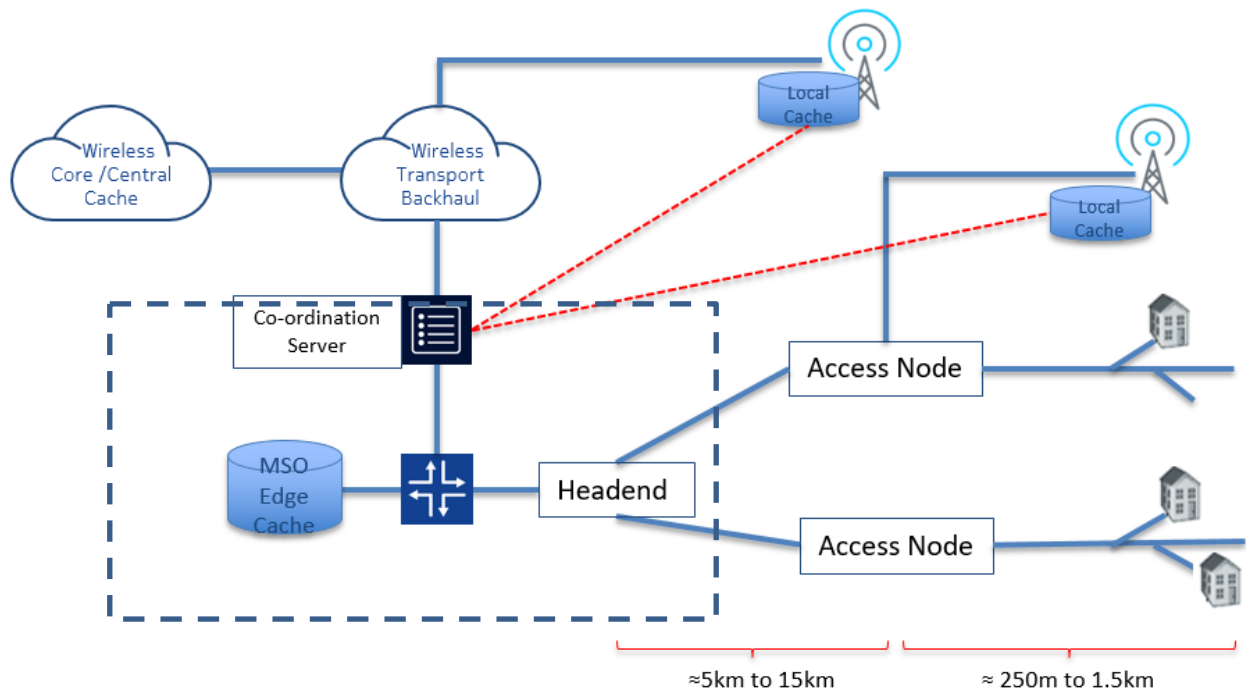


Figure 5 – Framework for Co-operative Caching

The proposed architecture thus provides two ways to implement the co-operative caching in the network:

1. Either we rely on the mobile operator backhaul as the breakout point for the streaming data, or,
2. In case the mobile operator decides to deploy a site over the MSO provided transport wherein all the streaming data is carried over the MSO network and rest of the services pass through a co-ordination server toward the mobile operator core.

The co-ordination server platform hosts the functionality required to run mobile streaming applications on top of the virtualization infrastructure. The platform hosts a set of services that can be consumed by the authorized applications. Some typical services provided by the platform include transport latency calculation function, location, and bandwidth manager. The co-ordination server platform provides visibility of the services available to the applications. If a service is provided, it can be registered in the list of services on the MEC platform so that it may redirect the request from the mobile node directly to the MSO cache server. The applications communicate with the services through well-defined application programming interfaces (APIs).

Co-operative caching between the mobile edge cache and the available fixed network MSO cache is one of the ways to support the required 5G densification while maintaining the latency, and to reduce the video traffic over the mobile backhaul. As discussed earlier, the MEC platform can provide the radio network information, latency and user location collected from the RAN (Radio Access Network) to the co-ordination server. This information is essential for the co-ordination caching server to make an optimized decision on the caching policy and resource allocation for the service type.

Figure 6 below depicts a caching framework defining the functional blocks of the coordinated edge caching system. However, to fully realize the concept, these functional blocks need to be interconnected in a sequential order, the service chaining within the system aims at interconnecting a set of network/service functions (multicast, server load balancers, HTTP header manipulation, etc.) to support

network applications. With service chaining, an operator is able to define and configure customized "service chains" in software without change at the hardware level. The service chaining helps addresses the requirement for both optimization of the network, through better utilization of resources and monetization, through the provision of services that are tailored to the service requirement context, Typically, these chains are applied to Layer 4-7 services. Here, service chaining, uses Software Defined Network (SDN) capabilities to create a chain of connected network services and connects them in a virtual chain, wherein it helps to dynamically apply or tear down single or multiple applicable services to the traffic. This capability can be used by network operators to set up suites or catalogs of connected services for use by different customers with different service and characteristics.





Figure 6 – Service Chaining in Co-operative Caching

High Level Business Considerations

Providing Cache as a Service can potentially be offered to mobile operators. An MSO as an streaming data redirector for mobile operators can upsell content and cache, also from a mobile operator point of view, based on current global traffic forecasts, we see an average of 70% traffic dominated by video on mobile networks refer to Figure below. Even if 50% of that traffic is redirected to a co-operative edge, mobile operators can reduce the backhaul traffic by 30-40% and potentially improve the QoE and support densification sites by utilizing the MSO fiber connectivity at last mile.



Source: OTT and VoD player interviews, Analysys Mason

Figure 7– Mobile Traffic by Content Type

Co-operative Edge Stakeholder	Benefits
MSO Operator	Sell Cache as a Service
	Position as OTT Traffic redirector for Mobile Operator
	Leverage unique position with the Content and application providers
Mobile Operator	Better QoE
	Better Network Resource Utilization
	Offload Streaming data to MSO for better capacity availability for bidirectional real-time services.

Figure 8 – Stakeholder Benefits

Mobile operators will continue to deploy MEC as integral parts of their network infrastructure. Although given the continuous pressure on CapEx and operational challenges, new approaches such as discussed in this paper may be welcomed.

In the co-operative cache space as depicted earlier, some new business models may arise that have a more direct and active role for MSOs, on the periphery side, and for content/application providers, on the cloud side. This in any case can be win-win situation for both from the co-operative infrastructure as provided in Figure 8, which provides the benefits which can be leveraged by stakeholders.

For example, a co-operative cache server that supports industrial applications in some localities may be better positioned with MSO networks than with mobile operator networks. The MSO may see a compelling business opportunity, and offer its services to mobile operators while a mobile operator might struggle to see a positive ROI or might not be able to assess the revenue potential. Similarly, a content or application provider may be willing to locate some of the infrastructure it needs at the co-operative edge of the network where it can serve to both and that is more effective – and potentially more cost effective – than a remote cloud location.

Smaller data analytics companies may be willing to locate processing and storage functionality at the edge in a combined environment to leverage the best of mobility and fixed access, where they do not need to own a host server but might pay only for the services they need. In this model, the MSO may deploy and pay for the initial edge hardware, but then it can monetize the investment by renting access to it to a mobile operator.

A model of this type can be mutually beneficial, in case to optimize network resources and performance.

Conclusion

This paper surveyed and provided a framework for the co-operative caching based on breaking out the streaming data traffic from the mobile edge to the fixed edge which is a one step towards integrating computing, caching and communication resources.

The issues of co-operation between the two edges and as well as some existing edge caching and computing platforms are presented. Co-operative cache edge goes beyond the centralized cloud model, which combines centralized and distributed processing, storage and control. Operators can leverage network flexibility to find the best edge location to maximize QoE and optimize network resource utilization, the main drivers for edge computing.

New business opportunities will accelerate a move to the edge, with an increased role of fixed asset owners, enterprises, and application and content providers. Traffic optimization at co-operative edge encourages a tighter co-operation of mobile operators with MSOs with clear benefits for the mobile operators.

Abbreviations

API	application program interface
gNB	next generation NodeB
HFC	hybrid fiber-coax
HTTP	hypertext transfer protocol
ICN	information centric networking
IP	internet protocol
IOT	internet of things
ISBE	International Society of Broadband Experts
ISP	internet service provider
MEC	mobile edge computing
MSO	multi service operator
OTT	over the top
QoE	quality of experience
QoS	quality of service
RAN	radio access network
ROI	return on investment
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined networking

Acknowledgments

- Martin Glapa, Partner & Bell Labs Fellow, Bell Labs Consulting, USA.
- Bill Krogfoss, Principal, Bell Labs Consulting, USA.
- Ben Tang, Principal, Bell Labs DMTS, Bell Labs Consulting, USA
- R.J Vale, Principal, Bell Labs Consulting, USA.

Bibliography & References

1. Online Edge Caching and Wireless Delivery in Fog-Aided Networks with Dynamic Content Popularity. Seyyed Mohammad reza Azimi, Osvaldo Simeone, Avik Sengupta and Ravi Tandon. IEEE Journal. 2018.
2. Modeling Operational Expenditures for Telecom Operators. Sofie Verbrugge, Sandrine Pasqualini, Fritz-Joachim Westphal, Monika Jäger, Andreas Iselt, Andreas Kirstädter, Rayane Chahine, Didier Colle, Mario Pickavet and Piet Demeester. Conference on Optical Network Design and Modeling. 2005.
3. Power at the edge. Monica Paolini, Senza Fili. 2017.
4. A Distributed Caching Architecture for Over-the-Top Content Distribution. Rui Dias*†, Adriano Fiorese†‡, Lucas Guardalben†, Susana Sargento*†. 14th annual conference on WONS. 2018
5. Cache in the Air: Exploiting Content Caching and Delivery Techniques for 5G Systems. Xiaofei Wang, Min Chen, Tarik Taleb, Adlen Ksentini, Victor C. M. Leung. IEEE Communications Magazine. February 2014.
6. Toward Smart and Cooperative Edge Caching for 5G Networks. Haitian Pang*y, Jiangchuan Liuy, Xiaoyi Fany, Lifeng Sun*. IEEE Journal on selected areas in communications. 2018.
7. Caching at the Wireless Edge: Design Aspects, Challenges, and Future Directions. Dong Liu, Binqiang Chen, Chenyang Yang, and Andreas F. Molisch. IEEE Communications Magazine. 2016.
8. A Survey on Mobile Edge Networks: Convergence of Computing, Caching and Communications. SHUO WANG¹, XING ZHANG¹, YAN ZHANG², LIN WANG¹, JUWO YANG¹, AND WENBO WANG¹. IEEE Access. 2017.
9. Edge Computing and the Role of Cellular Networks. Guenter Klas, Vodafone Group. The IEEE Computer Society. 2017.
10. Content-Exchanged Based Cooperative Caching in 5G Wireless Networks. Shu Fu, Peng Duan, and Yunjian Jia. IEEE. 2017.
11. Collaborative Edge Caching through Service Function Chaining: Architecture and Challenges. Lei Lei, Xiong Xiong, Lu Hou, and Kan Zheng. IEEE Wireless Communications. June 2018.

Critical Considerations For The Design Of A Robust And Scalable DAA Aggregation And Transport Network

A Technical Paper prepared for SCTE•ISBE by

Jon Baldry

Metro Marketing Director
Infinera

125 Finsbury Pavement, London, EC2A 1NQ
+44 7766 146 440
jon.baldry@infinera.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. DAA – A Once in a Generation Upgrade to Cable Networks.....	3
2. DAA – A Once in a Generation Set of Challenges.....	4
2.1. Challenge 1 – Secondary Hub Scalability	4
2.2. Challenge 2 – Fiber Deep	5
2.3. Challenge 3 – Limited Space and Power	5
2.4. Challenge 4 – Increased Automation and Control	5
3. Transport Network Innovation for DAA	6
3.1. Addressing Scalability of Secondary Hubs	6
3.2. Pushing DWDM “Fiber Deep”	7
3.3. Available Space and Power – The Ultimate Limiting Factor For Network Rollout	8
3.4. Advanced CORD Architectures – A New Networking Paradym	9
Conclusion.....	10
Abbreviations	10
Bibliography & References.....	10

List of Figures

Title	Page Number
Figure 1 - Generic DAA Transport and Aggregation Network	4
Figure 2 – Axes of Scalability in Next Generation Optical Networking	6
Figure 3 – Packet-Optical System With MPO-Based 10G Connectivity.....	7
Figure 4 – CORD-based Hyperscale Metro Networks	9

Introduction

The transition to distributed access architectures (DAA) and particularly Remote-PHY or Remote-MAC/PHY is undoubtedly one of the most significant architectural changes to hit the cable world. Whilst this has an enormous impact on the architecture of the last mile as it moves to an N+0 architecture there is a very significant knock-on effect to the DWDM optical network that supports this fiber-deep access network.

Cable MSOs across the globe are evaluating and planning for all aspects of DAA, including revamped DWDM-based optical infrastructure. As part of this evaluation MSOs need to consider some obvious and some less obvious impacts on the optical network, these include:

- **Scale** – DAA will drive massive bandwidth growth per home and therefore considerably high bandwidth within the aggregation and transport network. What advances in optical technology help drive down cost per bit in high-scale transport?
- **Fiber-Deep** – Pushing fiber and DWDM deeper into access networks brings additional challenges:
 - **Host independence** – Can DWDM optics deploy directly into 3rd party devices to avoid the need for additional DWDM termination hardware?
 - **Autotuneability** – Can DWDM optics help lower operational cost and rollout bottlenecks by learning their “color” from the network?
- **Limited space and power in secondary hubs** – Deploying real world networks quite often comes down to available space and power. How can the optical networking infrastructure help address this challenge?
- **Advanced CORD architectures** – The desire to move to CORD and Spine/Leaf architectures requires any Ethernet aggregation or switching to play a role in a wider Spine/Leaf architecture. How can this be achieved in modern packet-optical devices?

This paper will undertake an assessment of some of the recent trends in the optical networking and how they can be applied to address the considerations and challenges outlined above to help prepare cable MSOs for fiber-deep DAA.

Content

1. DAA – A Once in a Generation Upgrade to Cable Networks

An uninitiated person walking the corridors of the SCTE/ISBE Cable-Tec Expo in 2017 would be left in no doubt that the “distributed access architecture (DAA)”, whatever that is, was the main topic of conversation within the cable industry at that time. If they have decided to return and attend again this year then they will see the same, or possibly even more, excitement around DAA on the show floor.

DAA gives the cable industry the opportunity to defend its historic competitive advantage in residential markets and to migrate the numerous parallel networks required to support additional non-residential services into a single converged interconnect network (CIN). The migration to remote-PHY (R-PHY) or remote MAC/PHY (R-MAC/PHY) and the removal of expensive to maintain hybrid fiber coax (HFC) enables cable MSOs gives operators the opportunity to modernize networks enabling support for enhanced high-quality services while also reducing ongoing maintenance costs. The excitement around

DAA within the industry is understandable and by now hopefully it is also clear to our previously uninitiated visitor.

However, DAA isn't without its challenges, especially when looking at the optical transport portion of the new network. The generic DAA architecture pushes the edge of the digital network much closer to the end user, typically via a remote PHY device (RPD) and extends the WDM-based transport network that previously terminated in the secondary hub significantly deeper into the access fiber plant, as shown in Figure 1. This is part of wider *Fiber Deep* trend across the telecoms industry where fiber and associated transmission equipment is pushed deeper into the access plant and closer to the end user, such as fiber to the tower in wireless networks and fiber to the home/building in non-cable residential and business networks. Many of the challenges we'll discuss here for DAA also apply to the wider fiber deep trend and often the solutions used to overcome these challenges create the opportunity to converge networks into a CIN architecture.

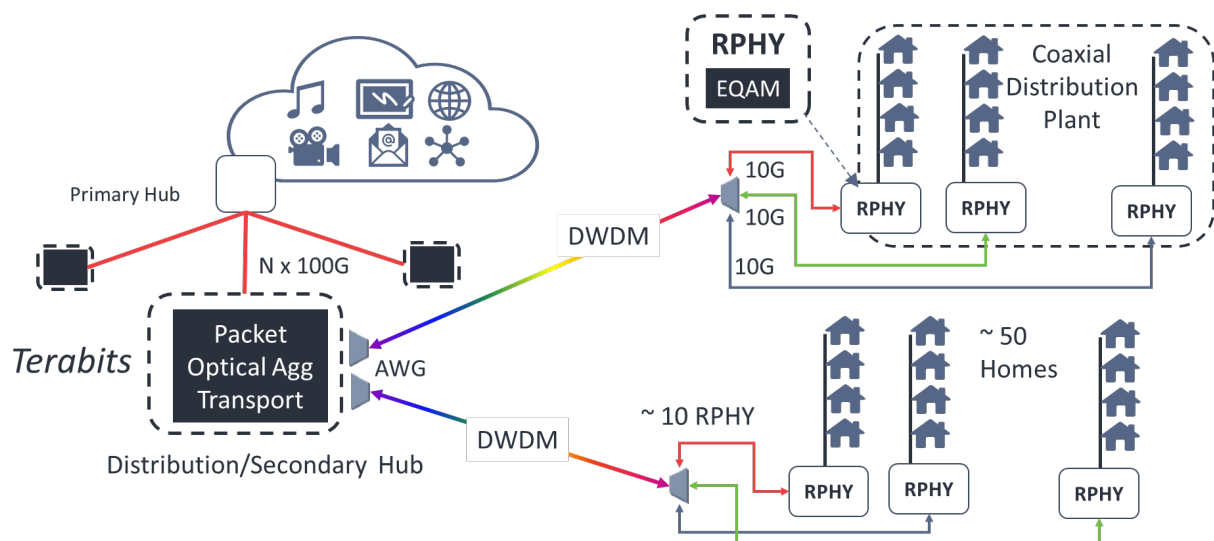


Figure 1 - Generic DAA Transport and Aggregation Network

2. DAA – A Once in a Generation Set of Challenges

The opportunity that DAA brings to cable networks globally is vast. The migration will bring a step-change to the quality and bandwidth of services that MSOs are able to offer while also modernizing the network to reduce ongoing maintenance costs. A change this big however brings new challenges that the previous network didn't need to consider or exacerbates existing challenges such as the inevitable limitations in available space and power.

2.1. Challenge 1 – Secondary Hub Scalability

DAA drive massive bandwidth into users' homes, which is the ultimate goal of DAA, and therefore will create a step change in bandwidth throughout the backhaul and transport network. DAA creates a new digital transport domain from the secondary hub to the RPD supporting 10G DWDM wavelengths per RPD. A typical secondary hub may well serve 300-400 RPDs and in some cases as many as 600-700 RPDs, generating an unusually large number of 10G circuits that require efficient aggregation into 100G+ wavelengths for terabit-level transport to the primary hubs. Each 10G RPD circuit is anticipated to be carrying around 2-3 Gbit/s of traffic on day one, allowing the 10G circuits to provide plenty of headroom

for future growth and allowing aggregation at the secondary hub to economically scale capacity to the primary hub as needed.

However, this enormous density of 10G circuits being terminated and aggregated in a single location creates another scalability challenge associated with managing the sheer volume of circuits and fibers in an economic and controlled manner.

2.2. Challenge 2 – Fiber Deep

Pushing DWDM from the secondary hub into the previously analogue optics domain closer to the end user is a key element of the DAA initiative. This enables the RPD to support the required high bandwidth 10G backhaul connection over distances of up to 60 or even 80 km. This longer reach enables the previous fiber plant to be redeployed for DAA with a chain of RPDs supporting a service area. Each RPD has a dedicated wavelength to/from the secondary hub and hardened DWDM filters can be used to support a wide variety of physical topologies.

The main challenge this brings is the operational challenge of handling the massive proliferation of DWDM end points. While each RPD is significantly simpler to install and maintain than a complete secondary hub, the previous end point of the digital DWDM network, we now must manage at least two orders of magnitude more end points. Each end point may only need a single DWDM wavelength but the largest DAA networks will ultimately have over a million RPDs. This creates a significant deployment management and logistics challenge and an additional ongoing sparing and maintenance challenge. In the ideal world we'd use automation and network intelligence to hide the complexing of DWDM from the RPD installation and maintenance with optics that behaved as if they were standard grey optics.

2.3. Challenge 3 – Limited Space and Power

As previously mentioned, to varying degrees the availability of space and power is always a concern in rolling out networks. There is always pressure in network design to optimize space and minimize power usage to keep ongoing operating costs down and to avoid expensive expansions and upgrades, if these are even possible. One advantage of DAA is that the approach includes the removal of older analogue equipment freeing up some space and power, but this can't be achieved until initial DAA equipment is deployed and ready to support a quick handover from the old network to the new on a spur by spur basis. If there is enough available rack space to support a clean initial DAA installation, then a dense solution with low power consumption is required to optimize available space and power. If there isn't enough space to support a separate clean installation on day 1, then compact and dense solutions become even more critical. Swapping equipment in and out of racks as networks evolve can lead to a fragmented and suboptimal use of the rack space and unfortunately there isn't a defrag button on the rack.

2.4. Challenge 4 – Increased Automation and Control

Operators globally are evaluating a range of approaches to the increased automation and control of the optical transport network required for DAA. Some operators are considering the central office rearchitected as a datacenter (CORD) approach that uses software defined networking (SDN) control and datacenter like operations to create an agile edge network. CORD encompasses the full DAA network from the RPD at the edge to content servers in the primary hub and any packet-optical transport network in between, requiring the transport solution to participate in the CORD network. This is a new paradigm for transport networks but one that is ultimately anticipated to become common place in many additional access network types such as 5G X-haul networks.

Other operators are looking at traditional network management and newer SDN control/orchestration systems for DAA management and control. Overall this range of requirements means that systems vendors providing DAA transport solutions need a lot of flexibility in management, control and orchestration options.

3. Transport Network Innovation for DAA

The optical networking industry is currently undergoing one of the highest, if not the highest, rates of innovation in its history. Innovation in advanced coherent optics is rapidly pushing the envelope of fiber capacity, packet-optical integration is bringing more sophisticated networking capabilities, SDN has brought a step change in multi-vendor network control to name just a few of the ongoing trends. Many of these can be applied to DAA networks to overcome the challenges outlined in this paper.

3.1. Addressing Scalability of Secondary Hubs

Recent advances in coherent optics have a clear role to play in the high capacity DAA networks to interconnect secondary hubs and primary hubs. Due to the large number of RPDs that will be aggregated in secondary hubs, networks collecting traffic from multiple secondary hubs will need to support terabits of traffic on day one and will need to scale significantly as bandwidth per user grows. Current networks support up to 10s of Terabits per second per fiber using quadrature phase shift keying (QPSK) and 16-quadrature amplitude modulation (16-QAM) modulation formats. Systems are in the process of migrating to 32-QAM and also increasing baud rates from 33 Gbaud and to 66 Gbaud. While WDM is a digital medium transmitting 1s and 0s, the engineering at the wavelength level is very much an analogue domain which involves adapting modulation format and baud rate to reach the optimum performance in terms of reach, total fiber capacity and cost per bit from both a capital and operational expenditure perspective.

Industry research is pushing these capabilities higher to 1024-QAM and 100 Gbaud [1][2], giving the potential to push fiber capacity closer and closer to the Shannon Limit in the future. Key to achieving these advances economically over useable distances is the use of techniques such as Nyquist subcarriers and advanced constellation shaping. These techniques are also used by some vendors to optimize performance in today's systems and can enable systems to outperform competitive solutions that have higher headline figures for modulation format and baud rate.

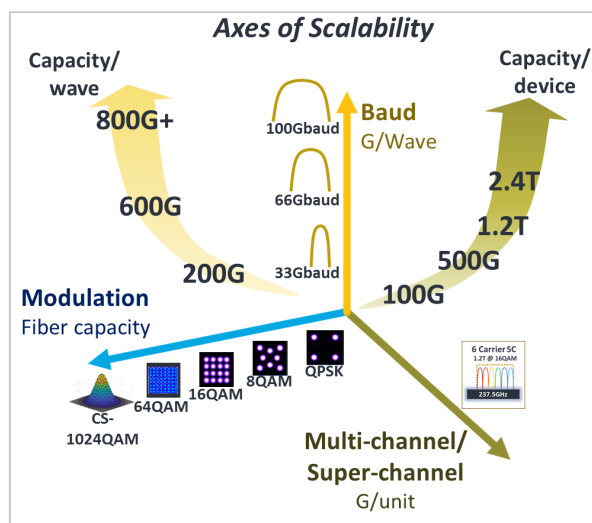


Figure 2 – Axes of Scalability in Next Generation Optical Networking

However, earlier in the paper we discussed the challenge of managing the scalability needed within the secondary hub itself. As DAA rollouts continue in the access network the secondary hubs will need to terminate hundreds of 10G circuits from RPDs and with the density of today's packet-optical aggregation systems this will be easily achievable within 1 rack. In fact, a single rack could support termination of over 1000 RPDs if needed to support a massive service area or in a consolidation scenario where multiple secondary hubs can be collapsed into a single location. Even putting this extreme example to one side, a secondary hub supporting 400 RPDs would need 800 fibers just to connect the individual 10G wavelengths from the DWDM filters terminating the RPD spurs and the 10G to 100G aggregation switch. These fibers dominate the fiber count within the configuration and could be a limiting factor in how easily a secondary hub can be upgraded to DAA and maintained.

One of the innovations that the optical industry has started to adopt to address this challenge is the use of multi-fiber push on (MPO) cables to drastically reduce the number of fibers needed in this part of the rack. MPO cables are already used for 100G connectivity so will be familiar to the installation crews. In this scenario each MPO cable will typically replace up to 20 individual fibers meaning the number of fibers needed for this 10G interconnect within the rack can be reduced by 95%. When all the other fibers within the rack are considered, the total number of fibers in the rack drops by 87% simplifying initial install and ongoing maintenance and allowing secondary hubs to scale smoothly as DAA is rolled out. Figure 3 shows a packet-optical system, with the side cover removed to expose individual SFP+ pluggable optics, that replaces 160 fibers with 8 MPO cables to support 80x 10G connections via the yellow cables. The unit also supports 8x 100G connections via the light blue MPO cables.



Figure 3 – Packet-Optical System With MPO-Based 10G Connectivity

3.2. Pushing DWDM “Fiber Deep”

The trend towards fiber deep networks is impacting a range of applications within metro networks, including DAA. Fiber deep impacts the network in several ways, such as requiring hardened DWDM filters and potentially hardened optical networking systems to be deployed in the access plant. The biggest impact on DAA networks is the rapid proliferation of DWDM endpoints in the network. In DAA the previous end point of the DWDM network was the secondary hub and each of these will now expand out to 300, 400, 600 or perhaps more RPDs. Obviously each RPD is significantly less complex than a

secondary hub and only requires a single 10G DWDM link, but nonetheless each location now requires DWDM capabilities and installation/maintenance crews that have some DWDM knowledge.

DWDM systems usually require some form of DWDM hardware at each end of the optical link to convert wavelengths to grey (uncolored) optics for handover to the client system. For a single wavelength site such as an RPD site it would be very advantageous if this could be removed and the DWDM optics be housed directly in the RPD. Other systems can house DWDM optics directly, but in these cases the host systems must be aware of the DWDM optics and fully manage them. For this capability to be useful in DAA it would be good to also make the optics host agnostic so that the host equipment does not need to manage the DWDM characteristics of the optics and treats them the same as grey optics. This removes the cost of any additional DWDM hardware needed to house the DWDM optics without adding any additional complexity to the RPD site.

An additional challenge of the proliferation of DWDM optics out into the depths of the access network is the complexity of installing and maintaining 100,000s of DWDM optics. Tuneable optics will allow operators to avoid the inventory, sparing and project management issues of deploying 100,000s of fixed wavelength optics out to the field but these then require some tuning by installation or maintenance teams that deal with RPD sites and may not have the training or time to also manage DWDM optics. Autotuneable WDM PON optics now offer operators a solution to this challenge by allowing remote optics to learn their color from the network. This means field staff can treat these optics in the same way as grey optics and no DWDM specific training is required and sparing costs can be simplified compared to fixed optics.

Earlier WDM PON technology was limited in terms of capacity to 1Gbit/s and to reaches of approximately 20 kilometers, making it unviable for DAA. However, recent innovations have led to higher speed 10G optics supporting up to 80 km, making these optics an option for DAA. Additionally, these new autotuneable optics are host independent enabling them to be deployed in any RPD that accepts third party pluggable optics. Operationally, these optics give the potential to drastically reduce the complexity of DAA rollouts as deployment and maintenance teams can avoid the need to deal with any DWDM specific concerns as RPDs are deployed and maintained.

3.3. Available Space and Power – The Ultimate Limiting Factor For Network Rollout

Once all the big decisions are made on overall network architecture the real-world issues of fitting this into existing network facilities come into play. In a lot, perhaps even the majority, of network deployment scenarios the question of available space and power becomes the overriding factor for the feasibility of the network deployment. As detailed earlier, advances such as intra-rack connectivity using MPO connectors can have a significant on face-plate density removing the need for multiple parallel fibers and associated “front-plate real-estate”.

Optical vendors are naturally focusing on high density solutions as a competitive advantage that has an obvious benefit to cable MSOs and other network operators. The denser the solution the easier it is to deploy, either in secondary hubs with space cleared for DAA or in the limited space that comes available as older pre-DAA systems are retired.

Power consumption often goes hand in hand with available space as a limiting factor in network rollouts. However, while the absolute power consumption is important from an ongoing cost perspective, power draw and fuse requirements can have a bigger impact. The downside of increased density is that more functionality/processing can be achieved in a smaller space and even though the power per Gbit continues

to decrease the overall power consumption can rise leading to the need for ever higher fuse ratings on power feeds. The optical industry continues to focus on the combination of high density and low power consumption/fuse ratings and cable MSOs should pay special attention to this in DAA networks.

3.4. Advanced CORD Architectures – A New Networking Paradigm

New fiber deep network architectures such as DAA or 5G mobile transport networks have created the opportunity for network operators to virtualize and optimize many components of the network. As outlined earlier this has led to a range of approaches from operators that packet-optical equipment vendors and others need to address. This includes traditional own built network management systems, own built SDN controllers/orchestrators and open interfaces for third party controllers directly to devices, especially any involved in Ethernet switching or IP routing.

As shown in figure 4, in a traditional ring/arc-based metro network all primary aggregation nodes are built to support the line speed of the ring/arc and access nodes are often daisy-chained off these. The move to a CORD network built using packet-optical spine and leaf switches breaks this linkage by using dedicated WDM wavelengths to scale each leaf switch independently. At the wavelength layer a ring-based fiber plant can be logically broken into separate rings with each leaf dual homed to a pair of spine switches. The packet optical network can then be managed by a CORD controller along with the other components of the wider network to efficiently support networks such as DAA or 5G transport.

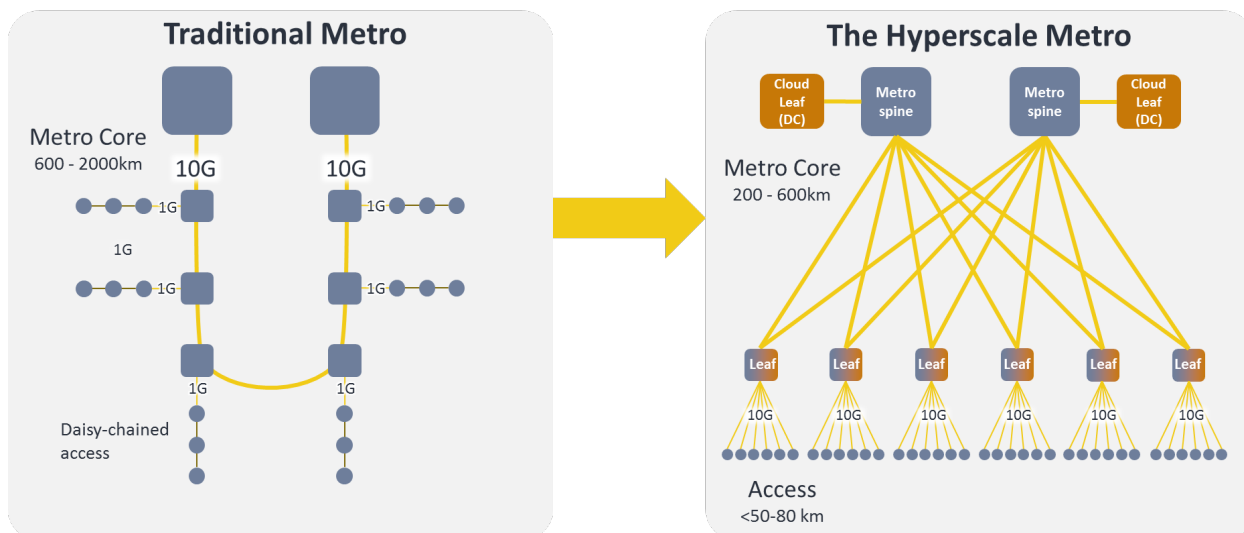


Figure 4 – CORD-based Hyperscale Metro Networks

In order to support CORD networks packet-optical transport systems will need to support a range of control options and will also need a degree of IP awareness to enable functionality such as segment routing. This should be achieved in packet-optical platforms without impacting the underlying transport performance characteristics such as low latency and SyncE and 1588v2 synchronization performance. This is especially the case in CIN networks where 5G mobile traffic will be carried in parallel to DAA traffic as the 5G traffic will need the sync and low latency performance to be optimal.

Conclusion

The industry's push to DAA is well underway and the challenges this creates for the optical network shouldn't be underestimated. At a high level the role of the transport network is simple, to transport bandwidth to/from an RPD in the most simple and economic way possible. However, DAA is perhaps unique in optical networking due to the sheer volume of RPD circuits that come together in the same place, the secondary hub.

The optical industry is innovating to help address these challenges and others found within other fiber deep applications with advances in fiber management, density, automation and control. By addressing wider fiber deep applications, such as 5G transport, this also creates the opportunity to support CIN architectures where the DAA capabilities need to expand to cover multi-service environments and differing performance requirements.

As operators embark on widescale deployments of DAA then operational and scalability issues will come to the center of attention so early visibility of these challenges and possible solutions to solve or mitigate against them is becoming increasingly important.

Abbreviations

CORD	Central Office Re-architected as a Datacenter
CIN	Converged Interconnect Network
DAA	Distributed Access Architecture
DWDM	Dense Wavelength Division Multiplexing
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase Shift Keying
R-MAC/PHY	Remote MAC/PHY
RPD	Remote PHY Device
R-PHY	Remote PHY
SCTE	Society of Cable Telecommunications Engineers
SDN	Software Defined Networking
WDM	Wavelength Division Multiplexing

Bibliography & References

- [1] *Constellation Shaped 66 GBd DP-1024QAM Transceiver with 400 km Transmission over Standard SMF*, Robert Maher, Kevin Croussore, Matthias Lauermann, Ryan Going, Xian Xu and Jeff Rahn [<https://www.infinera.com/wp-content/uploads/Infinera-tp-Constellation-Shaped-66-GBd-DP-1024QAM-Transceiver.pdf>]
- [2] *Multi-channel InP-based Coherent PICs with Hybrid Integrated SiGe Electronics Operating up to 100GBd,32QAM*, R. Going, M. Lauermann, R. Maher, H.Tsai, M. Lu, N. Kim, S. Corzine, P. Studenkov, J. Summers, A.Hosseini, J. Zhang, B. Behnia, J. Tang, S. Buggaveeti, T. Vallaitis, J. Osenbach, M. Kuntz, X. Xu, K. Croussore, V. Lal, P. Evans, J. Rahn, T. Butrie, A. Karanicolas, K.-T. Wu, M. Mitchell, M. Ziari, D.Welch and F. Kish [<https://www.infinera.com/wp-content/uploads/Infinera-tp-Multi-channel-InP-based-Coherent-PICs.pdf>]

Delta-Sigma Modulation for Next Generation Fronthaul Interface

A Technical Paper prepared for SCTE•ISBE by

Jing Wang, Ph.D.

Lead Architect, Core Innovation
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
j.wang@cablelabs.com

ZhenSheng (Steve) Jia, Ph.D.

Distinguished Technologist, Wired Technologies
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
s.jia@cablelabs.com

Luis Alberto Campos, Ph.D.

Fellow, Core Innovation
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
a.campos@cablelabs.com

Curtis Knittle, Ph.D.

Vice President, Wired Technologies
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
c.knittle@cablelabs.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
1. Challenges to C-RAN and CPRI	4
2. Motivation	4
3. State-of-the-Art.....	5
NG-RAN and Function Split Options.....	5
1. Evolution of RAN	5
2. Function Split Options	6
3. Comparison of Options 6, 7, 8, and 9	9
Experimental Demonstration.....	10
1. Operation Principles of Delta-Sigma Modulation	10
2. Experimental Setup	11
3. Experimental Results	13
4. Discussion	15
Conclusion.....	18
Appendix	19
1. State-of-the-Art of Delta-Sigma Modulator.....	19
Abbreviations	20
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 - Evolution of radio access network (RAN). (a) 3G RAN. (b) 4G cloud/centralized-RAN (C-RAN). (c) 5G next generation-RAN (NG-RAN).	6
Figure 2 - Function split options. (a) A complete list. (b) C-RAN with option 8 (CPRI) split. (b-d) NG-RAN with option 2 as HLS, and option 6 (MAC-PHY), 7 (high-low PHY), and 9 (high-low RF) as LLS.	7
Figure 3 - Detailed function block diagram of the PHY and RF layers.	7
Figure 4 - Architectures of analog RF transmitter (a) and digital RF transmitters (b, c).	8
Figure 5 - Architecture comparison of different low layer split (LLS) options, including option 6 (a), 7 (b), 8 (c), analog fronthaul (d), and option 9 (e).	9
Figure 6 - Operation principles of Nyquist ADC. (a) Each signal is digitized at baseband. (b) Input analog signal. (c) Nyquist sampling. (d) Multi-bit quantization.....	10
Figure 7 - Operation principles of bandpass delta-sigma modulation. (a) Oversampling. (b) Noise shaping. (c) BPF. (d) Cascaded-resonator feedforward structure.	11
Figure 8 - Experimental setup. (a) Xilinx Virtex-7 VX485T FPGA on VC707 development board with 4DSP FMC170 ADC. (b) 32-Pipeline architecture. (c) Optical testbed.....	12
Figure 9 - Experimental results of Case I. (a) Electrical spectra of input analog signal, OOK signal after delta-sigma modulation, and retrieved analog signal after BPF. (b) EVM vs received optical power. (c) Received constellation after 30-km fiber.	13

Figure 10 - Experimental results of Case II. (a) Electrical spectra of input analog signal, OOK signal after delta-sigma modulation, and retrieved analog signal after BPF. (b) EVMs vs received optical power. (c, d) Received constellations after 30-km fiber.	14
Figure 11 - Experimental results of Case III. (a) Electrical spectra of 10 LTE carriers. (b) EVMs of 10 LTE carriers.	15
Figure 12 - Experimental results of Case IV. (a) Electrical spectra of 14 LTE carriers. (b) EVMs of 14 LTE carriers.	15
Figure 13 - Comparison of bandwidth/bit efficiencies of CPRI, CPRI compression, and delta-sigma modulation.	17
Figure 14 - State-of-the-art of delta-sigma modulator for all-digital RF transmitter.	20

List of Tables

Title	Page Number
Table 1 - OFDM parameters of 4G-LTE and 5G-NR Signals Used in the Experiments.....	12
Table 2 - EVM Requirements from 3GPP TS 36.104 V15.2.0 [64]	13
Table 3 - Resource Utilization of Xilinx Virtex-7 VX485T FPGA.....	15
Table 4 – Comparison of Bandwidth/Bit Efficiencies of CPRI, CPRI-Compression, and Delta-Sigma Modulation	16
Table 5 – Comparison of Various Low Layer Split (LLS) Options	17
Table 6 - State-of-the-Art of All-Digital Transmitter based on Delta-Sigma Modulation	19

Introduction

The emerging video-intensive and bandwidth-consuming services, e.g., virtual reality, augmented reality, immersive applications, are driving the explosive growth of mobile data traffic [1-3], making radio access networks (RAN) the bottleneck of user experience.

1. Challenges to C-RAN and CPRI

During the 4G era, to enhance the capacity, coverage, and flexibility of mobile data networks, centralized/cloud-RAN (C-RAN) was proposed [4] to separate the baseband processing functions from base stations (BS) at cell sites, and consolidate them into a centralized baseband unit (BBU) pool, which not only simplifies each BS to a remote radio head (RRH), but also enables the radio coordination among multiple cells [5-8]. In this way, C-RAN architecture is divided into two segments, i.e., backhaul from 4G evolved packet core (EPC) to BBUs and fronthaul from BBUs to RRHs, and common public radio interface (CPRI) proposed by the CPRI cooperation, including Ericsson, Huawei, Nokia, and NEC, was adopted as the fronthaul interface [9].

However, it was quickly realized that, CPRI, as a digitization interface developed for narrowband radio access technologies (RATs), e.g., UMTS (CPRI version 1 and 2), WiMAX (v3), LTE (v4), and GSM (v5) [9], suffers from limited scalability due to its low spectral efficiency and requires tremendous data traffic in the fronthaul network segment. Moreover, CPRI features constant fronthaul data rate, which is independent to the actual mobile traffic, but scales with the antenna number, and therefore, cannot support statistical multiplexing of multiple traffic flows. All these features make CPRI the bottleneck of C-RAN, especially for the massive MIMO and large-scale carrier aggregation applications.

To circumvent the CPRI bottleneck, three strategies were developed, including the analog fronthaul, CPRI compression, and new function splits or next generation fronthaul interface (NGFI). The analog fronthaul technique transmits mobile signals in their analog waveforms using radio-over-fiber (RoF) links [10, 11]. It features high spectral efficiency, simple, low-cost system implementations, but is susceptible to nonlinear and noise impairments [12-14]. The CPRI compression solutions rely on the existing CPRI interface but manage to reduce the fronthaul data rate via compression algorithms [15-17] or nonlinear quantization techniques [18-20]. It requires additional hardware complexity and cost on both sides of BBU and RRH.

By rethinking the RAN architecture and reorganizing its function distribution [21], the next generation RAN (NG-RAN) architecture is proposed with new function split options other than CPRI [22-24]. These new function split options include option 6 and 7 proposed by the 3rd Generation Partnership Project (3GPP) telecommunications standard [25-27], and Ethernet CPRI (eCPRI) specification proposed by the CPRI cooperation [28, 29].

2. Motivation

Both CPRI compression and new function split solutions require a complete RF layer implemented in the analog domain at each remote cell site, which inevitably increases the hardware complexity and cost of each small cell and hinder the wide deployment of small cells in the 5G era. Different from the option 6, 7 or eCPRI, which moves the split point away from the PHY-RF layer interface to a higher level, we propose a new split option⁹, which lowers the split point into the RF layer. Compared with CPRI, it saves fronthaul data traffic by 50-75%; compared with other new function split options, such as 6, 7, eCPRI, it maintains the centralized architecture and significantly reduce the cost and complexity of remote cell sites, and thus facilitates small cell deployment in the 5G era.

The proposed option 9 function split implements all-digital RF transceiver by exploiting delta-sigma modulation, where the RF layer is split into high-RF layer centralized in the central unit and low-RF layer distributed in the remote cell site. It not only improves the spectral efficiency compared with CPRI, but also eliminates the need of analog RF devices, such as digital-to-analog converter (DAC), local oscillator (LO) and mixer at the remote cell site, making simple, low-cost, and energy-efficient small cell possible.

Meanwhile, the vision of software defined radio (SDR) is to push the AD/DA conversion as close as possible to the antenna, so that both baseband and RF processing are carried out into the digital domain for enhanced flexibility and compatibility to multiple radio access technologies (multi-RATs) with different PHY layer specifications. SDR also enables dynamic reconfiguration of function split, since 5G scenarios can have drastically different requirements in terms of data rate and latency, e.g., enhanced mobile broadband (eMBB), ultra-reliable low latency communication (uRLLC), and massive machine type communication (mMTC), which can significantly benefit from reconfigurable function split.

3. State-of-the-Art

As a cornerstone of SDR, all-digital RF transceiver based on delta-sigma modulation has attracted intensive research interest due to its low cost and flexibility to accommodate multi-RAT operations. Both transmitter [30-52] and receiver [53-59] designs have been reported, and various delta-sigma modulators, including lowpass [30, 32-34, 36-39, 41-45, 47, 48], bandpass [31, 35], and multiband [40, 49-52] designs have been demonstrated using either FPGA or CMOS. A state-of-the-art of delta-sigma modulation is presented in the appendix. In this work, we present a fourth-order bandpass delta-sigma modulator, which has the highest sampling rate and widest reported signal bandwidth for fourth-order modulation.

In [60, 61], we first proposed to use delta-sigma modulation to replace CPRI to improve the spectral efficiency of fronthaul, and successfully improve the fronthaul spectral efficiency by four times. In [62, 63], we first proposed that delta-sigma modulation can be used for N+0 fiber deep migration and transmitted 20 data over cable service interface specification (DOCSIS) 3.1 channels using delta-sigma modulation via a single-wavelength 128 Gb/s coherent optics link. All the modulators in these early works, however, were realized by offline processing, and so far, there is no real-time demonstration of delta-sigma modulation for NGFI application.

In this paper, we propose a new function split option 9 for NGFI enabled by delta-sigma modulation, and for the first time, a real-time FPGA-based delta-sigma modulator is demonstrated. The proposed option 9 split not only improves the fronthaul spectral efficiency, but also simplifies the small cell design and reduce the cost of small cell deployment in dense urban areas. Furthermore, all-digital RF transceivers based on delta-sigma modulation enables SDR and mobile network virtualization, which enhances the compatibility of NG-RAN with multiple RATs, including 4G-LTE, Wi-Fi, and 5G-NR, etc.

NG-RAN and Function Split Options

1. Evolution of RAN

Figure 1 shows the evolution of RAN from 3G, 4G toward 5G. In the 3G era, both baseband and RF processing are carried out in the all-in-one BS, which is distributed at each cell site, as shown in Fig. 1(a). After RF processing, the mobile signals are fed to the antennas via coaxial cables due to the short distance between the BS and antenna. In the 4G era (Fig. 1b), C-RAN architecture was proposed to separate the baseband processing functions from the BS, and consolidates them into a centralized BBU pool, so each

BS is simplified to a RRH. Since the fiber distance between the BBU and the RRH is extended to tens of kilometers, mobile signals are transmitted over digital fiber links via CPRI interface.

In the 5G era (Fig. 1c), to address the CPRI bottleneck, NG-RAN architecture is proposed with additional function split, with the baseband functions originally from the BBUs of C-RAN are now distributed into the central units (CU) and distributed units (DU). The NG-RAN architecture is thus divided into three segments, i.e., backhaul from the mobile edge computing (MEC) to CU, midhaul from the CU to the DU, and fronthaul (NGFI) from the DU to the remote radio unit (RRU), and there are two function split interfaces, high layer split (HLS) between the CU and the DU, and low layer split (LLS) between the DU and the RRU. For the HLS, option 2 has been adopted by 3GPP as a standard; whereas for the LLS, there is still debate among several different candidates, including option 6, 7 proposed by 3GPP, and eCPRI specification proposed by CPRI cooperation.




Figure 1 - Evolution of radio access network (RAN). (a) 3G RAN. (b) 4G cloud/centralized-RAN (C-RAN). (c) 5G next generation-RAN (NG-RAN).

2. Function Split Options

Figure 2 shows a comprehensive summary of function split options. Fig. 2(a) presents the block diagram of functions in different layers, including radio resource control (RRC), packet data convergence protocol (PDCP), radio link control (RLC), media access control (MAC), physical (PHY), and RF layers [23-26]. Function split options proposed by 3GPP are labeled in black, and options from eCPRI specification are labeled in blue. For the HLS, option 2 between the PDCP and RLC layers has been adopted by 3GPP as a standard; whereas for the LLS, there is still debate among several candidates, including option 6 (MAC-PHY), option 7 (high-low PHY), and eCPRI.

Fig. 2(b) shows the C-RAN architecture with option 8 (CPRI) split between the BBU and the RRH. Fig. 2(c) and (d) show the NG-RAN architectures with option 2 as the HLS between the CU and DU, and option 6 (MAC-PHY) or 7 (high-low PHY) as the LLS between the DU and the RRU. Fig. 2(e) shows LLS of option 9, where the high-RF layer is implemented in the digital domain and centralized in the DU, leaving only the low-RF layer in the RRU. In Figure 2, it should be noted that all existing LLS options, including 6, 7, 8, as well as eCPRI, all require a complete RF layer implemented in the analog domain at each remote cell site, including DAC, LO, mixer, power amplifier (PA), and bandpass filter (BPF), which increases the system complexity and cost of small cells.

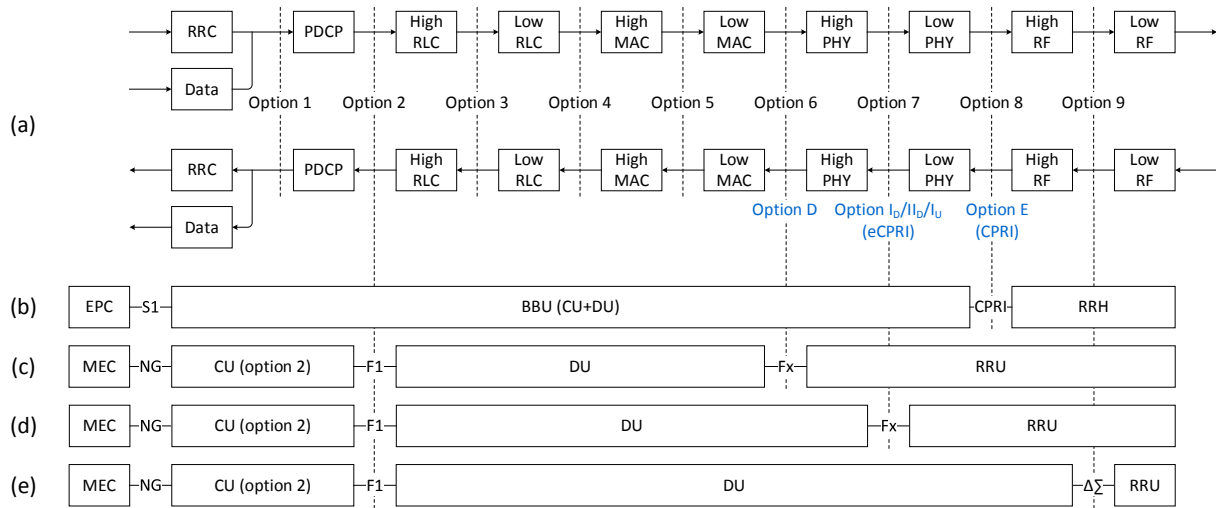


Figure 2 - Function split options. (a) A complete list. (b) C-RAN with option 8 (CPRI) split. (b-d) NG-RAN with option 2 as HLS, and option 6 (MAC-PHY), 7 (high-low PHY), and 9 (high-low RF) as LLS.

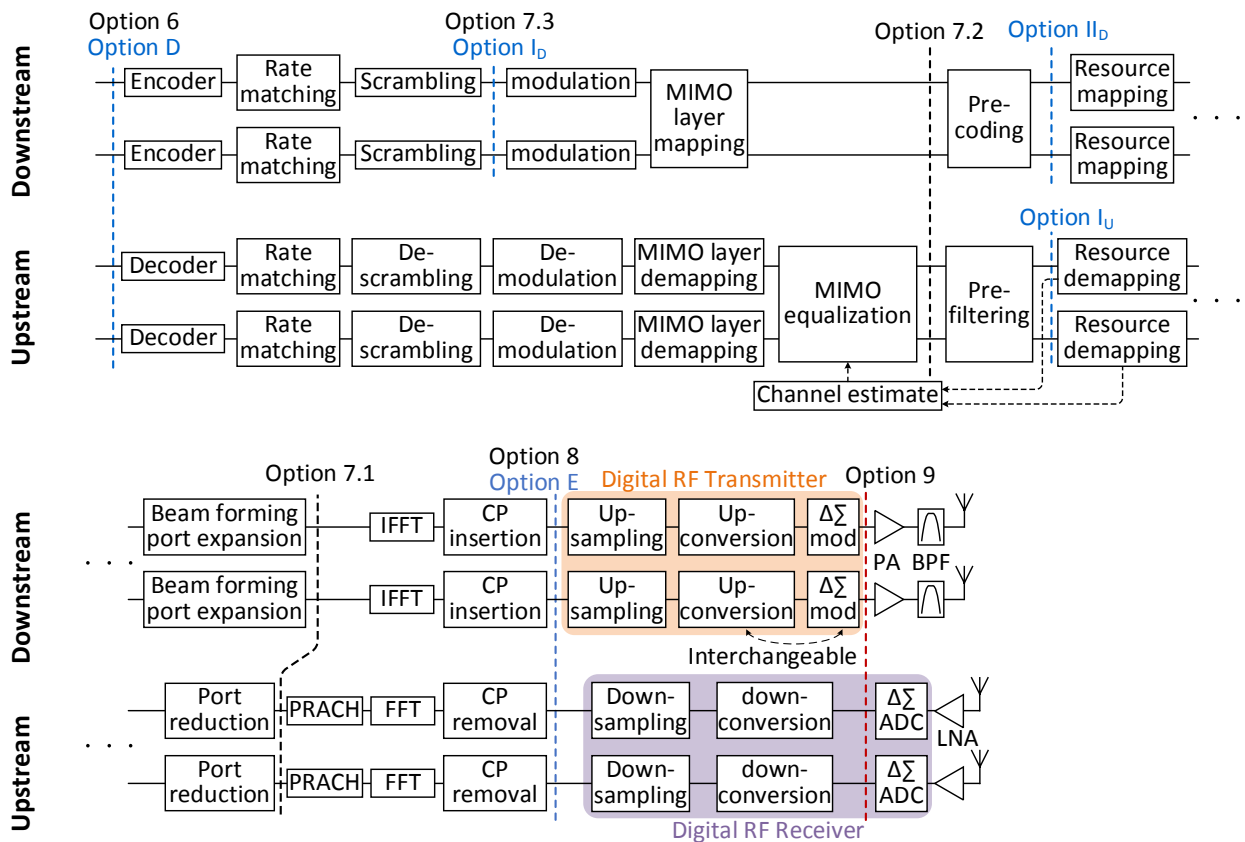


Figure 3 - Detailed function block diagram of the PHY and RF layers.

Figure 3 shows a detailed block diagram of functions in the PHY and RF layers. Function split options proposed by 3GPP are labeled in black, and options from the eCPRI specification are labeled in blue. The MAC-PHY split is defined as option 6 by 3GPP or option D in eCPRI specification; PHY-RF split is defined

as option 8 by 3GPP or option E in eCPRI specification. Within the PHY layer, both 3GPP and CPRI cooperation define three different options, 7.1, 7.2, 7.3, and I_D , I_U , respectively. Only 7.1 and 7.2 are bi-directional; all the rest are for one direction only. 7.3, I_D , I_U are for downstream, and I_U for upstream.

In this paper, we propose a new function split option 9, which lowers the split point into the RF layer. Compared with CPRI, it saves fronthaul data traffic by 50-75%; compared with other new function split options, such as 6, 7, eCPRI, it maintains the centralized architecture and significantly reduce the cost and complexity of remote cell sites, and thus facilitates small cell deployment in the 5G era.

In the RF layer of Figure 3, for downstream, an all-digital RF transmitter is used based on a bandpass delta-sigma modulator, and option 9 split takes place after the bandpass modulator, which encodes the discrete-time multibit signal into a one-bit data stream and transmits it from the DU to the RRU via digital fiber links. For upstream, a digital RF receiver based on a continuous-time delta-sigma ADC is used to digitize the received analog signal to discrete levels, and option 9 split takes place after the delta-sigma ADC, transmitting digital bits representing these discrete levels from the RRU back to the DU. With the help of delta-sigma modulation/ADC, the RF layer is implemented in the digital domain.

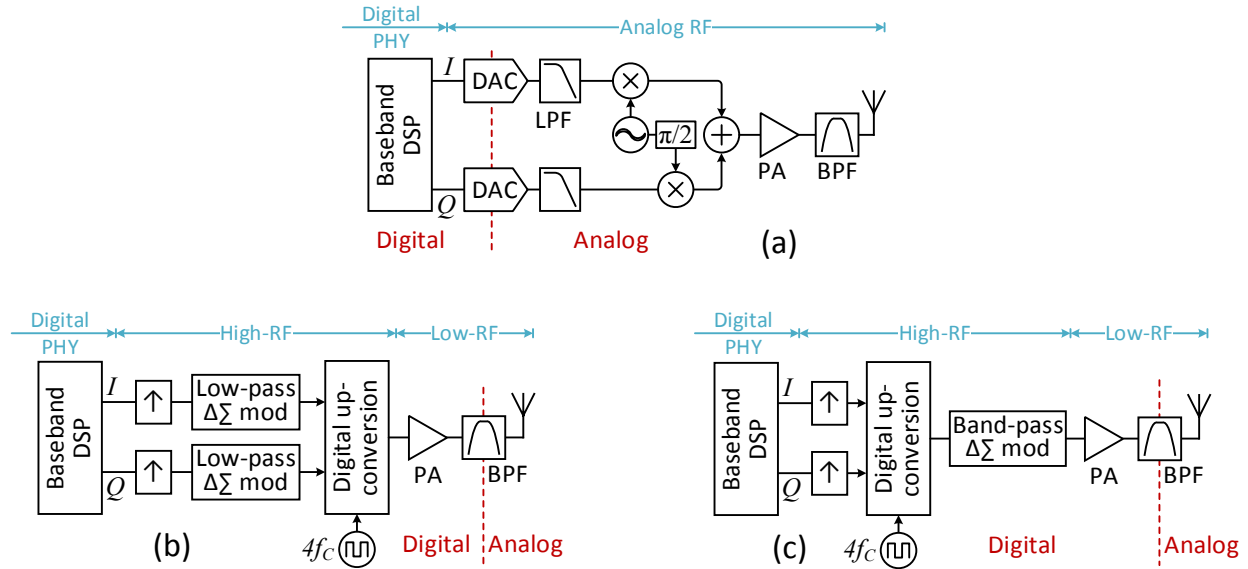


Figure 4 - Architectures of analog RF transmitter (a) and digital RF transmitters (b, c).

The architectures of digital RF transmitter are presented in Figure 4. Except for option 9, all other LLS options, including 6, 7, 8, and eCPRI, implement the RF layer in the analog domain, and require a complete RF layer at each remote cell site, which inevitably increases the system cost and complexity of each small cell. Fig. 4(a) shows an analog RF transmitter, which consists of DAC, LO, mixer, BPF, and linear PA. The DAC separates the digital processing of baseband signals from the analog processing of RF signals.

For option 9 split, on the other hand, the RF layer is implemented in the digital domain, and there is no need of analog LO, mixer, or linear PA. Fig. 4(b) and (c) shows the architectures of digital RF transmitter based on lowpass or bandpass delta-sigma modulators. In Fig. 4(b), baseband I and Q signals are first up-sampled, then encoded by two low-pass delta-sigma modulators, respectively, where the multibit baseband I/Q samples are converted to two one-bit data streams. A digital frequency up-converter then combines the I/Q bit streams and converts them to a radio frequency. The up-converted bit stream is transmitted from DU to RRU for wireless emission. In Fig. 4(c), the I and Q samples are first up-converted to a radio frequency, and then encoded by a bandpass delta-sigma modulator.

Since delta-sigma modulation utilizes noise shaping to push the quantization noise out of the signal band, in either case of Fig. 4(b) or (c), the analog mobile signal can be easily retrieved by a BPF at the RRU, which selects the desired mobile signal, while at the same time, retrieves its analog waveform by eliminating the out-of-band noise. Therefore, the conventional DAC used in Fig. 4(a) is now replaced by a simple low-cost BPF. This design also aligns with the view of digital RF transceiver to push the DAC as close as possible to the antenna, so both baseband and RF processing are carried out in the digital domain.

One advantage of all-digital RF transmitter is its flexibility and reconfigurability to different PHY layer specifications and carrier frequencies of multiple RATs. As a cornerstone to SDR, digital RF transceiver enables the virtualization of DU and RRU, making NG-RAN compatible with not only 5G-NR, but also 4G-LTE and Wi-Fi signals. Another advantage of all-digital RF transmitter is the signal fidelity. In Fig. 4(a), the analog RF signal is amplified by a PA with inevitable nonlinear impairments. But in Fig. 4(b, c), the BPF acts as a DAC, and the PA is placed before the BPF, which is still in the digital domain and dealing with digital bits. Therefore, switch-mode PAs can be used, which offers high power efficiency without nonlinear penalties. One limitation of digital RF transmitters is its high oversampling rate and high clock rate, which needs to be four times of the carrier frequency for digital up-conversion.

3. Comparison of Options 6, 7, 8, and 9

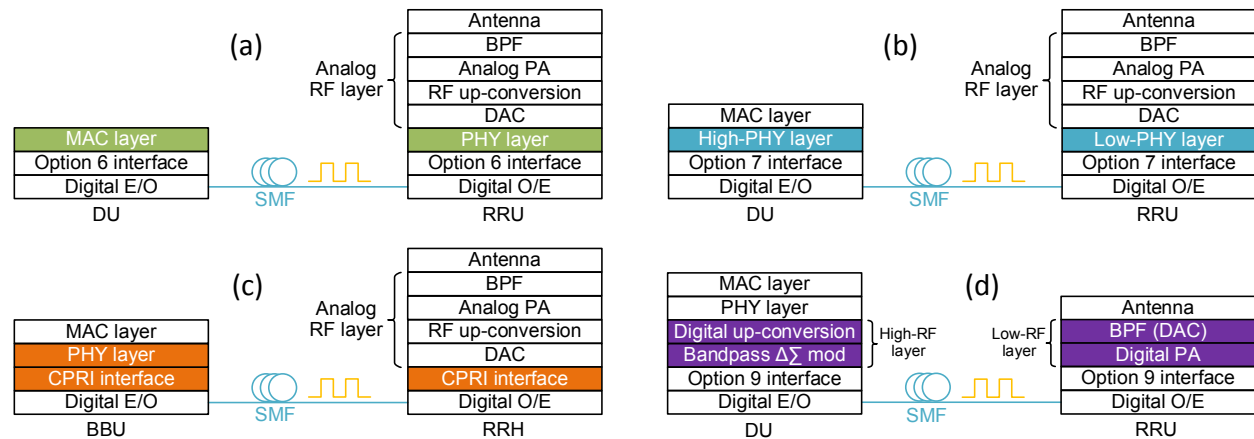


Figure 5 - Architecture comparison of different low layer split (LLS) options, including option 6 (a), 7 (b), 8 (c), analog fronthaul (d), and option 9 (e).

Figure 5 shows the architecture of different function split options, including 6, 7, 8, and 9. Fig. 5(a) shows the NGFI architecture with option 6 (MAC-PHY) split, where the MAC layer is centralized in DU, while both PHY and RF layers are distributed in the RRU. The baseband processing of the PHY layer is implemented in the digital domain; whereas the RF layer is implemented in the analog domain with DAC, LO, mixer, and PA. Since both PHY and RF layers are carried out at the RRUs, the complexity and cost of each cell site is high.

Fig. 5(b) shows the NGFI architecture with option 7 split. The high-PHY layer processing is carried out in the DU, and the rest low-PHY layer processing is implemented at the RRU. A complete RF layer is needed at each RRU, and it is implemented in the analog domain. Compared with CPRI, option 7 significantly reduces the fronthaul traffic, but also increases the cost and complexity of each cell site, which hinders the wide deployment of small cells.

Fig. 5(c) shows the fronthaul architecture with option 8 (CPRI) split, with the PHY layer centralized in the DU, and RF layer distributed at the RRUs. Like option 6 and 7, a complete analog RF layer is needed at

each RRU. CPRI has low spectral efficiency, requires tremendous fronthaul traffic, and has limited scalability for massive MIMO and carrier aggregation. Moreover, CPRI has a fixed chip rate (3.84 MHz), and can only accommodate UMTS (v1 and 2), WiMAX (v3), LTE (v4), and GSM (v5).

Fig. 5(d) shows the NGFI architecture with option 9 split, where both PHY and RF layers are implemented in the digital domain. PHY and high-RF layers, including digital up-conversion, delta-sigma modulation, are centralized in the DU; only low-RF layer functions, such as PA, BPF, antenna are left in the RRU. Since the BPF acts as an effective DAC, PA works in the digital domain, and switching-mode PA can be used with high power efficiency. Option 9 split enables a low-cost, DAC-free, and RF-simple design of RRUs, which significantly reduces the cost and complexity of cell site, and facilitates the wide deployment of small cells.

Since option 7 and 8 transmit digital baseband signals over NGFI, time division multiplexing (TDM) is needed to interleave the baseband I/Q components and components from different mobile signals. Time synchronization might be an issue considering the coexistence of legacy RAT and 5G-NR. On the other hand, option 9 transmits digital RF signal with I/Q components up-converted to the radio frequency, so frequency division multiplexing (FDM) can be used to accommodate multiband mobile signals.

Experimental Demonstration

1. Operation Principles of Delta-Sigma Modulation

Figure 6 and Figure 7 show the operation principles of Nyquist ADC and bandpass delta-sigma modulation, respectively. For a Nyquist ADC, each analog signal is digitized at baseband with a Nyquist sampling rate. The quantization noise is evenly distributed in the frequency domain. To reduce the quantization noise, multiple quantization bits are used for each sample, which leads to low spectral efficiency and large data rate after digitization and makes CPRI become the fronthaul bottleneck.




Figure 6 - Operation principles of Nyquist ADC. (a) Each signal is digitized at baseband. (b) Input analog signal. (c) Nyquist sampling. (d) Multi-bit quantization.

Different from Nyquist ADC, delta-sigma modulation trades quantization bit for sampling rate, using high sampling rate and only few quantization bits. After baseband processing, digital baseband signal is up-converted to radio frequency, then a bandpass delta-sigma modulation encodes the discrete-time multibit RF signal into a one-bit data stream. In Fig. 7(a), oversampling extends the Nyquist zone, so quantization noise can be spread over a wide frequency range. In Fig. 7(b), noise shaping technique pushes the quantization noise out of the signal band and separates the signal and noise in the frequency domain. After delta-sigma modulation, the signal waveform is transformed from analog to digital by adding out-of-band quantization noise. In Fig. 7(c), at RRU, a BPF filters out the desired signal, which not only eliminates the

out-of-band noise, but also retrieves the analog waveform as an effective DAC. Due to the noise shaping, the retrieved analog signal has an uneven noise floor.

In this paper, a one-bit bandpass delta-sigma modulator is implemented using a fourth-order cascaded resonator feedforward (CRFF) structure, shown in Fig. 7(d). There are four stages of feedback loops (z^{-1}), each two cascaded together to form a resonator. There is a feedback path in each resonator, g_1 , and g_2 . The outputs of four stages are feedforwarded with coefficients of a_1 , a_2 , a_3 , and a_4 to the combiner, then a one-bit quantizer acts as a comparator and outputs a one-bit (0/1) OOK signal.




Figure 7 - Operation principles of bandpass delta-sigma modulation. (a) Oversampling. (b) Noise shaping. (c) BPF. (d) Cascaded-resonator feedforward structure.

2. Experimental Setup

Using the CRFF structure in Fig. 7(d), a real-time one-bit bandpass delta-sigma modulation is demonstrated with Xilinx Virtex-7 VX485T FPGA on a VC707 development board, shown in Fig. 8(a). A FPGA mezzanine card (FMC170) from 4DSP is inserted in the high-pin count (HPC) connector on VC707 as the input ADC. The FMC170 ADC has a sampling rate of 5 GSa/s and 10 quantization bits per sample. The input analog signal is first digitized to 10 bits, then fed to the FPGA to perform delta-sigma modulation, which transforms the 10 input bits to one output bit. After delta-sigma digitization, the output 5-Gb/s OOK signal is outputted via a multi-gigabit transceiver (MGT) port on VC707. The 5-GSa/s sampling rate of FMC170 is contributed by 32 time-interleaved ADCs, each working at 156.25 MSa/s, so the FMC170 clock rate is 156.25 MHz. In each clock cycle, it outputs $32 * 10 = 320$ bits for 32 consecutive samples.

In Fig. 8(b), due to the speed limit of FPGA, a 32-pipeline architecture is designed to match the speed difference between the FPGA and FMC170. The input samples are de-serialized and sequentially filled into the first-in-first-out (FIFO) buffers in 32 pipelines. In each pipeline, once the input FIFO is filled up, delta-sigma modulation is performed, and the output bits are stored in an output FIFO. The output bits from 32 output FIFOs are serialized to a single bit stream. Since delta-sigma modulation is performed parallelly in all 32 pipelines, the operation speed of each line is relaxed to 156.25 MSa/s. Assuming each FIFO can store W samples with ΔW margin, since the input ADC has 10 bits per sample, each input FIFO has a size of $10(W + \Delta W)$ bits. The margin ΔW is allocated to each buffer for easy implementation. After delta-sigma modulation, the 10 input bits are transformed to one output bit, so the output FIFO has a size of $W + \Delta W$ bits.

Note that memoryless signal processing can be easily implemented by pipeline architecture, since the processing to each sample only depends on the current sample and has no relation with previous ones. After segmenting the input sample stream into several blocks, all blocks can be processed in parallel without performance penalty. On the other hand, delta-sigma modulation is a sequential operation with memory effect. The output bit not only depends on the current sample, but also previous ones, which makes it difficult to implement in a parallel way. There will be performance penalty to segment a continuous sample stream into several blocks, and the smaller block size is, the larger penalty will be. By making a tradeoff between performance penalty and the memory usage on FPGA, we choose a buffer size of $W = 20k$ with

margin of $\Delta W = 2K$. There have been several parallel processing techniques reported for high-speed, wide bandwidth delta-sigma modulators, including polyphase decomposition [44, 45], look-ahead time-interleaving [47, 48]. For a proof-of-concept experiment, here we only demonstrate the basic idea of pipeline processing with large buffer size. With the help of these parallel processing techniques, buffer size and processing latency can be significantly reduced.

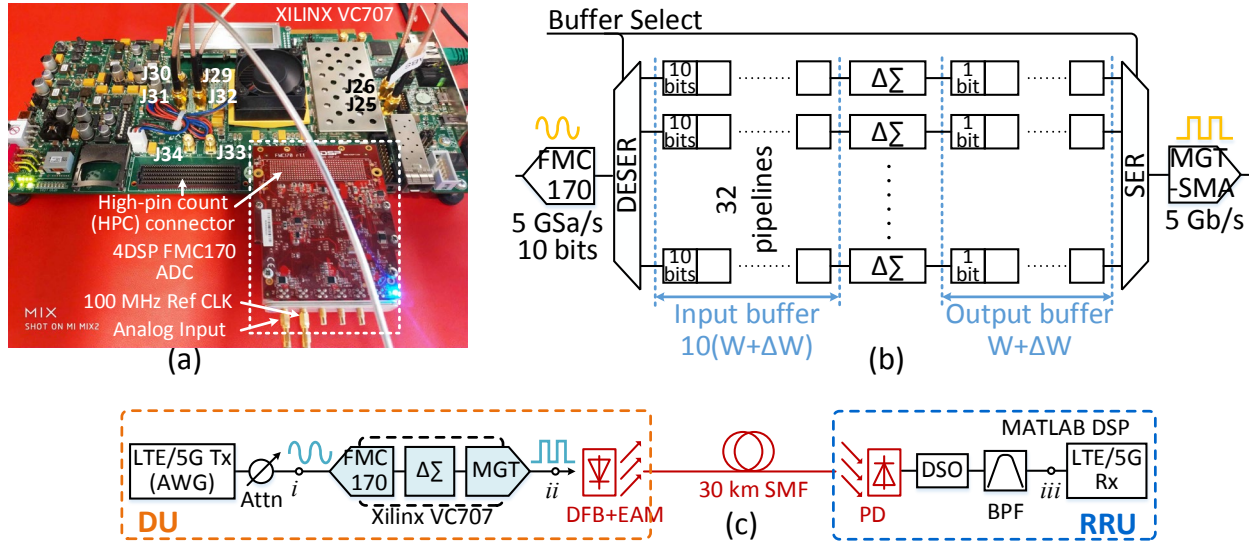


Figure 8 - Experimental setup. (a) Xilinx Virtex-7 VX485T FPGA on VC707 development board with 4DSP FMC170 ADC. (b) 32-Pipeline architecture. (c) Optical testbed.

Fig. 8(c) shows the experimental testbed. Carrier aggregated LTE/5G signals are generated by a Tektronix 7122C arbitrary Waveform Generator (AWG), then captured by the FMC170 ADC working at 5 GSa/s. The FPGA works as a one-bit bandpass delta-sigma modulator, transforming the 10 input bits to one output bit, and outputs a 5-Gb/s OOK signal. The OOK signal is delivered from DU to RRU via a digital fiber link, consisting of a 12.5 Gb/s Cyoptics DFB+EAM, 30-km single-mode fiber, and a 10 Gb/s Discovery optical receiver. 5-Gb/s error free transmission is achieved and the received OOK signal is captured by a 20 GSa/s Keysight data storage oscilloscope (DSO) MSOS804A and followed by real-time MATLAB DSP for bandpass filtering and LTE/5G receiving.

Four experimental cases are designed to verify the proposed all-digital transmitter based on delta-sigma modulation, and their OFDM parameters are listed in Table 1. 30 kHz subcarrier spacing is used for 5G signals with FFT size of 4096 and 122.88 MSa/s sampling rate. The number of active subcarriers is 3300, and the signal bandwidth of each 5G carrier is 99 MHz. The system performance is evaluated by the error vector magnitude (EVM) of received signals, and 3GPP requirements of different modulations are listed in Table 2 [64]. Note that the EVM requirement of 1024-QAM is first specified by TS36.104 V15.2.0 in 03/2018. Since this work was done earlier than that date, we use a stricter criterion of $EVM < 2\%$.

Table 1 - OFDM parameters of 4G-LTE and 5G-NR Signals Used in the Experiments

Case	Signals	Sampling rate (MSa/s)	FFT size	Subcarrier spacing (kHz)	Data subcarriers	Carrier number	Actual BW (MHz)	Modulation (QAM)
I	5G-NR	122.88	4096	30	3300	1	99	1024
II						2	198	256*2
III	4G-LTE	30.72	2048	15	1200	10	180	256*6, 1024*4
IV						14	252	1024*2, 256*4, 64*8

Table 2 - EVM Requirements from 3GPP TS 36.104 V15.2.0 [64]

Modulation	QPSK	16-QAM	64-QAM	256-QAM	1024-QAM
EVM (%)	17.5	12.5	8	3.5	2.5 (2)*

* The EVM requirement of 1024-QAM was first specified by TS 36.104 V15.2.0 in 03/2018. Since this work is done before that date, we used a stricter criterion of 2%.

In Case I, one 5G carrier with 1024-QAM and 99-MHz bandwidth is used, and the EVM performance of received signal is less than 1.25%. In case II, two 5G carriers with 256-QAM are used. Since the signal bandwidth is doubled to 198 MHz, the oversampling rate (OSR) is halved with reduced signal-to-noise ratio (SNR). So lower modulation format is used to accommodate the increased EVM, and less than 2.83% EVM is achieved for both 5G carriers. Case III and IV deal with LTE signals. In Case III, 10 LTE carriers are used with different modulations loaded on different carriers, depending on their SNR. As shown in Table 1, there are four carriers with sufficient SNR to support 1024-QAM, whereas the rest six carriers supporting 256-QAM. Similarly, in the 14 carriers in Case IV, there are two 1024-QAM, four 256-QAM, and eight 64-QAM. The reduction of modulation formats is due to the wider signal bandwidth and increased quantization noise.

3. Experimental Results




Figure 9 - Experimental results of Case I. (a) Electrical spectra of input analog signal, OOK signal after delta-sigma modulation, and retrieved analog signal after BPF. (b) EVM vs received optical power. (c) Received constellation after 30-km fiber.

The experimental results of Case I are shown in Figure 9. One 5G carrier with 1024-QAM and 99 MHz bandwidth centered at 960 MHz is generated by the AWG and converted to a 5-Gb/s OOK signal by the bandpass delta-sigma modulator on FPGA. Fig. 9(a) shows the RF spectra of input and output signals of the delta-sigma modulator. The input 5G signal at point i in the experimental setup (Fig. 8c) is labeled in blue; OOK signal after FPGA at point ii in red; the retrieved analog signal after BPF at point iii in yellow. For the retrieved analog signal, the adjacent channel leakage ratio (ACLR) is determined by the residual out-of-band noise after BPF. For easy implementation, a finite impulse response (FIR) Kaiser window filter with 40 dB out-of-band attenuation was used. By using a filter with higher out-of-band attenuation, it is not difficult to achieve the 44.2 dB ACLR requirement specified in 3GPP TS 36.141 [65]. Fig. 9(b) shows the EVM of the retrieved 5G signal as a function of the received optical power. Verilog simulation results, including floating point, fixed point, and pipeline, are also presented to show the step-by-step FPGA implementation and the performance penalty in each step. Compared with back-to-back transmission, there is no EVM penalty observed after 30-km fiber, and the received constellation is shown in Fig. 9(c).

The experimental results of Case II are shown in Figure 10. Two 5G carriers with 198-MHz total bandwidth and 256-QAM are converted to a OOK signal by the bandpass delta-sigma modulator on FPGA. Compared

with Case I, lower modulation formats are employed due to the doubled signal bandwidth and increased quantization noise. The electrical spectra of the input analog signal (point i in Fig. 8c), OOK signal (point ii), and retrieved analog signal (point iii) are presented in Fig. 10(a). EVMs of both carriers as functions of received optical power are shown in Fig. 10(b). After 30-km fiber transmission, EVMs of both carriers are less than 2.80% and 2.83%, satisfying the 3.5% requirements of 3GPP. Constellations after 30-km fiber are shown in Fig. 10(c) and (d).

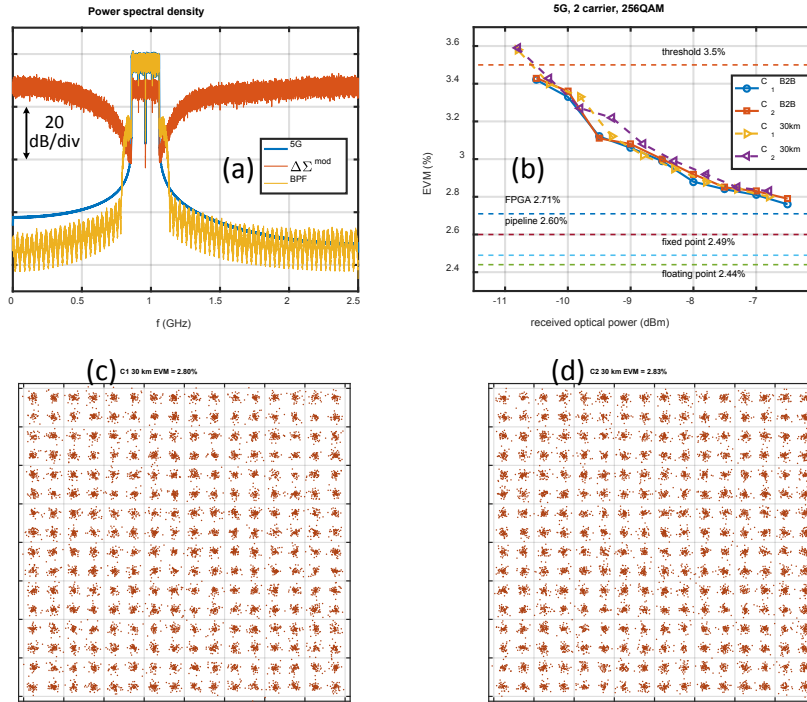


Figure 10 - Experimental results of Case II. (a) Electrical spectra of input analog signal, OOK signal after delta-sigma modulation, and retrieved analog signal after BPF. (b) EVMs vs received optical power. (c, d) Received constellations after 30-km fiber.

Figure 11 shows the experimental results of Case III, where 10 LTE carriers are used with different modulations assigned on different carriers according to 3GPP requirements. Fig. 11(a) shows the electrical spectra of input analog, OOK, and retrieved analog signals. Fig. 11(b) shows the EVMs of each LTE carrier. Within the 10 carriers, there are four carriers (2, 3, 8, 9) with EVM less than 2%, which can support modulation up to 1024-QAM; the rest six carriers (1, 4-7, 10) have EVMs less than 3.5%, and are able to support 256-QAM. The results of Case IV are shown in Figure 12. Due to the increased signal bandwidth, within the 14 carriers, there are only two carriers (3 and 12) with EVM smaller than 2%, and they can support modulation of 1024-QAM. There are four carriers (2, 4, 11, 13) with EVM less than 3.5% and used to carry 256-QAM; and the rest eight carriers (1, 5-10, 14) carry 64-QAM.

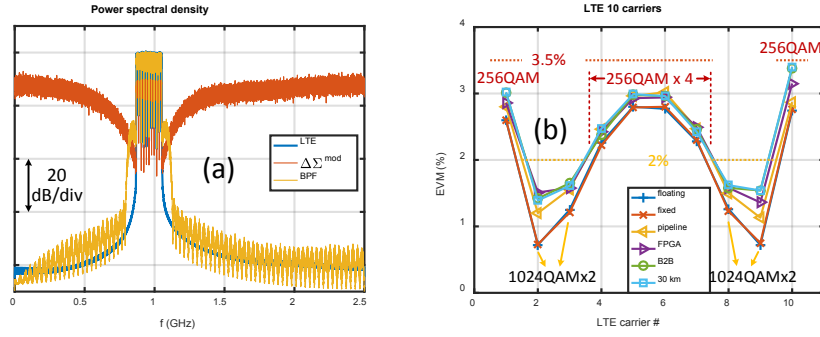


Figure 11 - Experimental results of Case III. (a) Electrical spectra of 10 LTE carriers. (b) EVMs of 10 LTE carriers.

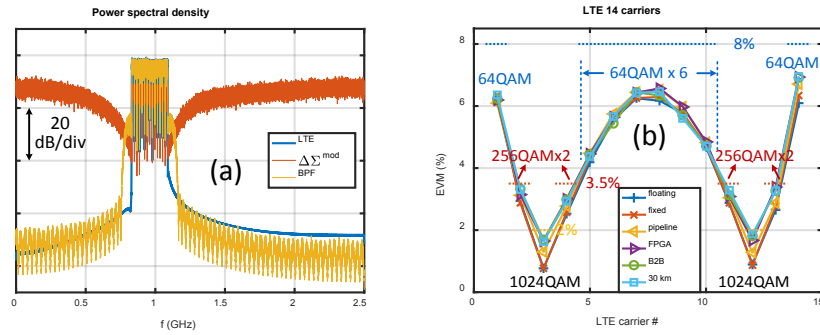


Figure 12 - Experimental results of Case IV. (a) Electrical spectra of 14 LTE carriers. (b) EVMs of 14 LTE carriers.

A summary of the resource utilization of Xilinx Virtect-7 FPGA is listed in Table 3. All four cases have similar resource usage, and the values listed are from Case II, 5G two-carrier aggregation. Note that 35.53% memory usage is due to the 22k buffer size in each pipeline. If time-interleaving technique is used, the memory usage can be significantly reduced.

Table 3 - Resource Utilization of Xilinx Virtex-7 VX485T FPGA

Resource	Utilization	Available	Utilization %
Logic cells	53362	485760	10.99%
DSP slices	64	2800	2.29%
Memory	13.18 Mb	37.08 Mb	35.53%
Transceivers	2	56	3.57%
I/O	181	700	25.86%
Max clock rate	156.25 MHz	650.20 MHz	N/A

4. Discussion

According to CPRI specification [9], a single 20-MHz LTE carrier requires $30.72 \text{ MSa/s} * 15 \text{ bits/Sa} * 2 = 921.6 \text{ Mb/s}$ fronthaul capacity without considering control word and line coding (8b/10b or 64b/66b). So CPRI can take up to 9.22 Gb/s or 12.9 Gb/s to support 10 or 14 LTE carriers, respectively. In this work, all LTE carriers are encoded by a delta-sigma modulator and transmitted through a 5-Gb/s OOK link, which saves 45.8% or 61.2% data rate compared with CPRI.

Table 4 lists a comparison in terms of spectral efficiencies of CPRI, CPRI compression, and delta-sigma modulation. Since CPRI has one control word for every 15 data words of IQ samples [9] and uses line

coding of 8b/10b or 64b/66b, for a fair composition, no control word or line coding is considered in Table 4. Since CPRI-based solutions have smaller quantization noise and higher SNR than delta-sigma modulation, it will be fair to introduce two measuring metrics, bandwidth efficiency and bit efficiency. Bandwidth efficiency is defined as the ratio between the fronthaul data rate and LTE signal BW, measuring the required fronthaul capacity per unit of BW. Bit efficiency is the ratio between fronthaul data rate and the net information rate carried by LTE signals, measuring the mapping efficiency from fronthaul traffic to real mobile traffic.

In this and our previous works [60, 61], delta-sigma modulation shows high BW efficiency, i.e., it only consumes small fronthaul capacity per unit of BW of LTE signals. On the other hand, CPRI-based solutions offer small EVM and high SNR, and therefore can support higher modulation and larger net information rate, so bit efficiency is introduced as a second metric. Although delta-sigma modulation has high bandwidth efficiency, its bit efficiency gain will not be as high as its bandwidth efficiency gain due to the high EVM and low modulations. In Table 4, it is assumed that all CPRI-based solutions carry the modulation of 1024-QAM. So far, the best bandwidth efficiency was achieved by delta-sigma modulation [60, 61], which was implemented by offline processing. The highest bit efficiency was achieved by our previous work [18, 19] using statistical compression of CPRI. Figure 13 illustrates the bandwidth and bit efficiencies of different solutions.

Table 4 – Comparison of Bandwidth/Bit Efficiencies of CPRI, CPRI-Compression, and Delta-Sigma Modulation

NGFI	CPRI-based solutions			Delta-sigma modulation				
	CPRI	Statistical Compression	Lloyd compression					
References	[9]	[18, 19]	[20]	[60]	[61]	[61]	This work	
Order	N/A			2	4	4	4	4
Sampling rate (MSa/s)	30.72	23.04	30.72	10,000	10,000	10,000	5,000	5,000
Bits	15 ¹	8	8	1	1	2	1	1
Fronthaul data rate (Gbps)	0.9216	0.36864	0.49152	10	10	20	5	5
LTE carrier #	1			32	32	32	10	14
LTE bandwidth (MHz)	18			576	576	576	180	252
Modulation	1024			64*18 16*14	256*16 64*16	1024*10 256*22	1024*4 256*6	1024*2 256*4, 64*8
Net information data rate (Gbps)	0.18			2.952	4.032	4.968	1.584	1.8
Bandwidth efficiency (MHz/Gbps) ²	19.5	48.8	36.6	57.6	57.6	28.8	36	50.4
Bandwidth efficiency gain w.r.t. CPRI	1	2.5	1.875	2.95	2.95	1.48	1.85	2.58
Bit efficiency ³	0.195	0.488	0.366	0.295	0.403	0.248	0.317	0.36
Bit efficiency gain w.r.t. CPRI	1	2.5	1.875	1.51	2.07	1.27	1.63	1.85

¹ CPRI uses 15 quantization bits and one control bit for each sample, and exploits line coding of 8b/10b or 64b/66b [9]. For a fair comparison, there is no control bit or line coding considered.

² Bandwidth efficiency = fronthaul data rate / LTE bandwidth, which measures the required fronthaul capacity per unit of bandwidth of LTE signals.

³ Bit efficiency = fronthaul data rate / net information data rate, which measures the mapping efficiency between the fronthaul traffic and mobile traffic.

Table 5 gives a comparison of various LLS options. Although the proposed option 9 splits at a lower level than option 8, it has improved bandwidth/bit efficiency and reduced fronthaul data traffic than CPRI. Compared with higher level split options 6, 7, and 8, it exploits an all-digital RF transceiver, centralizing high-RF layer at DU, replacing conventional DAC by a low-cost BPF, and eliminating the need of local oscillator and mixer at RRU. It not only makes low-cost, low-power, and small-footprint cell sites possible for small cell deployment, but also paves the road toward SDR and virtualization of DU/RRU for improved compatibility and reconfigurability among multi-RATs. Since option 9 splits deep in the RF layer, it has

very stringent latency requirement, which demands highly deterministic latency and makes it suitable for radio coordination applications.




Figure 13 - Comparison of bandwidth/bit efficiencies of CPRI, CPRI compression, and delta-sigma modulation.

Table 5 – Comparison of Various Low Layer Split (LLS) Options

3GPP / CPRI cooperation		6/D	7.3/I _D	7.2/I _D , I _U	7.1	8 (CPRI)/E	9
Architecture		Most distributed	More centralized on the right				Most centralized
RRU functions		PHY + RF layers	Low-PHY + RF layers			RF layer	Low-RF layer
RRU complexity		Highest	Medium (higher on the right)			Low	Lowest
		Complete RF layer implemented in analog domain at RRU, including DAC, LO, mixer					Only need PA and BPF
NGFI data		Baseband bits	Scrambled bits	Frequency domain I/Q samples		Time domain I/Q samples	Bits after ΔΣ modulation
Data rate		Lowest	1/10 of CPRI (higher on the right)			Highest	1/4~1/2 of CPRI [60, 61]
Data rate scalability		Traffic dependent Antenna independent, scale with MIMO			Traffic independent Antenna dependent, scale with antenna		
Ethernet compatibility	802.1CM	Future amendment	Class 2			Class 1	Future amendment
	P1914.1	Support	Support			N/A	N/A
	P1914.3	N/A	Structure-agnostic encapsulation Native RoE mapping			Structure-aware encapsulation Native RoE mapping	Structure-agnostic encapsulation
Latency requirement		Lowest	Higher latency requirement on the right				Highest

There are several IEEE standards addressing the Ethernet compatibilities of LLS options [66-68], shown in Table 5. As a collaborative effort of CPRI cooperation and IEEE 802.1 working group, IEEE 802.1CM specifies time sensitive network (TSN) profiles for fronthaul traffic over Ethernet bridged networks [66, 67]. Currently it supports two function split options, Class 1 for CPRI and Class 2 for eCPRI, and can address other split options, such as option 9, by future amendment. IEEE P1914.1, standard for packet-based fronthaul transport networks, defines the architecture and requirements of Ethernet-based mobile fronthaul traffic [68], including the Ethernet packetization of option 6 and 7. IEEE P1914.3 (previously P1904.3), standard for radio over Ethernet (RoE) encapsulations and mappings, defines 3 encapsulation methods of radio data into Ethernet packets [68], including structure-aware encapsulation, structure-agnostic encapsulation, and native RoE mapping. Structure-aware encapsulation maps CPRI frames to/from Ethernet frames with the help of knowledge of CPRI frame structure. It is optimized for CPRI and allows

CPRI to be structurally remapped to RoE. Structure-agnostic encapsulation offers a simple tunneling of radio data stream without knowledge of its frame structure, which is not restricted to CPRI and can support option 9. The third encapsulation method, native RoE mapping maps IQ payload data directly to Ethernet packets, and can support IQ samples from either time domain, such as option 8 (CPRI), or frequency domain, such as eCPRI or option 7.1, 7.2.

One major challenge to all-digital transceiver and SDR is the high processing speed. Delta-sigma modulation requires high oversampling ratio to achieve satisfying SNR/EVM performance. Moreover, digital frequency up-conversion needs a clock rate four times of the carrier frequency. To circumvent the speed limit of existing CMOS or FPGA, several parallel processing techniques have been reported, including polyphase decomposition [44, 45] and look-ahead time-interleaving [47, 48]. In this paper, for a concept-proof experiment, only basic pipeline technique is used. Given the wide frequency range of 5G from sub-1 GHz to millimeter wave, and various scenarios, e.g., eMBB, uRLLC, and mMTC, the proposed option 9 function split is expected to first find its applications in low frequency radio coordinate scenarios, such as low band 5G (T-Mobile 600 MHz). Given its highly deterministic latency, it also has high potential in uRLLC. By leveraging the low-cost, low-power, and small-footprint cell site enabled by all-digital RF transceiver, option 9 split can also be used to support low frequency narrowband IoT (NB-IoT) applications.

Conclusion

In this paper, we propose and demonstrate a new NGFI function split option 9 based on all-digital RF transceiver using delta-sigma modulation. Different from other low layer split options, e.g., 6 (MAC-PHY), 7 (high-low PHY), and 8 (CPRI), the proposed option 9 exploits the design of all-digital RF transceiver and splits functions within the RF layer, with high-RF layer centralized in DU, and low-RF layer left in RRUs. A proof-of-concept all-digital RF transmitter of LTE/5G signals is experimentally demonstrated using real-time bandpass delta-sigma modulation implemented with a Xilinx Virtex-7 FPGA. The delta-sigma modulator works at 5 GSa/s and can encode LTE/5G signals with bandwidth up to 252 MHz and modulation format up to 1024-QAM to a 5Gb/s OOK signal, which is transmitted from DU to RRU over 30-km fiber. To relax the FPGA speed requirement, a 32-pipeline architecture is designed for parallel processing. Four experimental cases are presented to validate the feasibility of proposed option 9, and 5G two-carrier aggregation and LTE 14-carrier aggregation are successfully demonstrated with the EVM performance satisfying the 3GPP requirement. A detailed comparison between CPRI, CPRI compression, and delta-sigma modulation, in terms of bandwidth and bit efficiencies is also presented.

Although it splits at a lower level than option 8, the proposed option 9 offers improved efficiency than CPRI and reduces the fronthaul traffic. Compared with higher level split option 6, 7 and 8, it exploits a centralized architecture with most RF layer functions consolidated in DU, eliminating DAC, LO, and mixer at RRU, and enables a low-cost, low-power, and small-footprint cell site for small cell deployment. Moreover, all-digital RF transceivers pave the road toward SDR and virtualized DU/RRU for multi-RAT compatibility, which has high potential in mMTC and NB-IoT applications. Given its deterministic latency, it is expected that the proposed option 9 split is more suitable for radio coordination applications than other higher level split options.

In the future, to target wider signal bandwidth and higher carrier frequency of 5G signals, more efficient time-interleaving pipeline processing architecture need to be investigated to relax the FPGA speed. Moreover, higher-order delta-sigma modulator with advanced noise shaping techniques, such as multiband operations for non-contiguous carrier aggregation or multi-RAT coexistence, will be investigated as well.

Appendix

1. State-of-the-Art of Delta-Sigma Modulator

Table 6 lists a state-of-the-art of delta-sigma modulators implemented by either CMOS or FPGA, including lowpass [30, 32-34, 36-39, 41-45, 47, 48], bandpass [31, 35], and multiband [40, 49-52] modulators. To relax the FPGA speed requirement, several time-interleaving or parallel processing architectures have been presented [38, 39, 43-45, 47, 48]. Figure 14 shows a summary of all-digital transmitters in terms of sampling rate and signal bandwidth. The right panel shows an overall summary; whereas the left panel (green circles) zooms into the early results with low sampling rate and signal bandwidth.

Table 6 - State-of-the-Art of All-Digital Transmitter based on Delta-Sigma Modulation

Reference	Sampling rate (GSa/s)	Bandwidth (MHz)	Fc (GHz)	Signal band	Implementation	Pipeline #	Signal type
[30]	0.0352	1.1	Baseband	Lowpass	CMOS 0.5 μ m	1	Continuous time Tx
[31]	0.7	< 1	0.175	Bandpass	CMOS 130 nm	1	GSM
[32]	2.625	200	5.25	Lowpass	CMOS 130 nm	1	Digital RF Tx
[33]	<3.6	10, 20	2.4-3.6	Lowpass	CMOS 90 nm	1	Digital RF Tx
[34]	5.4	5.6, 11.2, 20	2.4-2.7	Lowpass	CMOS 65 nm	1	Wi-Fi, WiMAX
[35]	2.6, 4	Up to 50	0.05-1	Bandpass	CMOS 90 nm	1	Digital RF Tx
[36]	0.05	0.25, 0.5	Baseband	Lowpass	Altera Stratix	1	OFDM, CDMA
[37]	0.045	1.25/1.23	2.45, 1	Lowpass	Altera Stratix	1	WiMAX, CDMA, EDGE
[38]	0.64, 0.8	3.84/7.68 (LTE) 4/8 (WiMAX)	2.1, 2.5	Lowpass	Altera Stratix II GX	8	WiMAX, LTE
[39]	0.025	1.6	1	Lowpass	Unknown FPGA	4	CDMA
[40]	3.9	5+5	0.8, 1.5	Dual-band	Unknown FPGA	1	Dual-band LTE
[41]	0.225	1.25+1.5	0.45, 0.9	Lowpass	Xilinx Virtex 6 HX380T on ML628	1	Dual-band WiMAX + SC-QAM ¹
[42]	0.15625	1.25+1.5	1.25, 0.78125	Lowpass	Xilinx Virtex 6 HX380T on ML628	1	Dual-band SC-64QAM + WiMAX
[43]	1/0.9 (1st/2nd order)	Up to 12.5	1, 0.9	Lowpass	Xilinx Virtex 6 HX380T on ML628	4	Single-carrier
[44]	3.2	6.1-122	1.6	Lowpass	Xilinx Virtex 6 VHX280T on ML628	16	Single-carrier
[45]	3.2	6-120	3.2	Lowpass	Xilinx Virtex UltraScale XCVU095 on VCU1287	16	Single-carrier
[46]	0.7	5	0.7	Envelope	CMOS 90 nm	1	LTE
[47]	10.4 = 0.325 * 32	20	5.2	Lowpass	Xilinx UltraScale XCVU095 on VCU108	32	Wi-Fi 802.11a
[48]	9.6 = 0.3 * 32	488	4.8	Lowpass	Xilinx UltraScale XCVU095 on VCU108	32	SC-64QAM
[49] ³	6.25	20+20	0.856, 1.45	Dual-band	Simulation + AWG ²	1	Dual-band LTE

Reference	Sampling rate (GSa/s)	Bandwidth (MHz)	Fc (GHz)	Signal band	Implementation	Pipeline #	Signal type
[50]	6.25	10+20	0.874, 1.501	Dual-band	Simulation + AWG	1	Dual-band LTE
[51]	2.15	5+10	0.244, 0.5	Dual-band	Simulation + AWG	1	Dual-band LTE
[52]	7	10+10+10	0.71, 1.75, 2.51	Triple-band	Simulation + AWG	1	Triple-band LTE
This work	5	99-252	0.96	Bandpass	Xilinx Virtex-7 VX485T on VC707	32	5G, LTE carrier aggregation

¹ SC-QAM: single-carrier quadrature amplitude modulation

² AWG: arbitrary waveform generator

³ The performance of references [30-48] are illustrated in Fig. 14. References [49-52] are not included since they are not implemented by CMOS or FPGA, but by offline processing.

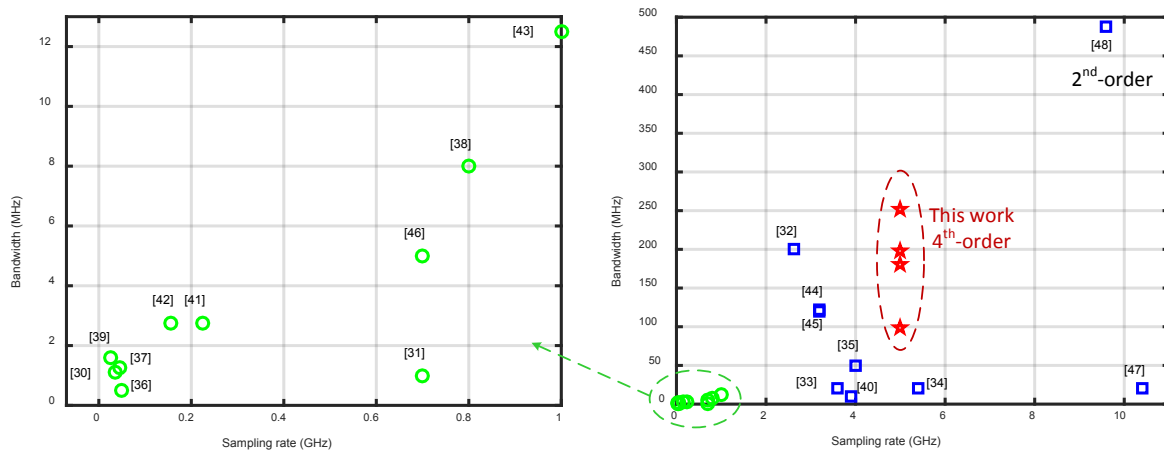


Figure 14 - State-of-the-art of delta-sigma modulator for all-digital RF transmitter.

In [48], a high speed lowpass delta-sigma modulator was demonstrated with 9.6 GSa/s sampling rate and signal bandwidth of 488 MHz. To accommodate this high speed, 32-pipeline FPGA architecture was used with 300-MHz clock rate in each line. This is the fastest delta-sigma modulator reported so far, but with only second-order modulation. In this paper, we present a fourth-order bandpass delta-sigma modulator with 5 GSa/s sampling rate and 252 MHz signal bandwidth. A 32-pipeline architecture was also employed with clock rate in each pipeline of 156.25 MHz. So far this is the fastest four-order modulator reported.

Abbreviations

ADC	analog-to-digital converter
AWG	arbitrary waveform generator
BBU	baseband unit
BPF	bandpass filter
BS	base station
CPRI	common public ration interface
C-RAN	centralized/cloud-RAN
CRFF	cascaded resonator feedforward
CU	central unit
DAC	digital-to-analog converter

DOCSIS	data over cable service interface specification
DSO	data storage oscilloscope
DU	distributed unit
eMBB	enhanced mobile broadband
EPC	ethernet packet core
EVM	error vector magnitude
FDM	frequency division multiplexing
FIFO	first-in-first-out
FMC	FPGA mezzanine card
HLS	high layer split
HPC	high-pin count
IoT	internet of things
LLS	low layer split
LO	local oscillator
MAC	media access control
mMTC	massive machine type communication
MEC	mobile edge computing
MGT	multi-gigabit transceiver
Multi-RATs	multiple radio access technologies
NB-IoT	narrowband IoT
SC-QAM	single-carrier quadrature amplitude modulation
SDR	software defined radio
PA	power amplifier
NGFI	next generation fronthaul interface
NG-RAN	next generation-radio access network
OSR	oversampling rate
PDCCP	packet data convergence protocol
PHY	physical
PON	passive optical network
RAN	radio access network
RAT	radio access technology
RF	radio frequency
RLC	radio link control
RoF	radio-over-fiber
RRC	radio resource control
RRH	remote radio head
RRU	remote radio unit
TDM	time division multiplexing
uRLLC	ultra-reliable low latency communication

Bibliography & References

- [1] C.-L. I, S. Han, Z. Xu, S. Wang, Q. Sun and Y. Chen, "New paradigm of 5G wireless internet," IEEE Journal on Selected Areas in Communications, vol. 34, no. 3, pp. 474-482, 2016.
- [2] A. Gupta and R. K. Jha, "A survey of 5G network: architecture and emerging technologies," IEEE Access, vol. 3, pp. 1206-1232, 2015.
- [3] M. Agiwal, A. Roy and N. Saxena, "Next generation 5G wireless networks: a comprehensive survey," IEEE Communications Surveys & Tutorials, vol. 18, no. 3, pp. 1617-1655, 2016.

- [4] China Mobile, "C-RAN the road towards green RAN (version 2.5)," White Paper, Oct 2011.
- [5] T. Pfeiffer, "Next generation mobile fronthaul architectures," Optical Fiber Communications Conference (OFC) 2015, paper M2J.7.
- [6] T. Pfeiffer, "Next generation mobile fronthaul and midhaul architectures," IEEE/OSA Journal of Optical Communications and Networking, vol. 7, no. 11, pp. B38-B45, 2015.
- [7] A. Checko, H.L. Christiansen, Y. Yan, L. Scolari, G. Kardaras, M.S. Berger, L. Dittmann, "Cloud RAN for Mobile Networks - A Technology Overview," IEEE Communications Surveys & Tutorials, vol.17, no.1, pp.405-426, 2015.
- [8] A. Pizzinat, P. Chanclou, F. Saliou, T. Diallo, "Things you should know about fronthaul," IEEE Journal of Lightwave Technology, vol.33, no.5, pp.1077-1083, 2015.
- [9] Common Public Radio Interface (CPRI) Specification V7.0 (2015-10-09). http://www.cpri.info/downloads/CPRI_v_7_0_2015-10-09.pdf
- [10] X. Liu, F. Effenberger, N. Chand, L. Zhou, and H. Lin, "Demonstration of bandwidth-efficient mobile fronthaul enabling seamless aggregation of 36 E-UTRA-like wireless signals in a single 1.1-GHz wavelength channel," Optical Fiber Communication Conference (OFC) 2015, paper M2J.2.
- [11] X. Liu, H. Zeng, N. Chand and F. Effenberger, "Efficient mobile fronthaul via DSP-based channel aggregation," IEEE Journal of Lightwave Technology, vol. 34, no. 6, pp. 1556-1564, 2016.
- [12] J. Wang, C. Liu, M. Zhu, A. Yi, L. Cheng, and G. K. Chang, "Investigation of data-dependent channel cross-modulation in multiband radio-over-fiber systems," IEEE Journal of Lightwave Technology, vol. 32, no. 10, pp. 1861-1871, 2014.
- [13] J. Wang, C. Liu, J. Zhang, M. Zhu, M. Xu, F. Lu, L. Cheng, and G.-K. Chang, "Nonlinear inter-band subcarrier intermodulations for multi-RAT OFDM wireless services in 5G heterogeneous mobile fronthaul networks," IEEE Journal of Lightwave Technology, vol. 34, no. 17, pp. 4089-4103, 2016.
- [14] J. Zhang, J. Wang, M. Xu, F. Lu, L. Chen, J. Yu, and G.-K. Chang, "Memory-polynomial digital pre-distortion for linearity improvement of directly-modulated multi-IF-over-fiber LTE mobile fronthaul," Optical Fiber Communication Conference (OFC) 2016, paper Tu2B.3.
- [15] B. Guo, W. Cao, A. Tao, D. Samardzija, "LTE/LTE-A signal compression on the CPRI interface," Bell Labs Technical Journal, vol.18, no.2, pp.117-133, 2013.
- [16] S. H. Park, O. Simeone, O. Sahin, S. Shamai, "Fronthaul Compression for Cloud Radio Access Networks: Signal processing advances inspired by network information theory," IEEE Signal Processing Magazine, vol.31, no.6, pp.69-79, 2014.
- [17] N. Shibata, T. Tashiro, S. Kuwano, N. Yuki, J. Terada, A. Otaka, "Mobile front-haul employing ethernet-based TDM-PON system for small cells," Optical Fiber Communications Conference (OFC) 2015, paper, M2J.1.
- [18] M. Xu, X. Liu, N. Chand, F. Effenberger and G. K. Chang, "Fast statistical estimation in highly compressed digital RoF systems for efficient 5G wireless signal delivery," Optical Fiber Communications Conference (OFC) 2017, paper M3E.7.
- [19] M. Xu, F. Lu, J. Wang, L. Cheng, D. Guidotti and G. K. Chang, "Key technologies for next-generation digital RoF mobile fronthaul with statistical data compression and multiband modulation," IEEE Journal of Lightwave Technology, vol. 35, no. 17, pp. 3671-3679, 2017.
- [20] M. Xu, Z. Jia, J. Wang, L. A. Campos, and G. Chang, "A novel data-compression technology for digital mobile fronthaul with Lloyd algorithm and differential coding," Optical Fiber Communication Conference (OFC) 2018, paper Tu2K.2.
- [21] C.-L. I, C. Rowell, S. Han, Z. Xu, G. Li and Z. Pan, "Toward green and soft: a 5G perspective," IEEE Communications Magazine, vol. 52, no. 2, pp. 66-73, 2014.
- [22] China Mobile Research Institute, "White paper of next generation fronthaul interface," version 1.0, 2015.
- [23] C.-L. I and J. Huang, "RAN revolution with NGFI (xHaul) for 5G," Optical Fiber Communications Conference (OFC) 2017, paper W1C.7.

- [24] C.-L. I, H. Li, J. Korhonen, J. Huang and L. Han, "RAN revolution with NGFI (xhaul) for 5G," *IEEE Journal of Lightwave Technology*, vol. 36, no. 2, pp. 541-550, Jan.15, 15 2018.
- [25] GSTR-TN5G, ITU-T Technical Report "Transport network support of IMT-2020/5G", Feb 2018. https://www.itu.int/dms_pub/itu-t/opb/tut/T-TUT-HOME-2018-PDF-E.pdf
- [26] 3GPP TR 38.801: "Study on new radio access technology: Radio access architecture and interfaces," V14.0.0, 2017-03. (Release 14)
- [27] 3GPP TS 38.401: "NG-RAN: Architecture description," V15.1.0, 2018-03. (Release 15)
- [28] eCPRI Specification V1.1 (2018-01-10). http://www.cpri.info/downloads/eCPRI_v_1_1_2018_01_10.pdf
- [29] Ericsson AB, Huawei, NEC, and Nokia, eCPRI presentation. http://www.cpri.info/downloads/eCPRI_Presentation_for_CPRI_Server_2018_01_03.pdf
- [30] S. Yan and E. Sanchez-Sinencio, "A continuous-time sigma-delta modulator with 88-dB dynamic range and 1.1-MHz signal bandwidth," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 1, pp. 75-86, 2004.
- [31] J. Sommarek, J. Vankka, J. Ketola, J. Lindeberg, and K. Halonen, "A digital modulator with bandpass delta-sigma modulator," *Proceedings of the 30th European Solid-State Circuits Conference 2004*, pp. 159-162.
- [32] A. Jerng and C. G. Sodini, "A wideband $\Delta\Sigma$ digital-RF modulator for high data rate transmitters," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 8, pp. 1710-1722, 2007.
- [33] P. Seddighrad, A. Ravi, M. Sajadieh, H. Lakdawala and K. Soumyanath, "A 3.6GHz, 16mW $\Sigma\Delta$ DAC for a 802.11n / 802.16e transmitter with 30dB digital power control in 90nm CMOS," *34th European Solid-State Circuits Conference (ESSCIRC) 2008*, pp. 202-205.
- [34] A. Pozsgay, T. Zounes, R. Hossain, M. Boulemnakhher, V. Knopik and S. Grange, "A fully digital 65nm CMOS transmitter for the 2.4-to-2.7GHz WiFi/WiMAX bands using 5.4 GHz $\Delta\Sigma$ RF DACs," *IEEE International Solid-State Circuits Conference (ISSCC) 2008*, pp. 360-361.
- [35] A. Frappe, A. Flament, B. Stefanelli, A. Kaiser and A. Cathelin, "An all-digital RF signal generator using high-speed $\Delta\Sigma$ modulators," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 10, pp. 2722-2732, 2009.
- [36] M. Helaoui, S. Hatami, R. Negra and F. M. Ghannouchi, "A novel architecture of delta-sigma modulator enabling all-digital multiband multistandard RF transmitters design," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 55, no. 11, pp. 1129-1133, 2008.
- [37] F. M. Ghannouchi, S. Hatami, P. Aflaki, M. Helaoui and R. Negra, "Accurate power efficiency estimation of GHz wireless delta-sigma transmitters for different classes of switching mode power amplifiers," *IEEE Transactions on Microwave Theory and Techniques*, vol. 58, no. 11, pp. 2812-2819, 2010.
- [38] M. M. Ebrahimi, M. Helaoui, and F. Ghannouchi, "Time-interleaved delta-sigma modulator for wideband digital GHz transmitter design and SDR applications," *Progress in Electromagnetics Research B*, vol. 34, pp.263–281, 2011.
- [39] S. Hatami, M. Helaoui, F. M. Ghannouchi and M. Pedram, "Single-bit pseudoparallel processing low-oversampling delta-sigma modulator suitable for SDR wireless transmitters," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 22, no. 4, pp. 922-931, 2014.
- [40] T. Maehata, K. Totani, S. Kameda and N. Suematsu, "Concurrent dual-band 1-bit digital transmitter using band-pass delta-sigma modulator," *European Microwave Conference 2013*, pp. 1523-1526.
- [41] N. V. Silva, A. S. R. Oliveira and N. B. Carvalho, "Evaluation of pulse modulators for all-digital agile transmitters," *IEEE MTT-S International Microwave Symposium (IMS) 2012*.
- [42] N. V. Silva, A. S. R. Oliveira and N. B. Carvalho, "Design and optimization of flexible and coding efficient all-digital RF transmitters," *IEEE Transactions on Microwave Theory and Techniques*, vol. 61, no. 1, pp. 625-632, 2013.

- [43] R. F. Cordeiro, A. S. R. Oliveira, J. Vieira and N. V. Silva, "Gigasample time-interleaved delta-sigma modulator for FPGA-based all-digital transmitters," 17th Euromicro Conference on Digital System Design 2014, pp. 222-227.
- [44] R. F. Cordeiro, A. S. R. Oliveira, J. Vieira and T. O. e Silva, "Wideband all-digital transmitter based on multicore DSM," IEEE MTT-S International Microwave Symposium (IMS) 2016.
- [45] D. C. Dinis, R. F. Cordeiro, A. S. R. Oliveira, J. Vieira and T. O. Silva, "Improving the performance of all-digital transmitter based on parallel delta-sigma modulators through propagation of state registers," IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS) 2017, pp. 1133-1137.
- [46] M. Tanio, S. Hori, M. Hayakawa, N. Tawa, K. Motoi and K. Kunihiro, "A linear and efficient 1-bit digital transmitter with envelope delta-sigma modulation for 700MHz LTE," IEEE MTT-S International Microwave Symposium (IMS) 2014.
- [47] M. Tanio, S. Hori, N. Tawa, T. Yamase and K. Kunihiro, "An FPGA-based all-digital transmitter with 28-GHz time-interleaved delta-sigma modulation," IEEE MTT-S International Microwave Symposium (IMS) 2016.
- [48] M. Tanio, S. Hori, N. Tawa and K. Kunihiro, "An FPGA-based all-digital transmitter with 9.6-GHz 2nd order time-interleaved delta-sigma modulation for 500-MHz bandwidth," IEEE MTT-S International Microwave Symposium (IMS) 2017, pp. 149-152.
- [49] S. Chung, R. Ma, S. Shinjo, and K. H. Teo, "Inter-band carrier aggregation digital transmitter architecture with concurrent multi-band delta-sigma modulation using out-of-band noise cancellation," IEEE MTT-S International Microwave Symposium (IMS) 2015.
- [50] S. Chung, R. Ma, K. H. Teo and K. Parsons, "Outphasing multi-level RF-PWM signals for inter-band carrier aggregation in digital transmitters," IEEE Radio and Wireless Symposium (RWS) 2015, pp. 212-214.
- [51] S. Chung, R. Ma, S. Shinjo, H. Nakamizo, K. Parsons and K. H. Teo, "Concurrent multiband digital outphasing transmitter architecture using multidimensional power coding," IEEE Transactions on Microwave Theory and Techniques, vol. 63, no. 2, pp. 598-613, 2015.
- [52] S. Chung, R. Ma, S. Shinjo, K. Yamanaka and K. H. Teo, "A concurrent triple-band digital transmitter using feedforward noise cancellation for delta-sigma modulation," 12th European Microwave Integrated Circuits Conference (EuMIC) 2017, pp. 400-403.
- [53] I. Galton, "Delta-sigma data conversion in wireless transceivers," IEEE Transactions on Microwave Theory and Techniques, vol. 50, no. 1, pp. 302-315, 2002.
- [54] M. R. Miller and C. S. Petrie, "A multibit sigma-delta ADC for multimode receivers," IEEE Journal of Solid-State Circuits, vol. 38, no. 3, pp. 475-482, 2003.
- [55] C. Wu, E. Alon and B. Nikolić, "A wideband 400 MHz-to-4 GHz direct RF-to-digital multimode $\Delta\Sigma$ receiver," IEEE Journal of Solid-State Circuits, vol. 49, no. 7, pp. 1639-1652, 2014.
- [56] M. Englund, K. B. Östman, O. Viitala, M. Kaltiokallio, K. Stadius, K. Koli, and J. Ryyänen, "A programmable 0.7–2.7 GHz direct $\Delta\Sigma$ receiver in 40 nm CMOS," IEEE Journal of Solid-State Circuits, vol. 50, no. 3, pp. 644-655, 2015.
- [57] H. Shibata, R. Schreier, W. Yang, A. Shaikh, D. Paterson, T. C. Caldwell, D. Alldred, and P. W. Lai, "A DC-to-1 GHz tunable RF $\Delta\Sigma$ ADC achieving DR = 74 dB and BW = 150 MHz at f_0 = 450 MHz using 550 mW," IEEE Journal of Solid-State Circuits, vol. 47, no. 12, pp. 2888-2897, 2012.
- [58] L. Bettini, T. Christen, T. Burger and Q. Huang, "A reconfigurable DT $\Delta\Sigma$ modulator for multi-standard 2G/3G/4G wireless receivers," IEEE Journal on Emerging and Selected Topics in Circuits and Systems, vol. 5, no. 4, pp. 525-536, 2015.
- [59] R. F. Cordeiro, A. Prata, A. S. R. Oliveira, J. M. N. Vieira and N. B. De Carvalho, "Agile all-digital RF transceiver implemented in FPGA," IEEE Transactions on Microwave Theory and Techniques, vol. 65, no. 11, pp. 4229-4240, 2017.
- [60] J. Wang, Z. Yu, K. Ying, J. Zhang, F. Lu, M. Xu, and G.-K. Chang, "Delta-sigma modulation for digital mobile fronthaul enabling carrier aggregation of 32 4G-LTE / 30 5G-FBMC signals in a

- single- λ 10-Gb/s IM-DD channel," Optical Fiber Communication Conference (OFC) 2016, paper W1H.2.
- [61] J. Wang, Z. Yu, K. Ying, J. Zhang, F. Lu, M. Xu, L. Cheng, X. Ma, and G.-K. Chang, "Digital mobile fronthaul based on delta-sigma modulation for 32 LTE carrier aggregation and FBMC signals," IEEE/OSA Journal of Optical Communications and Networking, vol. 9, no. 2, pp. A233-A244, 2017.
 - [62] J. Wang, Z. Jia, L. A. Campos, C. Knittle, and G. Chang, "Optical coherent transmission of 20x192-MHz DOCSIS 3.1 channels with 16384QAM based on delta-sigma digitization," Optical Fiber Communication Conference (OFC) 2017, paper Th1K.1.
 - [63] J. Wang, Z. Jia, L. A. Campos, L. Cheng, C. Knittle and G. K. Chang, "Delta-sigma digitization and optical coherent transmission of DOCSIS 3.1 signals in hybrid fiber coax networks," IEEE Journal of Lightwave Technology, vol. 36, no. 2, pp. 568-579, 2018.
 - [64] 3GPP TS 36.104: "Evolved Universal Terrestrial Radio Access (E-UTRA): Base Station (BS) radio transmission and reception," V15.2.0, 2018-03. (Release 15)
 - [65] 3GPP TS 36.141: "Base Station (BS) Conformance Testing," V15.3.0, 2018-06. (Release 15)
 - [66] <https://1.ieee802.org/tsn/802-1cm/>
 - [67] <http://www.ieee802.org/1/files/public/docs2018/cm-farkas-overview-0718-v01.pdf>
 - [68] <http://sites.ieee.org/sagroups-1914/public-documents/>

Deploying IP Video Services, Architectures and Technologies from the Head End to the Home network.

Video Full IP - Architecture

A Technical Paper prepared for SCTE•ISBE by

Eduardo M. Panciera Molanes

Chief of Architecture - Access and Service Platforms
Telecom S.A.
Agüero 2392, CABA, Argentina
epanciera@teco.com.ar

Adrian Grimaldi

Sr, Expert in Applications and Services
Telecom S.A.
Agüero 2392, CABA, Argentina
agrimaldi@teco.com.ar

Norberto Harmath

Sr. Expert in Access Services
Telecom S.A.
Agüero 2392, CABA, Argentina
nharmath@teco.com.ar

Gaston Diaz

Expert in Access Services
Telecom S.A.
Agüero 2392, CABA, Argentina
gadiaz@teco.com.ar

Marcos Aberastury

Product and Strategy Manager
Telecom S.A.
Agüero 2392, CABA, Argentina
maberastury@teco.com.ar

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	5
1. Definition of Managed and Unmanaged Services and Convergent Video System.	5
2. Driver to move towards Full IP Video Systems – a Convergent TV Platform.....	6
3. Where are we and which is the goal?	9
4. Video Distribution and Audience Behavior.....	12
4.1. DTV , Video over Internet and Hybrid distribution.....	12
4.2. Head End - Video Origin Server and CDN.....	15
4.3. Linear TV and Content on demand based on Hybrid distribution.	17
4.4. Impact of the Unicast in the Access Network - Audience Behavior.	19
5. Changes to reach Video Full IP system.....	27
5.1. Video Head End	29
5.2. Transport network	32
5.3. Access Network.....	34
5.3.1. Multicast in DOCSIS.	34
5.3.2. Multicast Channels – Dedicated, Shared or combination of both.....	36
5.3.3. CM Capabilities and channel assignment to CMs.	37
5.3.4. Uncorrected error and Partial Services - Multicast Resiliency in DOCSIS Networks	39
5.3.5. Video QoS for Unicast and Multicast.	44
5.4. Home Network.....	47
Conclusion.....	52
Abbreviations	54
Bibliography & References.....	57

List of Figures

Title	Page Number
Figure 1 – Managed and Unmanaged services.....	6
Figure 2 – Different kind of network to provide the video services	7
Figure 3 – Home Network (1).....	7
Figure 4 - MVC Model.....	8
Figure 5 – Technologies evolution helps the migrations towards IP	9
Figure 6 – The Hybrid scheme.....	9
Figure 7- Hybrid ecosystem solution.....	10
Figure 8- DTV and ABR Streams.....	11
Figure 9 – Introduce the Video Full IP service.....	12
Figure 10 – DTV Video services distribution.....	13
Figure 11 – Internet Video Services.....	14
Figure 12 – ABR Streaming	14
Figure 13 – Video Origin Server	15

Figure 14 - CDN	17
Figure 15 – Hybrid video distribution, DTV, and IP ABR	18
Figure 16 – Video CDN Traffic March 23th – 2018.....	19
Figure 17 – CDN HTTP Request by Device Type.	20
Figure 18 – Connected devices/hour	21
Figure 19- Connected devices/hour during the world cup	22
Figure 20 - Gbps per Service Group for Linear TV based on Unicast ABR.	23
Figure 21 – Channels Viewers Distribution.....	25
Figure 22 – Zipf Distribution	25
Figure 23 – High and Low concurrency events/programs	26
Figure 24 – IP Multicast Distribution	27
Figure 25 – Hybrid vs Full IP Video	28
Figure 26 – Today HE Architecture.....	30
Figure 27- Several encoding systems vs Unified encoding system.....	31
Figure 28 – Unified HE Architecture	32
Figure 29 – Multicast in the transport network.....	33
Figure 30 – Example of Multicast Forwarding in DOCSIS 3.0.....	35
Figure 31 – Dedicated or Share QAMs for Multicast	36
Figure 32 – Decrement of Multicast Gain because of diversity in CM capabilities.	37
Figure 33 – Multicast distribution in the Service Group	38
Figure 34 - Multicast Traffic issues due to Partial Service.....	40
Figure 35 – CM’s channels in Partial Service Mode per frequency.....	41
Figure 36 – Reconfiguration of channels for Multicast to reduce Video Issues per PS.....	41
Figure 37- Multicast Resiliency by Capacity reduction of Multicast SG.....	42
Figure 38 - Multicast Resiliency by Multicast to Unicast in CM with Partial Service.....	43
Figure 39 - Different kinds of traffic in DOCSIS	44
Figure 40 – Service Flow Configuration.....	45
Figure 41 – SCN definition.....	47
Figure 42 – Hybrid Home Network (A) vs Full IP Home Network (B).	48
Figure 43- Cable Modem Residential Gateway WiFi D3.0 IPTV ready.	49
Figure 44 – Cable Modem Residential Gateway WiFi D3.1 IPTV ready.	50
Figure 45 –Multicast Functions that CPE must implement.....	51
Figure 46 – WiFi Traffic capture with WMM.....	52

List of Tables

Title	Page Number
Table 1 - Resume of video distribution mechanism for Linear and Cod.....	18
Table 2- ABR Profiles.....	21
Table 3 – Distribution mechanism to distribute TV services in managed and unmanaged devices.....	27
Table 4 – Changes from DTV to Hybrid and form Hybrid to Full IP.....	28
Table 5- WMM “P” bit mapping.	52

Introduction

In recent years we have seen the growth of IP-based content distribution services, not only in terms of services over controlled networks where we can highlight the IPTV deployments made by TELCOs companies, but also the proliferation of content distribution in unmanaged networks, also known as OTT (Over the Top).

But the new paradigm of TV consumption does not impact only in the distribution, but also and more important, in a better user experience (UX), with advanced User Interface (UI), new integrated applications and functions, new ways to consume the video, not only linear and on-demand but also different flavors of network-based time shifted video (like cloud reverse EPG (REPG), cloud Digital Video Recording (cDVR), Pause live TV, and others), new ways to show the information to the end user, where the video itself is not the only main piece, also the metadata that exposes to the end user enriched information and assists them with improved visual approach, searching and recommendations, or generating extra information, during special events, that allows end client interaction with applications related to the content.

In addition to this, there are also a new set of devices like PC, Tablets, Smartphones, Smart TVs, Consoles Games, etc., these are in general known as Consumer Electronic (CE) devices. They are starting to be used more and more to consume video services, and even more, there is an interaction between those CE devices and the traditional STB that are using in general in the cable operations.

The cable industry must be prepared to move in this direction, and the purpose of this paper is to explain how our company is transitioning to video over IP delivery. We have already moved from legacy Digital TV (DTV) to Hybrid (DTV+IP) system and now, finally, we are starting to deploy Full IP Video delivery. The explanation will include the drivers that generate these necessities of the migration towards IP world, as well as the technologies and architectures involved, starting in the headend using High Efficiency Video Coding (HEVC) compression technologies allowing not only High Definition (HD) definition but also Ultra High Definition (UHD), the System Delivery Platform based on cloud which controls advanced functionalities like HTML based user interfaces, network DVR, reverse EPG and others, the media transport -not only over DOCSIS access network but also over FTTH¹-; and in the end the challenges and opportunities to deliver video in the home network using WiFi connectivity for customer electronic devices and for IP Set Top Boxes.

This document will describe which are the new necessities in terms of TV services' user experience, the definitions and first steps to move from Legacy TV system to a new ecosystem that allows satisfying those necessities. Then it will enumerate and explain some drivers to evolve the ecosystem towards a Full IP system. It also will take a look at which are the possible IPs mechanism to distribute the different types of video services, and finally, it will go through every domain, from the Head End to the Home Network explaining which are the changes that the operator will introduce finally to move to Full IP video.

¹ The document will mention FTTH in some cases, but the focus is specially HFC networks, however, most of the concepts are similar.

Content

1. Definition of Managed and Unmanaged Services and Convergent Video System.

Before starting with the explanation of the different systems over which TV services can be provided and then the explanation of why we are moving to an IP video ecosystem, we will make a distinction between what we will call managed and unmanaged networks as well as managed and unmanaged devices.

Managed networks are those ones in which the transmission of information/data can be guaranteed by managing the quality of service over it. For example, fixed telephony network is a managed network, as it is the IP network with quality of service policies over which telephony services are provided, on the other side, the Internet is an unmanaged network.

Managed devices are those ones which are controlled by the operator. The operator is who specifies the hardware (HW), software (SW) and applications that must be installed and executed over those devices, the end user just use them for the specific services that the devices were designed, and the user cannot install other applications that are not in the operator catalog or store. STBs are examples of this managed devices.

On the other hand, unmanaged devices are those ones that the end client can control, they can install a different kind of applications or even more they can modify and install the Operating System of those devices. For example, Personal Computers (PC), Tablets, Smartphones are examples of unmanaged devices.

In the video world, we could also add an extra level a separation between managed and unmanaged content, where the managed content is the one that the operator can ingest in the video system, it can configure the different video parameters, it also manages metadata and right of the content based on different technical and commercial agreement. Contrarily the unmanaged content is the one that operator cannot control, for instance, content that can come from OTT providers or content that could come in general from 3rd parties CDNs (Content Delivery Networks).

In Figure 1 it can appreciate the different combinations of uses cases that could be possible, where the complete managed services it shows the combination at the top of the picture (framed in red), this is the case of the Legacy DTV system or even more the Legacy IPTV system.

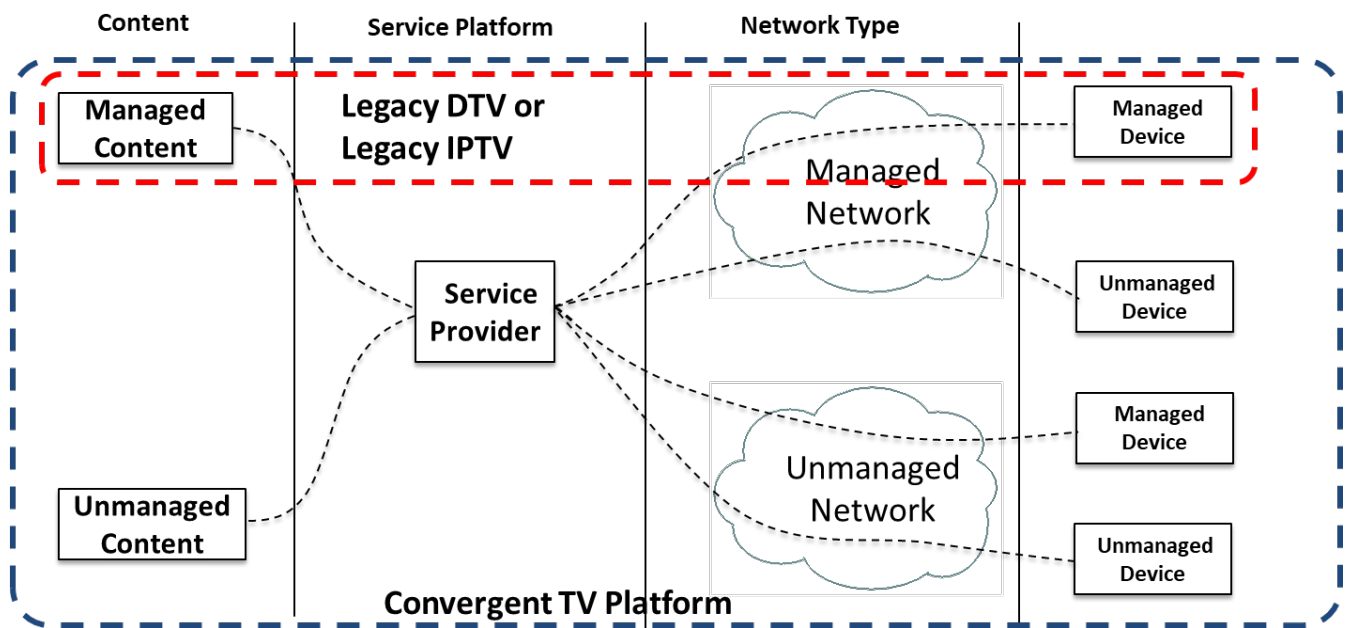


Figure 1 – Managed and Unmanaged services.

2. Driver to move towards Full IP Video Systems – a Convergent TV Platform

As it could be appreciated in the Figure 1, the traditional DTV or IPTV systems are very limited in term of type of devices where the services can be consumed, networks where the home network or more generally the devices can be connected, and the content sources (or applications) the end client can consume or use on their devices. On the other side, the legacy DTV system has strong limitations in terms of UX. We can split the problem into different planes, the first plane is related to the media distribution and devices where the media² is consumed, and the second plane is related to the user experience and control plane.

Regarding the media distribution the DTV networks generate huge limitation, that is where the customer can consume the content, consuming video on this system require a managed device that is an STB, which has to be connected to the COAX in the home. Today the operators provide services not only over HFC networks but also are deploying their services over xPON networks, Mobile or wireless Networks, xDSL and why not directly over the Internet, when the user moves temporally to other areas where the Operator that user belongs it does not have a managed network. All those networks have a common factor “IP” as a common mechanism to connect the services.

On the other hand, the end user does not consume video services only with STBs, but also in mobiles and stationary unmanaged devices (smartphones, tablets, PC/MACs, console games, smart TVs) which are mostly connected in home networks using WiFi, stationary devices that are most of the time connected at home, but mobile devices could be connected outside the home, in other 3er Party WiFi networks or 4G networks.

² With the term of media in general we are meaning video, audios, subtitles, etc.

As you can see, leaving side DTV system, all the networks aforementioned are IP based network. The Figure 2 and Figure 3 represent the different networks where the services must be provided including also the Home Network.




Figure 2 – Different kind of network to provide the video services



Figure 3 – Home Network (1)

On the other hand, we must consider the control plane that allows managing the service and user UX. In traditional DTV system, the control is limited to add or remove some video services packages Live TV, premium channels, Subscription VOD (SVOD), to buy video content based on Pay Per View (PPV) or Transactional VOD (TVOD). At the same time the presentation of the information to end user is limited and most of the functionalities are based on local applications on the STB that could or not exchange very basic information with DTV video Platforms, in addition to this is very complicated to modify or develop new services or applications, for instance the applications can base on Multimedia Home Platform (MHP), Open Cable Application Platform(OCAP), JavaTV or even more proprietary solution, that most of the time requires very specific programming skills to do changes or new functionalities, and the development time is in terms of years.

Nevertheless, OTT providers use much simpler and well know technologies that allow developing and improving their services in an agile way. They are “web technologies”, and in this world the UX is based on very well know technologies like HTML, JavaScript, Document Object Model(DOM), CSS, Webservers and Web Application servers, open source Databases, etc. and it uses lightweight protocols such as HTTP, REST vs RESTFUL open API, etc. The control plane is a distributed system that run-in end client devices and in the cloud, most of the logic is resolved not at the end user devices but in the cloud. In that way, a new functionality or change can be introduced very quickly. Figure 4 shows the architecture Model View Controller that decouple the application in three layers: 1) "Model" that represent the data-related logic that the application and finally the user work with, the "Presentation" or "View" that is the component that generates the logic of the UI, how the end client will see the information in the screen and how they interact with the application, and 3) the "Controller" that links the Views with the Models, the business logic occurs at the Controller, for instance in case of video applications it could manage which catalog has to be presented to a specific user with specific set of video packages and/or entitlements. Or for instance, the metadata (that could be stored in the data Models) is needed not only to have control over the media but also to provide information about it to the user. We could create extra Models to enrich with different sources of information, take information from the user behavior and generate recommendations, and to combine with other not traditional metadata sources to present a richer and improve EGP and/or Catalogs.

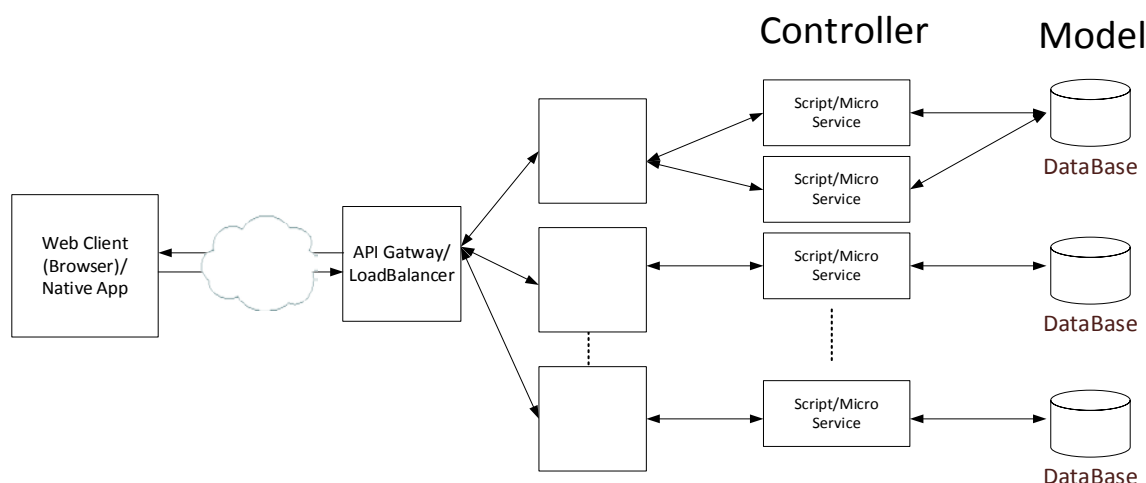


Figure 4 - MVC Model

This kind of architecture gives us some benefits over traditional DTV architecture and technologies.

- Much better and user-friendly EPG
- Elastic demand
- Open interfaces to change, improve or develop new functionalities.
- Easy to integrate with other platforms that allow integrating new ways to consume video (CuTV, REPG, etc.)
- Easy to integrate others 3rd party applications (for instance with another OTT video service that complement offer that Communication Service Provider (CSP) has)
- Personalization not only of the UI but also to the complete service at all
- Introduce companion devices functionalities.
- Its applicable not only to traditional CE devices but also to managed devices like STBs

An extra issue that we have in our DTV implementations, is that the system that we have deployed (more than 10 years before) is based on Motorola Digicipher II (DCII) scrambler, because of different legal and commercial constraints in the LATAM region, the only STB vendor can provide devices for this DTV system is the owner of the DCII technology, that generate a “lock-in” with the STB vendor leading to expensive STB’s constraints³.

Finally, the evolution of technologies helps the migration towards IP. The new encoding technologies like HEVC video generates a better performance in network usage and the make easy to have the extra network capacity to deploy new lineups (that requires simulcast with the today's DTV lineup). The requirements of 4K and UHD content require new STBs to support that kind of content, so to protect future investment it is better to use technologies that have a long-term support, and then to use IP STBs instead DTV STBs. New STBs are based on more open OS like Linux and/or Android TV, which allow working with the web models that were described above.

³ SmartCards cannot be sold in the region, if they could the costs are not good enough for the business



Figure 5 – Technologies evolution helps the migrations towards IP

All the drivers that were mentioned previously show a clear path to move the TV services directly over IP. For that in HFC networks, we must use DOCSIS, but that is not free, we need available bandwidth in the spectrum to allocate more DOCSIS capacity, now not only for High-Speed Internet services, and other data services, but also to transport the TV services. That could be the very first issue to deploy a FULL IP Video Service, but as it mentioned the system can be split into two layers, data plane and video plane, and actually most of the features that generate an impact in the UX, new functionalities, etc. are not relative with the media itself but with the signaling and control layer of the service, the media that is the heaviest in terms of network resources could be kept until the resource (spectrum in terms of MHz) is available, over the today networks (DTV). So here is where the Hybrid system appears, where the control plane is based completely over IP and the video could be based on IP (generally Content on Demand⁴ and some low audience linear channels) and the linear TV is based on DTV. This concept is depicted in the Figure 6, where the control plane (in red) is represented by the HTTP/HTTps protocol to communicate the devices with IP Video Service Platforms and HTML4/5 for UI, meanwhile, the video could be transported over IP Network or over traditional QAM like it was done in HFC network.

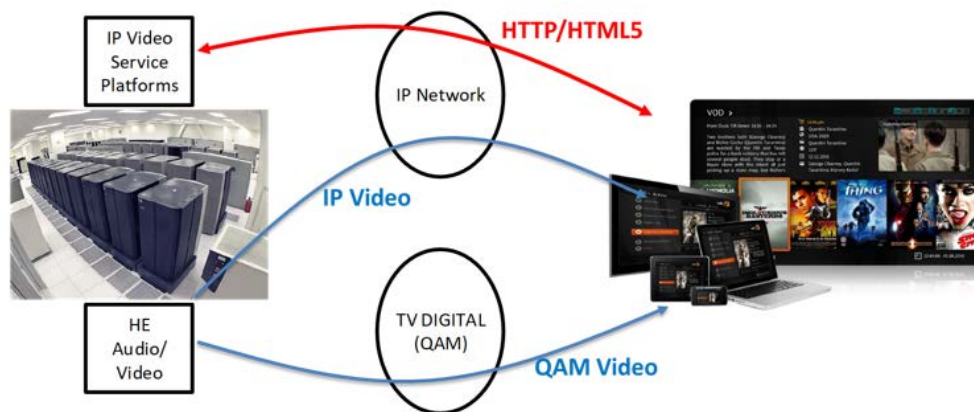


Figure 6 – The Hybrid scheme.

3. Where are we and which is the goal?

Telecom (former Cablevision Argentina) started to roll out a project, that was called “@TV” during 2015 to deliver the best in class entertainment services to its subscribers, to increase market penetration and to gain competitive advantage. @TV takes advantage of the latest technology solutions including the IP Video Service Delivery Platforms based on web technologies, Conditional Access (CA) and Digital Rights Management (DRM) system that allow control and protect content; Origin Server and Content

⁴ aka VOD or Video on Demand

Delivery Network (CDN) to prepare the media and their distribution on any kind of device to provide the next generation IPTV multiscreen video experience based on cloud.

@TV offers the following services: Linear/Live TV, diverse types of VOD models, Time-Shifted TV (TSTV), Network DVR (nDVR), Restart, Reverse EPE (REPG), advance UX, recommendations, companion device functions, bookmarks between devices, profiles in the same account, etc.

The first phase of this deployment was based on unmanaged devices (Android, iOS, PC, MAC, and Chromecast) and on managed devices Hybrid STB (DCII Hybrid STB).

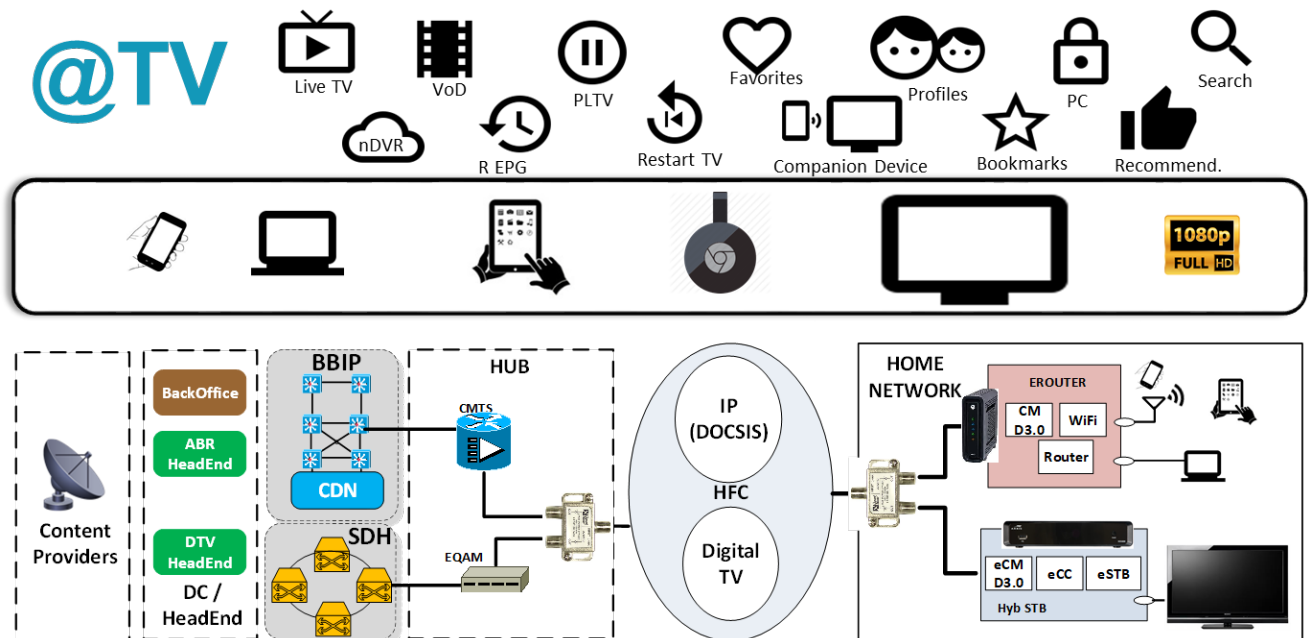


Figure 7- Hybrid ecosystem solution

To do this there is an ecosystem that is described in Figure 7, which is divided into platforms and networks, from the content reception, until it is consumed by customers in their homes and devices.

- The HE systems that receive the contents and process them to accommodate the formats to the needs of the different types of devices connected to different access networks
- The Backoffice Systems that oversee managing the video services that the user receives, protecting the contents, etc.
- The transport network that carries the contents from the HE to the access. The CDN is what allows the distribution of video using streaming technologies, like ABR, in a scalable way, it will have more explanation about that in the next sections.
- The different types of access networks to which the homes and/or devices of the clients are connected.
- And finally, the home that is challenging ecosystem where administration is not completely controlled, for example, the location of the devices and customer-provided equipment (CPE) in the home impact the quality of the service.

At this stage the System provides two different mechanisms to stream video:

- 1- DTV which is based on TV MPEG2-TS protected with Motorola Digiipher II (that is shown in the Figure 8 below as Legacy CA), as today this is a legacy mechanism that is used to stream Linear Content for Digital TV services; on the Home Network side Hybrid STB will receive these services.

The encryption is done in the Edge QAM equipment, so the video signal that is transported in MPEG2-TS is modulated, and it is sent to the access network HFC, then Hybrid STB demodulate the signal and decrypt the signal when it is authorized to (the STB has an embedded CableCard -DCII). The Hybrid STB is also connected in a secure way with the STB Controller (Legacy CA DCII) in the Head-End (HE) through legacy Out of Band (OOB) or DOCSIS STB GATEWAY (DSG) technology. The video distribution from the HE to the Edge Quadrature Amplitude Modulation (E-QAM) is based on L2 IP network transported over Synchronous Digital Hierarchy (SDH).

- 2- IP Streaming based on ABR transport and protected with DRM. That mechanism will be used to stream linear services and On Demand services. These streams are received from the content providers (Live and on Demand) and then ingested into the Origin Server. The Origin server receives the Live Content in several profiles using Adaptive Transport Stream (ATS) after transcoding process and it also receives on Demand Content in several profiles of MPEG4 Part12/14 after transcoding or directly from providers. The videos are stored in a NAS in a common format and delivered using Just in Time (JITP) method, that is when a given user get a piece of content the playout server of Origin Server packages and encrypts the content in real time. DRM assists JITP with asset encryption and delivers DRM license to client devices via a secured path. That is depicted as IP Unicast ABR in the Figure 8.

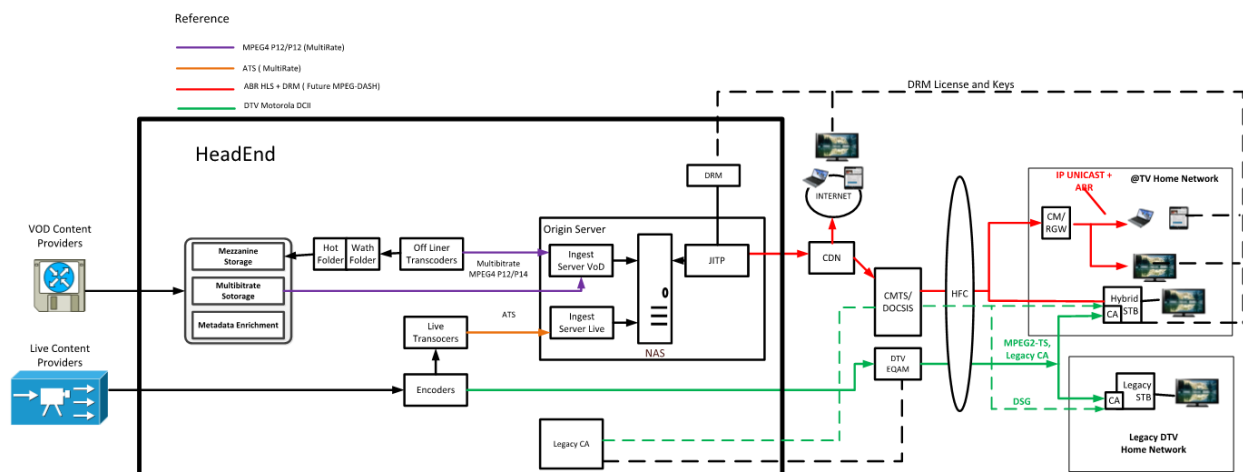


Figure 8- DTV and ABR Streams

To deploy a quick solution, the DTV encoders were reused as a source of Live TV Transcoders which generate the ATS for the Origin Server. As it will be described later in this document, it is not possible to generate the source for a Full Video IP transmission from the DTV HE and a new HE is required for that.

It is important to point out, such as it can be noted in Figure 7, that the home network is connected to two different access networks DTV and DOCSIS. The STB has an internal embedded CM (eCM) that provide the IP connectivity to the STB, this eCM does not provide connectivity to other devices inside of Home Network, it is needed a second CM that is in an EROUTER (showed as CM/RGW in Figure 8) which provide internet connectivity to other devices in the same home.

The goal is to add the BO functions, HE element, transport and access mechanism and Home network capabilities that allow distributing video directly over IP, using an IP STB connected in the home network

as “any other⁵” device, replacing the QAM by IP, so in the Figure 6 the ecosystem could avoid the QAM network. In the Figure 9 now appears the Full IP TV system combined with the previous one. There are some new components that are shown in the picture, for instance, a new Head End, an IP STB that is connected to the home network, there are also IP STB connected not only to HFC networks but also to FTTH networks or directly over Internet as a Managed devices over unmanaged networks, the idea is to reach those steps in a gradual way, with the same ecosystem, protecting the existent investment.

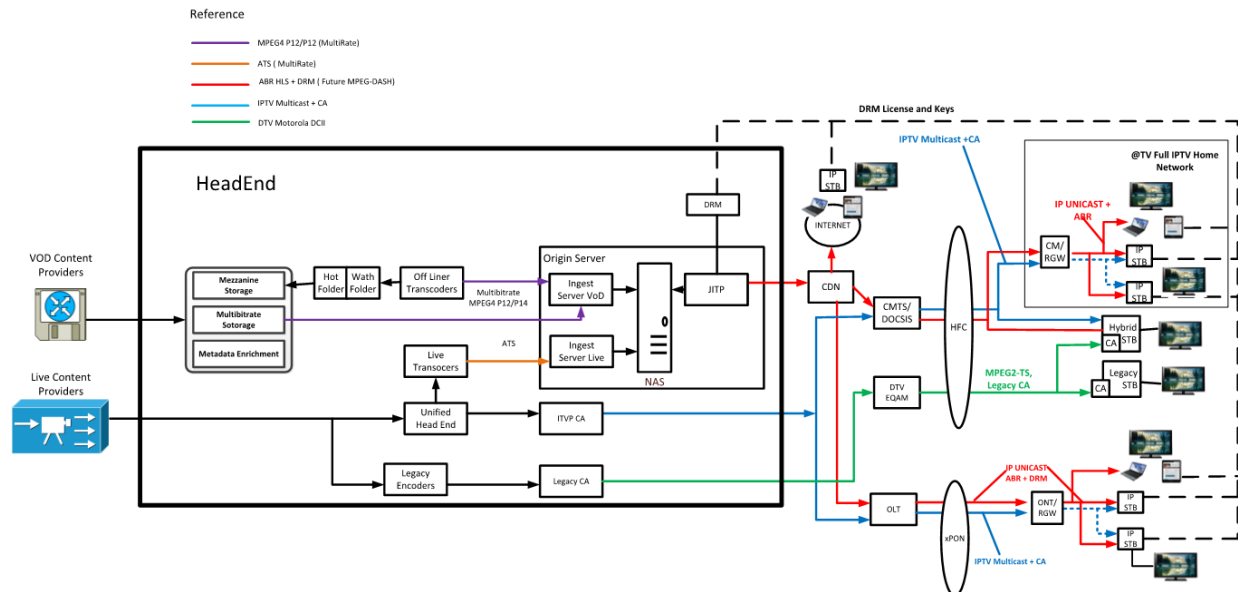


Figure 9 – Introduce the Video Full IP service

The document will explain what are the changes that we should introduce to meet with this goal, but first in next two sections we should understand a few more details about the video transmission mechanism and the audience behavior and which are the implications to replace the QAM network by all IP (DOCSIS in case of HFC).

4. Video Distribution and Audience Behavior.

4.1. DTV , Video over Internet and Hybrid distribution.

From the transport and access point of view, we can identify two kinds of services that the system must distribute over the network: Linear content and Content on demand (CoD), in the second one we could include VOD, CuTV, REPG.

The linear service is a kind of video traffic that is constantly requested by the end users, all of them are consuming the same content simultaneously with high concurrency of the users in many parts of the network. Linear TV services could be Live content (sports, news, etc.) but also movies, programs, etc. that are transmitted in a scheduled way. On the other hand, the second one (CoD) is requested by user request, and the concurrency is considerable lower than linear.

⁵ We should remember that if we have to consider a Managed Service as was explained in Section -1 , the connectivity of those IP STB must be managed in any way, the document will cover that point in the next sections.

This type of service is translated in practical terms into a greater or lesser consumption of resources of the transport and access network, the main one being the consumption of bandwidth in the access network where the resources are in general limited.

In the DTV system the linear content is based mostly in Broadcast distribution⁶ and the CoD is Unicast, both are digital video that is generated in the HE, distributed over IP network and modulated on the access network using Edge Quadrature Amplitude Modulation (EQAM), which receive video re-packetizes into MPEG transport stream and digitally modulates the transport streams onto RF carrier. Conceptually this distribution scheme is depicted in the Figure 10.

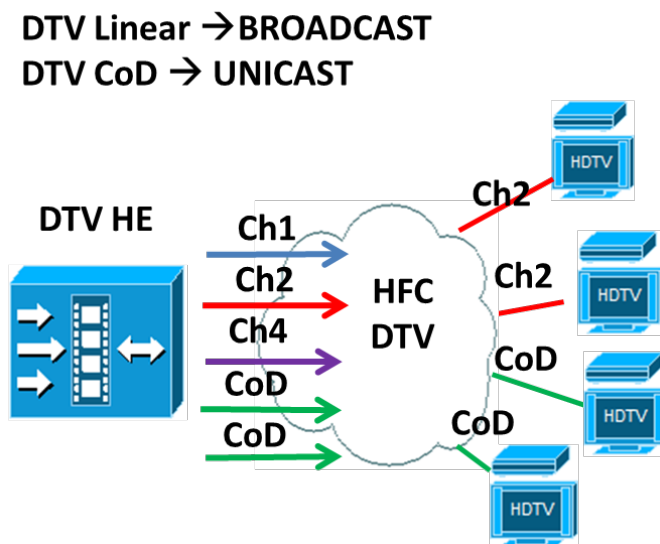


Figure 10 – DTV Video services distribution.

In the case Linear TV for the Broadcast mechanism, the network is sizing by the number of channels. All of the linear content is available in the network even when no user is consuming it. For instance, this case is shown in Figure 10, channel 1 and 4 is in the network but there are not clients that are tuning this channel. Let's note that that means that are networks resources that are being wasted but they are not really used by STBs. So, in DTV for Linear services the HE sends the channels to network, and all the channels are in network even when nobody consume those signals, when a given STB wants to consume a specific channel, it should look for in the network, for instance, the frequency where the channel is modulated, the program inside the transport stream, etc.

For CoD the content is pushed to the network only based on user request, and the even when two or more clients are watching the same content, there is a copy of the content for each one, let's remember that in CoD every client could to start, pause, rewind, do fast forward the content, so even when the "movie" could be the same, it could be shifted in the time and then that requires different copies for different users, that is because Unicast is used.

In contrast to DTV systems, that are managed networks, the Internet is a completely unmanaged network, and so the mechanism to distribute video is completely different. First at all those mechanisms must be based in very standard IP protocols, they use TCP/IP as transport and HTTP or HTTPS to encapsulate the

⁶ There exists another more efficient mechanism in DTV that is Switched Digital Video (SDV), but that is not the case in our company and it is not explained in this document.

video, that allows transport it easily over internet, even more, it could work on devices connected behind Network Address Translation (NAT) or Port Address Translation (PAT) that are widely used in IPv4 Home Networks. Besides this layer 3 IP networks do not allow to propagate broadcast between layer 2 domains, so for both kind of services, Liner and CoD Unicast is used. The Figure 11 illustrate that case, when a device has to consume a linear channel, that device ask for this content to network (in fact the devices get the content from a Content Delivery Network (CDN) that it is described later) and then the content is routing over the network until it reach de device, if other device tune the same channel then the same happens and therefore there are two copies of the same content on the network, as you can see the sizing of the network for linear services grows up proportionally to the devices that are consuming those services. Same happens for CoD.

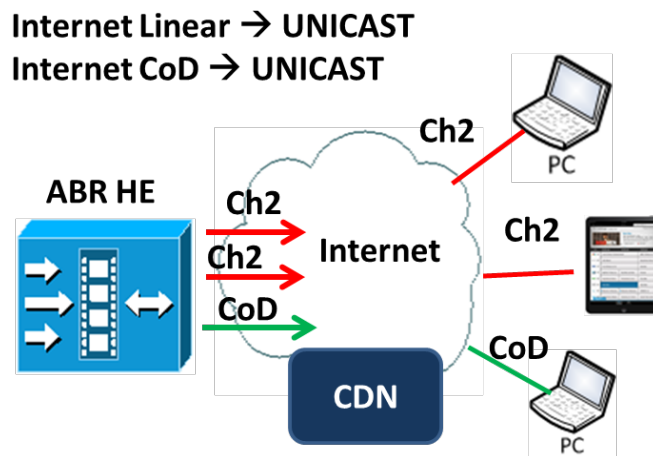


Figure 11 – Internet Video Services

There are different video streaming techniques, like HTTP download, Adaptive Bit Rate (ABR) that also uses HTTP/HTTPS as transport, Peer to Peer Video, etc., but the most widely used is the ABR, which allow to play out the content almost in real time, the quality of the video is adapted to the network quality, CPU devices capacity and other factors that could depend on the player implementations and flavor of ABR mechanism (2).

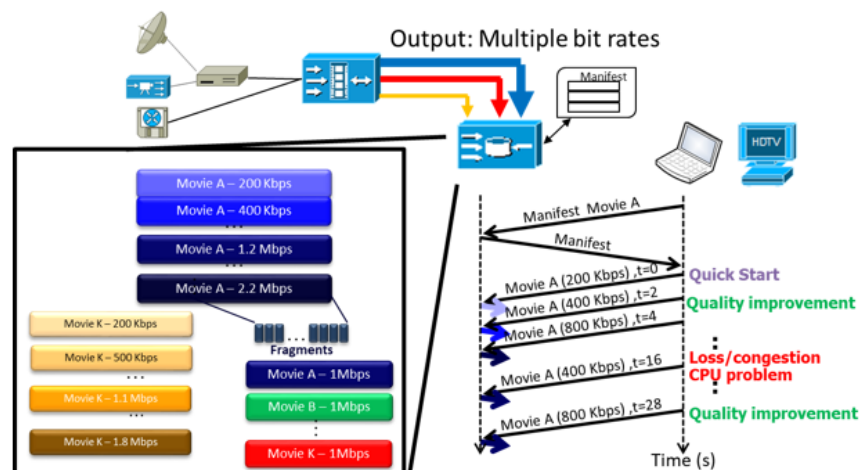


Figure 12 – ABR Streaming

Figure 12 shows conceptually how the ABR streaming works, first at all it must be generated different copies of the same content, different resolutions and bit rates, that are called profiles, then every copy is segmented into small chunks, each segment only has just some video's seconds, and then a special file that is called playlist or manifest is generated, this file contains the URL of every piece of segment of every profile, so when the end device has to consume a video, it starts getting this manifest, and read the URL of the first chunk (generally the lower bit rate profile to have a quick start) and according to the player strategy in function of the network quality, CPU's device capacity or others it could start to get better and better profiles and improve the video quality.

Let's understand now how the IP ABR streaming is originated and distributed in more detail.

4.2. Head End - Video Origin Server and CDN.

In the Figure 8 a block called Origin Server (OS) is the HE component that is in charge of to generate the different profiles of Linear and CoD video, and to make them available to be consumed by the clients through the CDN. Figure 13 is a zoom over this component.

Video OS works as an HTTP Origination Server (Web Server with Video Enhancements), its goal is: to ingest, store in NAS and deliver VoD content using HTTP ABR format, a mechanism to receive, buffer and deliver Live content (that comes from broadcasters) using HTTP ABR format, and to provide a way to reach end-user devices that support that kind of format and delivery mechanisms. Video OS is made up of several components to get more performance and efficiency.

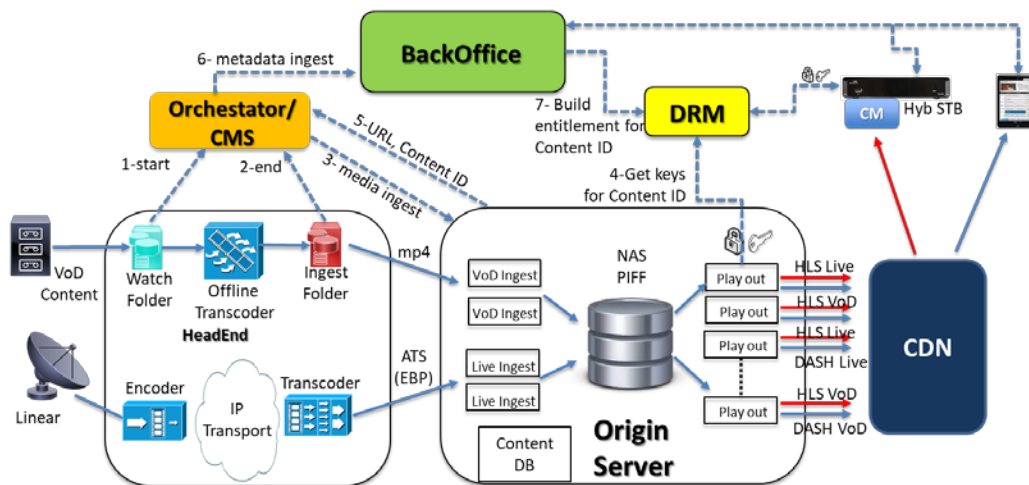


Figure 13 – Video Origin Server

VoD Ingest Server ingests video that could be in MP4 (format created by the Moving Picture Experts Group (MPEG)) as a multimedia container, or IIS Smooth Streaming Media Video (ISMV), or other formats, and stores it into NAS in Protected Interoperable File Format (PIFF) format (pivot format), for doing so, it receives an order from Orchestrator to queue a job that fetches content from Watch/Ingest Folders using FTP or HTTP, once the ingest process ends up the URL which points to VoD content is passed back to the Orchestrator/Content Management System (CMS).

Live Ingest Server receives video from On-line transcoders which deliver them using some kind of Adaptive Transport Stream over IP Multicast (3) (CableLabs ATS-EBP is preferred, but it could be others), this video is kept into a live buffer (time-base programmable size) that is useful to create and

continuous update live manifest/playlist files with a programmable window length and to create DVR assets (those that are used for nDVR, CuTV, REPG, nPLTV, and TSTV in general services) , for doing so, it receives in advance an order from the video BackOffice (BO) through Planner Content DB to queue a job that makes permanent those segments/chunks when they are within the live buffer and are going to be part of the DVR asset. URLs for live content points out to chunks/segments in a live buffer and are populated in the BackOffice when creating channels, URLs for DVR assets are passed back to BO Planner Content DB. Those URLs are requested by the devices when they want to play out a given content.

Playout Servers receive requests for playlists and contents from CDN or from end user devices directly (it doesn't matter if is a VoD, Live or TSTV service), create playlist/manifest on the fly, take content from live buffer or NAS, encrypt and package it on the fly and deliver to CDN or end-user device. For doing so, Video Origin Server communicates with Digital Right Manager (DRM) using a protected path to create an encryption key for that content at ingestion/buffering time, and it receives that key (and the rotated keys) through using same protected path at delivery time for encrypting, packaging video and adding URL that points to key server (generally inside the DRM) into playlist/manifest, so content that leaves Video Origin Server is DRM protected. Playout Servers are also capable to deliver diverse types of ABR formats, Figure 13 shows HTTP Live Streaming (HLS) (4) and MPEG-DASH (5).

As it was described previously ABR is a unicast scheme video distribution and sizing of the system (OS and network distribution) is based on the number of devices, therefore using OS as the system that deliver the stream directly to the devices could generate a big issue for scalability, for instance in terms of I/O capacity of the NAS to be read by the Playout servers. But that is not the only problem, the video must be distributed from the HE to the end client using an IP transport Network and using Unicast. That has two major issues:

- 1) The distance and therefore delay that could exist between the end user device and the HE, as ABR use HTTP/HTTPs as transport and those use TCP the delay could generate a limitation for the bit rate that the connection could reach.
- 2) Again, the sizing of the transport network to support every Unicast stream between clients and HE systems.

CDN fixes these two major issues (and others), instead of that end devices get the content directly from OS, they get the content through cache layers. Figure 14 shows conceptually how the CDN works, the end device instead to get the content directly from OS, it gets from Edge Cache, if the Edge cache has the content (what is considered as a hit) that is being requested then it delivers directly, if the content it is not in the Edge cache, that means a cache miss, and then the Edge looks for it in the Intermediate Cache (IC), that is another layer of the CDN, where it can be again a hit or miss, with a hit it is served by the IC to the Edge and finally to end device, or if it is a miss the IC asks for the content finally to the OS.




Figure 14 - CDN

Each time that a content is serviced by a Cache node, it stores the content by a given expiration time so next time that other devices get for the same content it can respond directly. For linear content where the content is consumed in real time, when a device asks for a given channel the next device that asks for the same one there is a high probability of a hit and the traffic over the IP Backbone (BBIP) is proportional to the Edge cache nodes deployed and not to the number of end-user devices. With the same logic, the traffic that is requested to the OS is just proportional to the number of IC nodes. At the same time if the Edge Node is deployed at the border of the access network, as it is depicted in the Figure 14, then the delay between these nodes and the end client devices could be minimized, which mitigate the problem of the delay in TCP.

There are diverse types of CDN topologies, with more or fewer numbers of intermediate layers, but at the end, it must have an Edge cluster, that is a set of Edge cache node that works all together caching content and serving a same set of customers. This cluster should be placed as close as possible to the end user devices, for instance at HUB or in a regional site that connects several HUBs, the topology will depend on the distances and amount of traffic that the Edge cluster must deliver to the access network. The way, in which the end devices determine which is the Edge cluster, and particularly the Edge Node inside the cluster, is based in a CDNs control layer, that works as routing layer for the CDN, and this routing layer must have logic that must be a function of devices' IP source, geolocation, load of the Edges nodes, it could also be aware of the cache content to redirect the device to the edge that already has the content and increase the hit probability, it may use the user-agent of the device player, and others that in general depends on the CDN implementations.

4.3. Linear TV and Content on demand based on Hybrid distribution.

If ABR is used in a managed network, and the Communication Service Provider (CSP) could guarantee the quality of the network, the same technique that is used on Internet to distribute the video can also be used to distribute the video in those controlled networks and even when the ABR has the capacity to adapt the quality of the video to the network conditions, if the quality of the network allows the ABR uses the best profile then the quality of the video could be guaranteed. On the other hand, if the STB could support this ABR mechanism then it can be used the same technique for Managed or Unmanaged networks and for Managed or Unmanaged devices. Even more, the same DRM can be used for each of those

combinations. The **Figure 22**Figure 15 shows the case where there is a Hybrid STB that can receive Linear TV from DTV HE, but also have an embedded CM that allows receiving IP ABR streaming for both Linear TV or CoD.

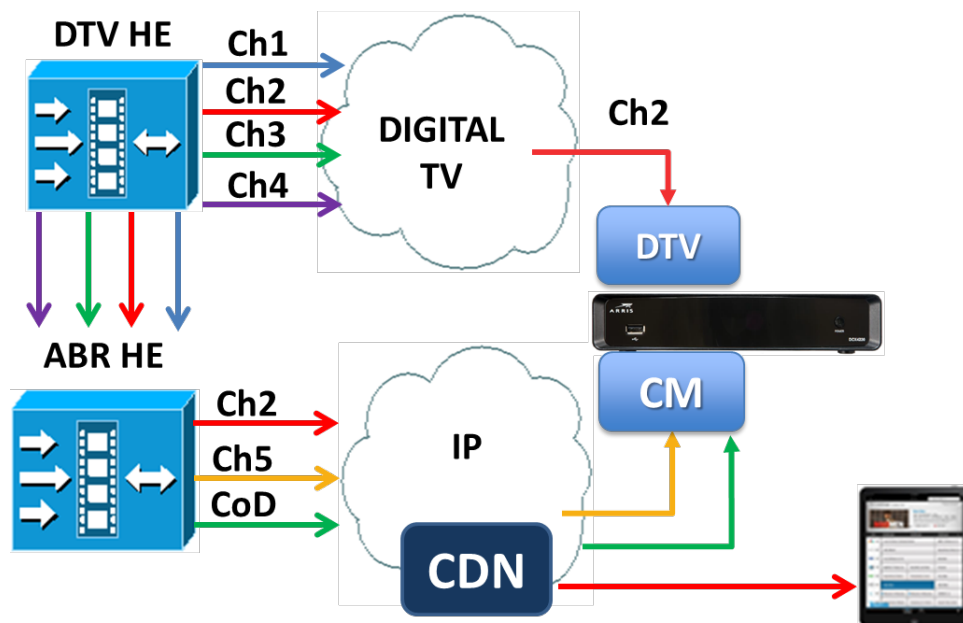


Figure 15 – Hybrid video distribution, DTV, and IP ABR

As a resume Table 1 shows which are the available distribution mechanism for video in a hybrid ecosystem such as presented up to here.

Table 1 - Resume of video distribution mechanism for Linear and Cod

	TV Digital	IP
Live TV	Broadcast (BW -> Nro Ch)	Unicast (BW -> Nro Clients)
CoD	Unicast (BW -> Nro Clients)	Unicast (BW -> Nro Clients)

The interrogation that it could be asked is, what's happens if in the in the Figure 15, instead of using the DTV the system configure the STB to receive all the channels over IP, based on the ABR mechanism that it has been explained above? In that scenario the complete ecosystem it could be considered based on a FULL IPTV video, all the BO interfaces are in IP and the video distribution too. However, that could have some issues. One of them is the delay in the video that ABR introduce, as it was explained the ABR mechanism is based in chunks/segments of some video's seconds which are buffered, the player needs to receive some of those segments previously to start the play out, and that generate delay against the real signal (between 20 and 40 sec) that it is not acceptable for live content like sport events. However, there are some new Low Latency ABR technologies that are considering this issue and probably could be addressed in the future (6). Alternatively, as it was aforementioned the CDN helps with the scalability of a Unicast distribution on the BBIP, from the Headend until the access (edge) but how does the access should be sized to support the video unicast distribution?

So, to answer this question we are going to analyze the audience behavior.

4.4. Impact of the Unicast in the Access Network - Audience Behavior.

Fixed access networks can be divided into two kinds: 1) point to point access network like xDSL mostly using in TELCO companies, where there are dedicated resources from the access to node until end client and 2) share access networks, like DOCSIS or FTTH, where the resources are shared by all the clients that are connected to those access network. In the first case, point to point access network if it is considered that CDN it collocated at the side of the access node, the video unicast traffic for each client, at access level, will impact in the client itself because each client has a dedicated link. Contrary to this, in share access network the video distribution based on Unicast for a given client will impact in all other clients that share the same access network. So, the question that shows up is if the Video Unicast distribution mechanism is scalable in such kind of networks. And the answer is that it will depend on the number of devices that consuming video using Unicast and the capacity of the access network, and the number of devices consuming a given type of content depends on the audience behavior, so to analyze that it will use a real case that it was taken from a day where there was a very important sport event and the traffic in the CDN reach its maximum value until that day.

The picture showed in Figure 16 is the traffic of the CDN, dark green and light green is the traffic of two different clusters that are used for unmanaged devices⁷ and Hybrid STB, meanwhile, the orange traffic is the cluster dedicated to unmanaged devices that are Chromecast. It is important to differentiate this traffic from the rest of other unmanaged devices because the Chromecast is connected to a big screen, like an STB, and the user behavior is different.




Figure 16 – Video CDN Traffic March 23th – 2018

The picture shows two peaks between 16:00 and 19:00 hrs., that is the time slot where the event was broadcasted, every peak corresponds to the first and the second half of the match soccer match. Another highlight of the graphics is that those peaks overpass the peak in the TV prime time that is from 21 to 24hs. (at least in Argentina).

⁷ PCs, MACs, Android and iOS devices.

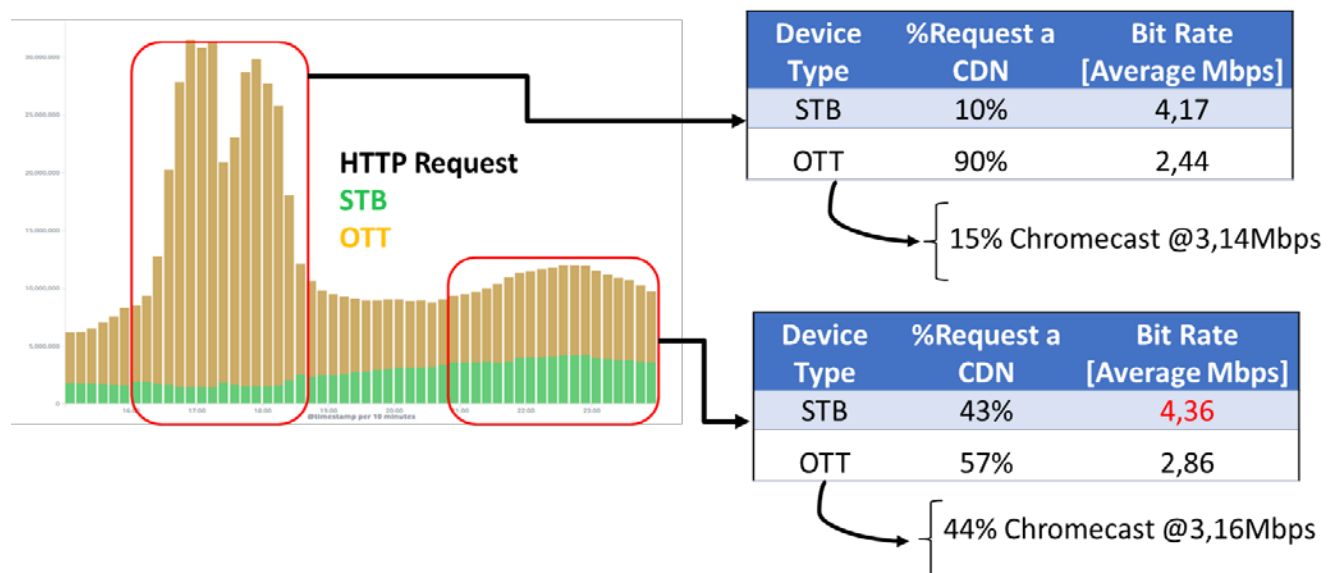


Figure 17 – CDN HTTP Request by Device Type.

If we analyze which devices were those that consume this Unicast traffic from the CDN it is observed that during this sports event 90% of the devices were unmanaged (OTT) and just 10% were managed (STB), that can be appreciated in Figure 17. Let's remember that Hybrid STBs consume linear services from DTV system and not from the CDN, so most of the STB HTTP request is because CoD and not because linear TV consumption, contrary to that most of the OTT HTTP request were because of the live sports events. During the prime time the amount of HTTP request from OTT devices and STB are quite similar (43% for STB and 57% for OTT). It is important to point out that the sporting event was during working hours where the people are not in their home so they consume the video using their OTT devices.

Another highlight is that the average Bit Rate in the OTT device is lower than for STBs, that is because STBs are connected in a managed network where the throughput is assured, while the OTT most of the time are connected in wireless networks, which are unmanaged networks, such as home networks where WiFi is regularly a best effort shared network, 4G or even worst in 3G mobile networks.

In Table 2 there are the different set of video profiles for various HD and SD kind of movies or live programs, there exist a different kind of content complexity in terms of for instance movies with soft or abrupt changes of scene, sports events, news programs, etc. There is a wide variety of bit rates and resolution from the lower 400kbps up to the higher 7528kbps. Most of the time the content is in HD @1280x720 and sometimes @1920x1080, and the average that STB is in 4.36Mbps what means that they are using the highest of the profile that is mainly because there are QoS to support the distribution mainly in the access.

Another high spot to point out is the case of Chromecast devices. They are generally connected at home WiFi network, as you can see during the sport event just the 15% of OTT devices were Chromecast, while during the prime time almost half of the OTT devices were Chromecast, let's remember that Chromecast is a device that is connected to a big screen (in HDMI of TV set) and receive the stream from network through WiFi, and in general that provide better throughput than in 3G/4G networks and that is because the average bit rates in Chromecast devices are bit higher than in OTT devices in general.

Table 2- ABR Profiles

HD Low Complexity (HDLCL)						HD Generic Complexity (HDGC)						HD High Complexity (HDHC)					
Bitrate (kbps)	H	V	FPS	Profile H264		Bitrate (kbps)	H	V	FPS	Profile H264		Bitrate (kbps)	H	V	FPS	Profile H264	
530	424	240	29.97	Main	2.1	604	424	240	29.97	Main	2.1	678	424	240	29.97	Main	2.1
1040	640	360	29.97	Main	3	1192	640	360	29.97	Main	3	1344	640	360	29.97	Main	3
1480	854	480	29.97	Main	3	1726	854	480	29.97	Main	3	1972	854	480	29.97	Main	3
2100	1024	576	29.97	Main	3.1	2420	1024	576	29.97	Main	3.1	2740	1024	576	29.97	Main	3.1
3000	1280	720	29.97	Main	3.1	3450	1280	720	29.97	Main	3.1	3900	1280	720	29.97	Main	3.1
5500	1920	1080	29.97	Main	4.1	6514	1920	1080	29.97	Main	4.1	7528	1920	1080	29.97	Main	4.1

SD Low Complexity (SDLCL)						SD Generic Complexity (SDGC)						SD High Complexity (SDHC)					
Bitrate (kbps)	H	V	FPS	Profile H264		Bitrate (kbps)	H	V	FPS	Profile H264		Bitrate (kbps)	H	V	FPS	Profile H264	
400	320	240	29.97	Main	2.1	454	320	240	29.97	Main	2.1	508	320	240	29.97	Main	2.1
650	426	320	29.97	Main	3	730	426	320	29.97	Main	3	810	426	320	29.97	Main	3
780	480	360	29.97	Main	3	894	480	360	29.97	Main	3	1008	480	360	29.97	Main	3
1010	576	432	29.97	Main	3	1158	576	432	29.97	Main	3	1306	576	432	29.97	Main	3
1200	640	480	29.97	Main	3.1	1366	640	480	29.97	Main	3.1	1532	640	480	29.97	Main	3.1

Following, it will analyze CDN and Access Network capacity required to support a Full IP video distribution, instead of DTV based on Broadcast, using IP Unicast ABR for Linear channels.




Figure 18 – Connected devices/hour

The Figure 18⁸ shows the connected devices per hour during the day that it is analyzing. As it illustrates there is a peak of OTT devices during the sportings event, but this peak reaches just 6% of the total of active OTT devices in the system. That is the amount of OTT devices that mostly contribute to reach the

⁸ 27/03/2018: 338,383 active STBs/ 1,169,663 active OTTs

140-150Gbps that it shows in Figure 17. At it can expect there are a lot of more OTT devices that STBs in the system, the average is 3.5 OTT devices per each STB, Nevertheless the behavior is clearly different, even when the STBs are less than OTT devices there are a lot of more STB devices connected and consuming video at the same time than OTT. At the prime time, there is a 56% of STB concurrency. Something important to point out is that this number of STB concurrency is in a regular day, but if the prime time it is overlapped with very popular sport event this STB concurrency could reach 80%, that can be appreciated in the Figure 199, that is the same graphic of connected devices/hour but during the “FIFA world cup”. This event was on Saturday when most customers are in their home and it can be seen how the STB concurrency it is close to 80%, and even more, soccer is a very popular sport in Argentina and due to the others world cup’s match during the day, the concurrency was almost flat the rest of that day.




Figure 19- Connected devices/hour during the world cup

Therefore, we can use these figures to estimate which should be the DOCSIS Service Group (SG) capacity to support the Linear TV based on ABR Unicast, so let's assume:

- Average Bit Rate per STB = 4.36Mbps
- Max Bit Rate HD per STB = 7.5Mbps
- Max Bit Rate UHD per STB = 15Mbps
- HHPP/SG from 500HHPP/SG to 32HHPP/SG, Penetration 60%
- Two different cases of concurrency, for the average of 56% and for the case of very popular event 80%
- 1.5 STB per Home

⁹ 16/06/2018: 403,407 active STBs/ 1,570,135 active OTTs




Figure 20 - Gbps per Service Group for Liner TV based on Unicast ABR.


Figure 20 gives the capacity needed in the SG for different sizes of services areas (HHPP), for three video bit rates (today average, Max for HD and Max for UHD) and two scenarios of concurrency (56% and 80%), and the graphics give an idea of the boundaries of the SG capacity.

A real HFC plant with 24SC-QAM could provide almost 900Mbps, with 32SC-QAM channels and 1x96Mhz DS OFDM block it could reach almost 2Gbps and if it considers 2 more 196Mhz DS OFDM block it could get approximately 5/6Gbps depending on the modulation profiles.


A conclusion here, in those scenarios and conditions the picture shows it is almost impossible to provide Linear TV using Unicast ABR just with 24SC-QAM, the network must support at least the first version of D3.1 and even with that, the SG size should be reduced at least until 128HHPP/SG, and that is only to support the linear TV services over DOCSIS, but the network must also support HSI services so possibly the SG must be reduced even more, 64 o 32 HHPP/SG. That is something that will happen during next years and very gradually, but today our network is not in those conditions, most of them are in 500 or 256HHPP/SG, and even when at plant point of view, it supports D3.1, the EROUTERS deployed are 24x8 D3.0. So, that way to provide Linear TV it could result very expensive and inviable in terms of investment. The example here was done based on DOCSIS, but similar analysis and results could be done in xPON, where for instance GPON or XG-PON networks that could achieve almost 2.5Gbps or 10Gbps in the downstream direction respectively.

The question that arises is: is the behavior the same for all the channels? As it can expect during a very popular live event the audience will tune the channel and at the same time, other channels will have less number of viewers. So, the other point that we should analyze is the audience behavior. If the channels are ordered from highest audience to the lowest audience, it gets a graphics like the one shown in the Figure 21¹⁰, which shows the channels viewer distribution.

¹⁰ The number of channels in the lineup is almost 200 channels the Figure do not show all of them at the tail of the graph.



(a)



(b)




Figure 21 – Channels Viewers Distribution

This picture depicts which is such distribution during the event day (27/03/2018) at three different hours, (a) 12 AM previously to the event, (b) 5 PM during the event, and (c) 10 PM during prime time. Cases (a) and (b) shows that the 80% of the audience just watch 30 channels and but during the event the audience is concentrated in the channels that were broadcasting the event and then 80% of the audience is just in 17 channels. In general, the audience behavior can be fit to Zipf's distribution (7) (Figure 22), which has a form of power-law like following:

$$P_i = \frac{i^{-\alpha}}{\sum_{i=0}^N i^{-\alpha}}$$

Where the percentage usage varies as a power of channel rank, where N is the number of channels and α is a shape factor (8).




Figure 22 – Zipf Distribution

Share access network, like DOCSIS or xPON network, are divided into groups of services: N1, N2, N3, N4, etc. Conceptually, it could say that there will be channels that will have a high probability to be tuned

in every network N_i , and others that are not so popular with low probability to be tuned and then only some clients in only some networks will be connected consuming those channels. It is represented in the Figure 23, where channel 1 is a popular channel and 2, 3, 4 and 5 are not so popular, and it is possible to split those channels into at least two different sets, high concurrency events/programs, and low concurrency events/programs.

In the aforementioned explanations it was shown that Unicast Linear video distribution is not scalable for networks with more than 512/256 HHPP per SG, however one of the causes of that is the concurrency; if the audience behavior, for a given set of channels, is such as the concurrency is low, the unicast is an option that can be used. The Figure 21 a, b and c show that most of the channels have low probability to be tuned (low concurrency). So, there are a lot of channels that can be distributed by Unicast.




Figure 23 – High and Low concurrency events/programs

The problem that it should be addressed is the set of channels with the highest popularity because depending on the size of the SG it could have not enough capacity to support the linear services and the other data services (HSI, etc.). So, for those set of channels Multicast distribution is used. This mechanism allows replicating the same stream of information between several clients that are requiring it, with just one stream for all them. The concept of Multicast gain that is the relation of the amount of the clients that are watching the same channels. That generates a most efficient mechanism to distribute the linear video and then saves in the HFC/xPON networks resources. Then, let's suppose to have a lineup of five channels, in the Figure 24 there are 3 clients that are tuning the channel 1 and other that is tuning channel 2, the channel 1 it will appear on the network just then first client tune it, then when a second client wants to join to the same channel, which is called multicast group, then it just generates a JOIN message to this multicast group, and then multicast mechanism in the network forwards the traffic up to this second client, without necessity to generate a new video stream in the whole network.




Figure 24 – IP Multicast Distribution

Table 3 gives a brief of what was discussed during this chapter, about the different mechanisms to distribute linear TV and CoD services.

Table 3 – Distribution mechanism to distribute TV services in managed and unmanaged devices.

	Digital TV	Unmanaged devices (OTT)	Managed devices (STB)	
			Not popular channels	Popular Channels
Live TV	Broadcast	Unicast	Unicast/Multicast	Multicast (Unicast only small SG)
CoD	Unicast	Unicast	Unicast	Unicast

As it was explained previously in broadcast transmission the sizing of the network depends on the number of channels and it is independent of the number of clients connected; contrarily, in unicast, it depends on the number of clients connected consuming video but it is independent of the number of channels. In Multicast distribution, the network sizing will depend on the Multicast Gain, which is a function of the number of the client connected and the number of the channels.

When the number of viewers is low the Multicast Gain is low, which is also another justification for why to go to Unicast distribution in the SG with few numbers of clients. For the same the numbers of viewers bigger number of channels less saving in Multicast distribution, the Multicast Gain decreases, then if we kept the multicast for the most popular channels it could generate a big saving in terms of resources.

5. Changes to reach Video Full IP system

In a FULL IPTV system, the Broadcast distribution that is used for Linear TV service must be replaced by an IP mechanism, and as it has been explained in aforementioned, for linear TV service there are two mechanisms that have to be used, Multicast for the most popular programs and Unicast for less popular programs. Then, the goal is to replace the QAM Video in Figure 6 by a complete IP distribution. This change will be translated in modifications in every domain of the complete ecosystem from the

HE/Datacenter to the Home Network, going through IP transport network (CORE and BBIP) and the Access network. Figure 25 shows the Hybrid and Full IP video ecosystem and it highlights in red the modifications that must be introduced over the complete system. During the next section is going to explain deeper which are the changes.

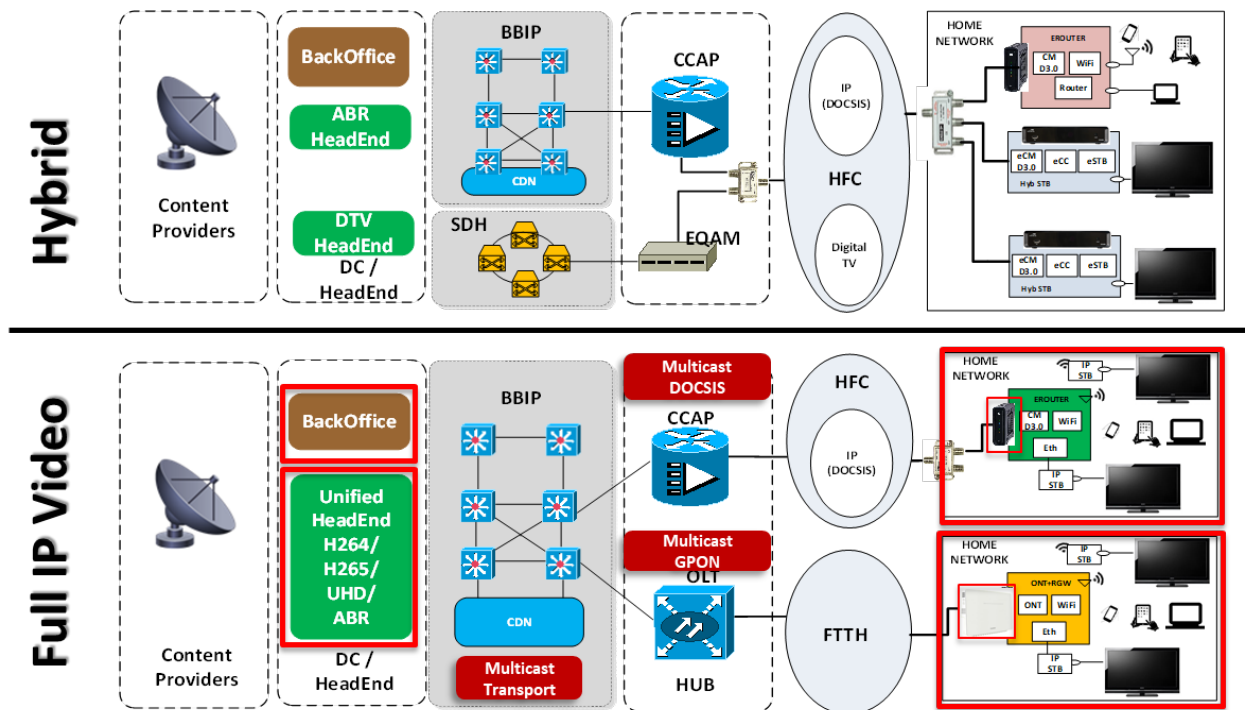


Figure 25 – Hybrid vs Full IP Video

Something important to point out is that the technologies are not the only change. There are several other changes at the process level in the company, for instance just to mention one of them the installation process in the home network, in the hybrid system there are not so many changes in this aspect, the technician kept installing the STB as a legacy DTV STBs, but now the IP STB is connected using new techniques, Ethernet, WiFi, PLC extenders, etc. And the new skills and tools are needed in the toolbox of the technician. In general, those changes are as difficult as the technology itself. Table 4 enumerates the main aspects to evolve from DTV to Hybrid, and then to Full IP.

Table 4 – Changes from DTV to Hybrid and form Hybrid to Full IP.

Domain	Legacy DTV	Hybrid System	Full IP Video
Head End	It manages the encoding system for Linear TV and VoD generating the MPTS for Linear and SPTS for DTV distribution.	Linear TV signals are reused, the current HE works with the Hybrid STBs. New ABR transcoding is needed to support CoD for any devices and Linear TV for unmanaged devices. It is transported over Unicast IP.	A new HE is required to encode the content in an appropriate format for the IPTV system. A new concept of Unified HE appears. New technologies like HEVC, Virtualization, UHD, are exploited.

Domain	Legacy DTV	Hybrid System	Full IP Video
	Conditional Access system where the different Linear TV packages are defined. Legacy VoD Back Office to manage the VoD distribution over QAM.	New IP Back Office is deployed to support the new User Experience, new CE devices, searching, the recommendation engine, etc. There is an integration with the Legacy DTV CAS to consolidate the Linear TV packages. A DRM is also needed to manage the rights in Unicast ABR distribution.	Some changes are required in other to signaling the Multicast Content and the DTV CAS is replaced by IPTV CAS.
IP Core and Backbone	Media is transported with IP over SDH or DWDM in L2 connectivity from HE to the Access networks.	It reuses the same transport mechanism used for Linear DTV. There are no configurations changes in IP Core or Backbone. A new CDN is deployed to distribute the Unicast ABR content. A new class of service must be configured in the network to prioritize Unicast IP ABR and the BO's signaling.	Linear TV requires IP transport on the IP Core and Network, Media is transmitted at Layer 3 (no more Layer 2 connectivity between HE and Access) Multicast routing protocol like PIM (or other) must be enabled in the IP Core and Backbone that connects the HE with the Access. A new class of service must be configured in the network to prioritize Linear IP Video over Multicast.
Access	Typically EQAM is used.	No changes for DTV. Specific DOCSIS Service Flows could be configured to assure the video ABR QoS.	Linear TV over MPTS signals modulated with EQAM is changed by Multicast IP SPTS transported over DOCSIS or GPON. Multicast must be enabled in DOCSIS or GPON.
Home Network	Linear TV or VoD are consumed in Legacy DTV STBs	A new Hybrid STB is required. Hyb. STB installation is the same that Legacy STB. New unmanaged devices can be used to consume the video services through the regular EROUTER that also provides HSI service. More than one CM in the HN, one for the EROUTER, another for every STBs.	A new IP STB is required. There are specific functions to enable or to develop in the ROUTER/ONT/RGW. Multicast must be managed in the HN. New installation scheme, ethernet, WiFi, WiFi extension. A new process of logistic, provisioning must be developed in the operator. The concept of COAX outlet connection disappears from the HN. HN management is a MUST (TR-069). There are no requirements to split the coax, the CM represents a unique Point of Entry and the HFC network termination.

Now it will go through the changes on HE, BBIP, and it will describe with bit more detail the Access Networks and Home network¹¹.

5.1. Video Head End

Starting at the Headend it is necessary to introduce encoders that could generate the video signals in appropriate formats, contrary to the DTV signals where the video uses statistical video multiplexing optimized for QAM to save resources, in IPTV system normally requires Constant Bit Rate (CBR) video. Another change that, could be introduced because the state of the art of the encoding technologies, is the possibility to use High-Efficiency Video Coding (HEVC or H265) which allow to save up to 50% of bit rate to encode video with same quality as H264. Taking advantage of the change a new concept in the head end architecture is also introduced, that is what it is called Unified Headend.

¹¹ BO also require some changes but that are not going to be described here, but briefly those changes are to support of video's URI with multicast address, and the IP Conditional Access (CA) integration.

Figure 26 depicts the HE Architecture today, where exists several encoding silos isolated. With this architecture every time that a new encoding technology is added the ecosystem adds a new HE system silo, so there are silos for MPEG-2, silos for H264, another for ABR transcoding, and then it will be necessary to add other for IPTV.

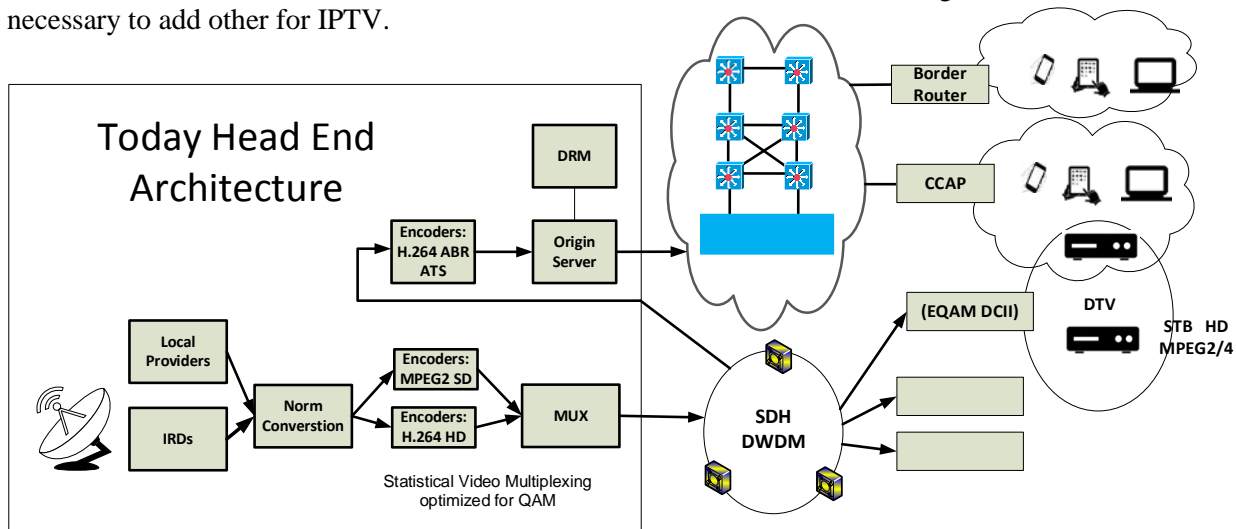


Figure 26 – Today HE Architecture

This scheme has some disadvantages, for instances:

- Every silo has its own infrastructure, a signal is carried from the source (Content Providers) up to every encoder, so there are Serial Digital Interface (SDI) cables between the origin of the signal and every encoder of every silo that must process the signal. When a new signal is added, then an operative task for cabling the new signal is required, the signal is connected to an expensive matrix and then from here cables for each silo. That is showed in the Figure 27 on the left.
- Each silo has their own roadmap of technology, and sometimes some features are available only in an HE system but not in all.
- The redundancy is local in each silo, so the same signal is redundant several times.
- In some cases, different silos are managed for different teams.

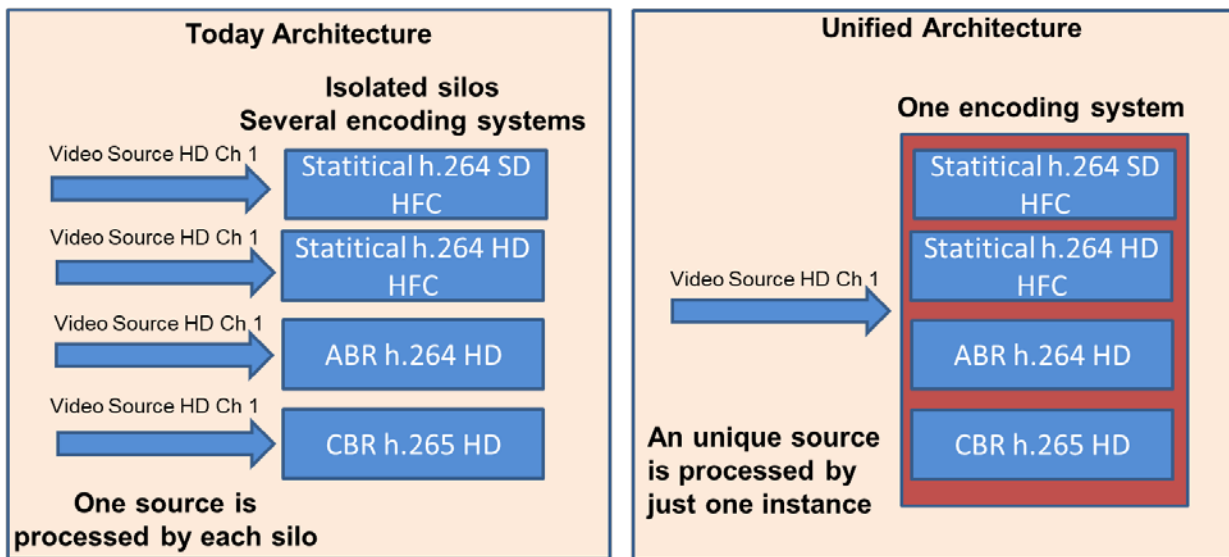


Figure 27- Several encoding systems vs Unified encoding system

The main idea of the Unified Headend is to have just one encoding system with several compression instances, then the system receives only one copy of the signal and it is proceeded by the correct instance. The different instances are virtualized and run over Commercial Off the Shelf (COTS) servers, the input could be IP or SDI, and depending on the server capabilities it is possible to run an instance in a server and that can deliver all the required formats. The advantages of this architecture are:

- This reduces the necessary infrastructure.
- Reduce the points of failures.
- The backup systems are shared.
- Improvement on the encoding algorithm could be deployed change the software version.
- Add a new signal means deploy a new virtualized instance, which is very flexible.
- Prepares the infrastructure towards a migration from the headend to the datacenter.

Figure 28 shows how the Unified Architecture is integrated into the ecosystem. The different outputs of the system are used to feed the different devices, on different networks. For example, one instance could deliver:

- The H.265 profile is used for IP STB connected in managed networks, the output of this instance goes through the scrambler (SCR) that together with the Conditional Access (CA) it is used to encrypt the video for IP STBs.
- The Multiple Bit Rate profiles in H264/H265 that generate the ATS feed for the OS, and the go through the CDN when a managed or unmanaged device get a channel in unicast distribution.
- MPEG2 or H264 Profiles used for DTV distribution like DVB, in cases of Analog reclamation (DVB EQAM) or for Analog modulators, in cases where it is necessary to provide analog TV.
-

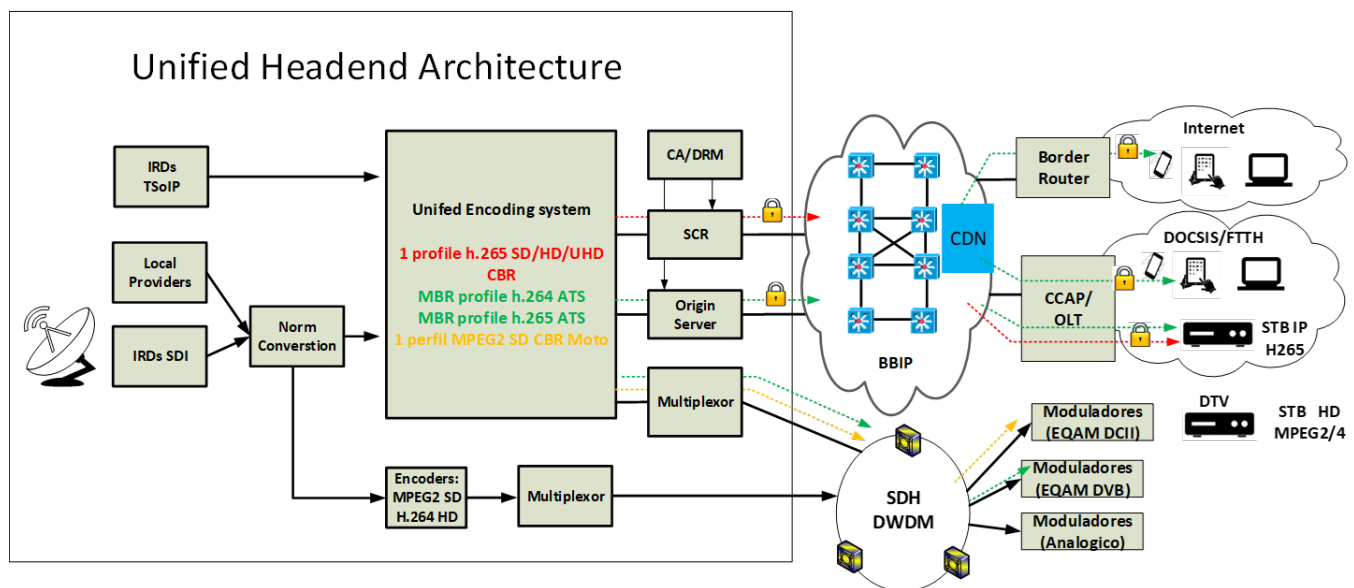


Figure 28 – Unified HE Architecture

5.2. Transport network

The signals that are generated from the HE must be transported to the access networks, for unicast distribution it was aforementioned how that is done in a scalable way using a CDN. But now we have to add a mechanism to forward the most popular Linear TV channels using Multicast.

Figure 29 shows how the Multicast mechanism works, it is based on a concept of groups of receivers (in this case the IP STBs) that are interested in receiving the same information, specifically here the video stream, the information is addressed with a Multicast IP, that comes from a special segment of IPs defined by the Internet Assigned Number Authority (IANA) in the range from 224.0.0.0 through 239.255.255.255 (9). These addresses are used as destination address in IP datagram, but a difference of other IP address there is not hosted on the network that has those address, these could be seen as “the address of the information” that the receiver wants to get, in our case it will represent a specific channel, and it will be represented as “Gi”. On the other hand, the information is originated from different sources, in these cases, they are the encoders or the scramblers, that are represented in an IP network by an IP Address “Sj”.

IP MULTICAST

Multicast (Group=Chi, Source=Sj)

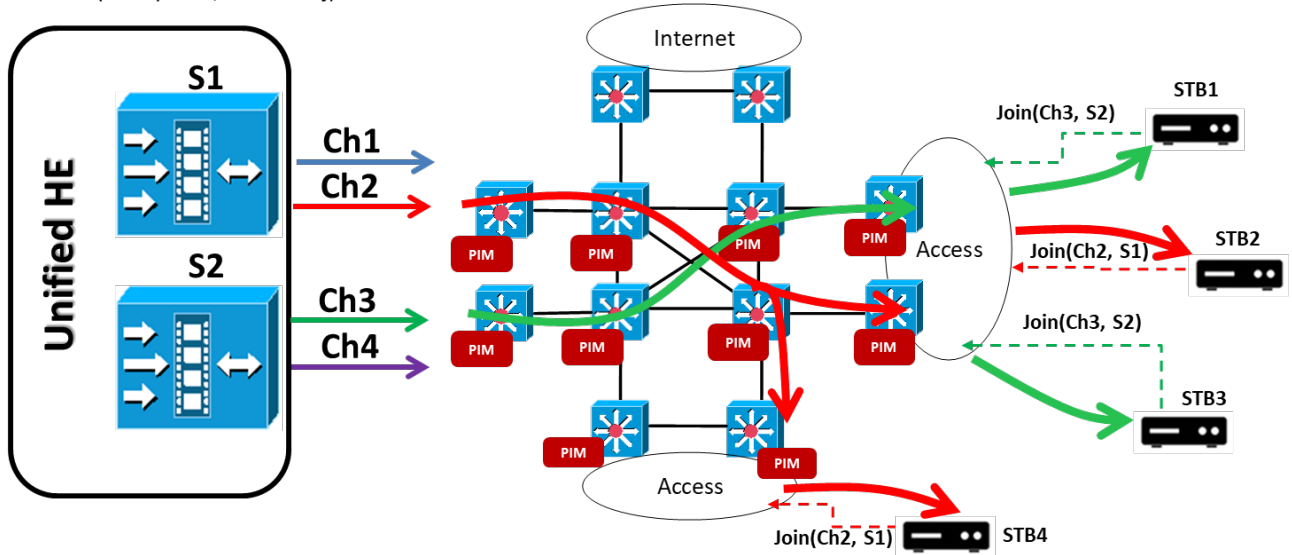


Figure 29 – Multicast in the transport network.

When an STB need to tune a channel, it must register to the multicast group G_i that carry this information. STBs use Internet Group Management Protocol (IGMP) which allows to a host ask for a membership to a local multicast router, this communication is done in a local network, at Layer 2. Local routers listen to IGMP information and sent out, with a given frequency, queries to host (in this case the STBs) that are connected in its local network, to discover which are the multicast groups that are active or not on that Layer 2 network. There are different versions of IGMP, but in this case, the IGMP version 3¹² (10) is used. Let's suppose that an STB1 tunes the Ch3, then it sends a Join IGMPv3 message where it informs to the local router that it wants to register in the multicast group $G = \text{Ch3}$ ¹³ (which is an IP multicast address). When the local router receives this Join and realizes that it is not forwarding this multicast address, then it must request to the source the information.

In unicast routing, the forwarding's tables in the routers are based on the network's destinations, and those networks are manually configured or automatically learned with routing protocols. For multicast routing, the whole network must be configured to indicate to routers how to generate the forwarding tree between the source and the receivers. There are different options for that but here Protocol Independent Multicast (PIM) is used, and particularly when the source S_j of the information is known (that is the case of IGMPv3) PIM Source Specific Multicast (PIM-SSM). With PIM-SSM if local router receives a subscription to a given (G_i, S_j) , then it forward that request to the PIM neighbors routers using the normal routing tables based on the S_j until that the request reaches the router that is connected to the source S_j , with this mechanism a route tree is built between the routers that were asked for the G_i and then the information is forwarded through the interfaces of those routers, finally reach the local router who forward the traffic to its local network and the STB1 receive the information. When a second STB (STB3 in Figure 29) tunes the same channel Ch3, it sends the Join to the local router but in this case, the

¹² Unlike of IGMPv1 or v2, IGMPv3 must define which is the IP address of the source where the information must be requested.

¹³ The video is carried on SPTS and is transmitted using UDP, or RTP over UDP, then the client also informs the UDP port.

multicast has been already forwarded in the local network and the STB3 simply read the stream from the connected network.

When an STB change the channel, what happens is that firstly the STB sends a Leave message, which indicates to the local router that this specific receiver doesn't want to receive the information, and then the STB generates again a new Join with the new (Gi-new, Sj-new) for the new channel. When the last STB in the local network sends the leave for a given (Gi, Sj) the local router stops forwarding traffic of that multicast group toward the local networks. Eventually, that process continues in the rest of the network's routers and if there are no other receiver registered to this (Gi, Sj) the multicast groups it is not forwarded in the network until next Join.

5.3. Access Network

Once that multicast traffic reaches the access network the way in which every type of access network manages this traffic could be different. Multicast traffic in Layer 2 networks are mainly specified in Ethernet networks, but here the video must be distributed on DOCSIS or GPON14, which define a specific mechanism to forward the multicast traffic. Unlike DSL, DOCSIS uses a shared medium but provides tools to ensure Quality of Service (QoS) to each subscriber. In this section, it is going explain how DOCSIS support multicast traffic comparing the hybrid with Full IP video in terms of QoS and Service Flows (SF), and finally some issues that could arise with Multicast in DOCSIS.

5.3.1. Multicast in DOCSIS.

In DOCSIS® 3.0 (D3.0) MAC and upper layer specification (11) defines a flexible scheme for multicast to support Any Source Multicast (ASM) (9) and Source Specific Multicast (SSM) (10), as aforementioned the last one is what it used in this paper to distribute a broadcast linear TV services. Unlike previous DOCSIS® versions, where it works based on snooping IGMP messages (only IGMP v1 and v2) at CM level, in D3.0 the CM is not aware of the multicast, the responsibility of the Multicast control over DOCSIS, relies on the CMTS. That imposes the first constraint that is that the deployed services architecture requires “as a MUST” D3.0 in both sides HE and Home Network, it is possible to have CM < D3.0 connected in the same SG but, they must not be used to deploy Linear TV based on Multicast¹⁵.

DOCSIS® 3.0 introduces the Downstream Service Identifier (DSID), this is a variable of 20-bit length that is embedded in the Downstream Service Extended Header of all packets that belong to a downstream bonded service flow and/or a multicast group service flow. The CMTS tags all the multicast packets with destination to same multicast group with a same value in the Downstream Services ID (DSID), from CMs point of view, they use the same tag to enable or not the multicast forwarding. This is a variable of 20-bit length that is embedded in the Downstream Service Extended Header. So, for multicast forwarding, in the CMTS MUST be configured the MAC domain as “Multicast DSID Forwarding” (MDF). The CM declares itself that it has MDF capabilities in the registration request messages, and the CMTS answers that the MAC domain, where this CM is connected, has MDF configured using the registration response message.

¹⁴ The scope of this document does not include explanation of multicast over GPON because the limitations on the document length.

¹⁵ Nowadays the networks of the company are completely in D3.0 and it partially updated to D3.1, D2.0 CM are being withdrawn.

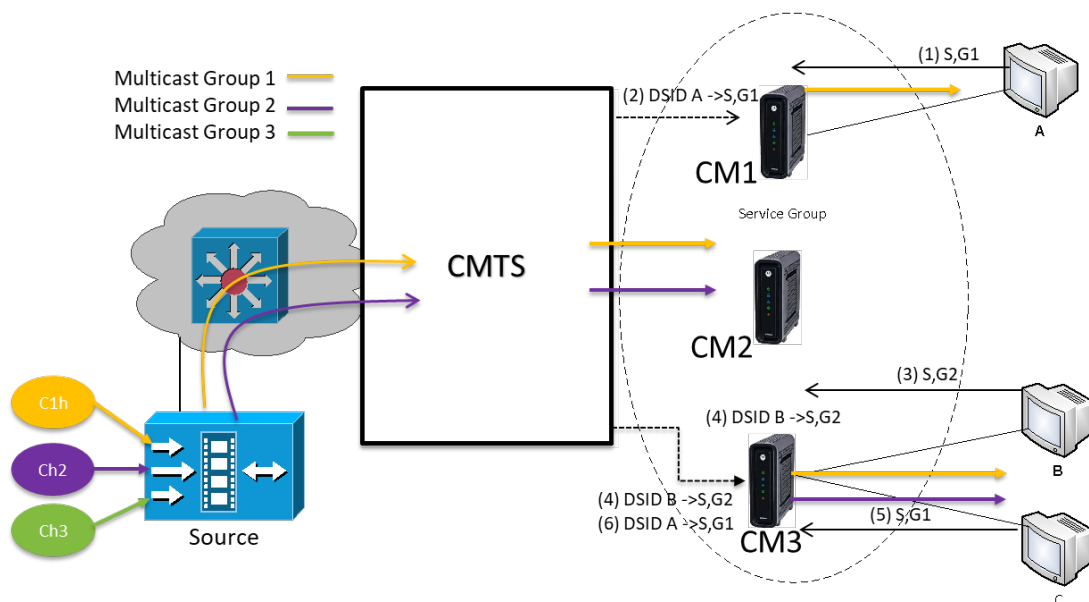


Figure 30 – Example of Multicast Forwarding in DOCSIS 3.0.

Multicast client protocol signaling (IGMPv3) is used for the clients (in this case the STBs) to join and to leave dynamically to the multicast group streams (Multicast Group Session).

The CMTS replicates the multicast traffic in the MAC domain in the set of channels that are called Downstream Channel Set (DCS) which is an identifier of a single Downstream Channel or Downstream Bonding Group (DBG). The CMTS replicates the Multicast traffic, which receives at the network side, towards the MAC Domain using those DCS, obviously, in D3.0 they are based in DBG.

In the Figure 30 a home device (which could be an IP STB) subscribes the channel 1 (Ch1) using a Join to (S, G1), this message goes through the CM transparently and it is received by the CMTS. If the CMTS does not have the multicast traffic, it requests it to the network, then CMTS will receive the multicast traffic and it will replicate this traffic to the DOCSIS MAC domain (cable side) in a DCS (which is a DBG) using the DSID A, following this, CMTS will signal to the CM1 with DSID A and add the client mac address, using a Downstream Bonding Change Request (DBC-REQ) to inform it to read the DBG where the multicast traffic is replicated, and then CM1 forwards the traffic (S,G1) to end device that had requested originally the channel 1. It can see that the CM1 use the DSID as a filter to know which multicast traffic must forward or not to its Ethernet LAN ports.

Following with the Figure 30 explanation, when a devices B connected to the CM 3 asks for a channel 2, the same process happens (it will use a new DSID B) and the device B receives the multicast G2. At this point, in the DOCSIS mac domain there are two multicast streams present, Finally the device C connected to the same CM3 and it subscribes the Ch1, but it had been requested by device A in CM 1, so the multicast is already in the MAC domain, thus in this case, the CMTS just inform the DSID A to CM3, the client mac address and then it forwards the traffic to its LAN.

If some device, for example Device C, change the TV channel, this results in the CM 3 LEAVE the old Ch1 and JOIN the new channel. Then, again CMTS signals to the CM 3 using a DOCSIS MAC Management Dynamic Bonding Change transaction to create/change/delete the parameters necessities for the CM to Leave the Ch1 and to join the new channel. Is important to note that is necessary for the CM

always generates the "LEAVE" message because the CM has a limited number of bonded multicast DSID. The CMTS oversees signaling and add/change or delete the DSID and the client MAC.

5.3.2. Multicast Channels – Dedicated, Shared or combination of both.

The DOCSIS channels (QAMs¹⁶) used to transport the multicast traffic can be dedicated or shared by another kind of traffic. Figure 31 (A) and (B) show scenarios for dedicated and share QAMs respectively.

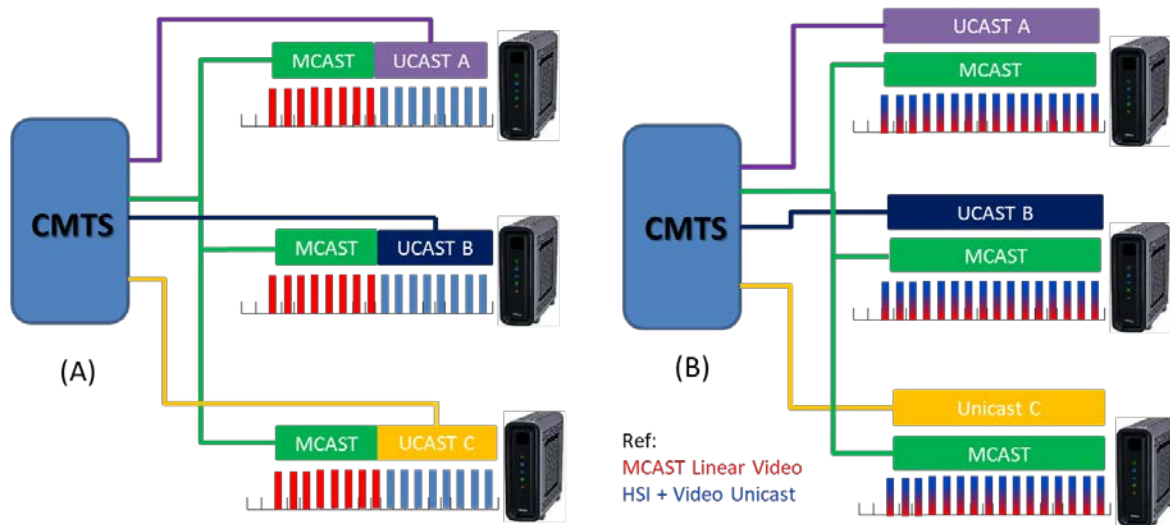


Figure 31 – Dedicated or Share QAMs for Multicast

The advantages of dedicated QAMs is that in those scenarios the Linear Video traffic in multicast is physically isolated from the rest of the other services, that could simplify the capacity planning of the services group associated with multicast, because the only traffic that they carrier is the multicast traffic for the Linear video (of most popular channels), and thus it be able to calculate and track the usage of those recourses in a simple way. In the same way, other advantages that could appear in this case is that, because the services are very well divided into different QAMs, it could be: the troubleshooting, the monitoring of those resources.

Contrarily, the advantages of shared QAMs is that the resources in term of bandwidth on the service group are better utilized during the beginning of the IPTV deployment. The video traffic will not be too much, let's remember that multicast video traffic it is a function of the number of channels in the lineup but also the number of the devices. Thus, to use dedicated carriers for multicast would mean a waste of the available spectrum. On the other side, because the statistical behavior of the data, the most efficient architecture is to share the channels between unicast and multicast traffic. If the sizing of the network is done in a correct way, resources are better utilized. But it is not ruled out to use dedicated carriers for a future, when the traffic will be greater and there is a better knowledge regarding planning.

The architecture's recommendations is to use a combination of both as it can see in next point.

¹⁶ The document uses the term of DOCSIS channel or QAM indistinctly.

5.3.3. CM Capabilities and channel assignment to CMs.

Regarding the configuration, the way to handle multicast video and data through the carriers is by configuring the DOCSIS service flow attribute mask¹⁷ in the corresponding downstream and bonding groups. Each service flow creation request attribute masks (Required or Forbidden) may be defined as part of a provisioned service class on the CMTS. Every channel and bonding groups have a resource attribute mask denoting the capabilities of that resource. This resource attribute mask is compared to the attribute masks of the service flow request provisioned in the service class. The CMTS attempts to assign service flows to channels or bonding groups such that all required attributes are present and no forbidden attributes are present, doing that is possible to match what kind of traffic must be carried by which downstream groups of channels (DBC).

Besides this, there are different kind of CM with a different number of tuners already deployed in the network: D3.0 4x4, 8x8, 16x4, 24x4, 24x8 and D3.1 32x8 + 1 OFMD Block. This diversity could generate an extra issue for capacity planning because depending on the combination of the amount of the channels used for the DOCSIS SG and the CM's capabilities (number of tuners) it could be necessary to replicate several times the same multicast group traffic in the same SG. Figure 32 depicts this case where it is necessary to generate 3 different multicast streams that are using a different DCS in the same SG, which means a decrement in the multicast gain.




Figure 32 – Decrement of Multicast Gain because of diversity in CM capabilities.

It must be mentioned that inside the home network, as it will discuss in the next section, the Customer Premise Equipment (CPE) is no simple CMs, but it is EROUTERS, which means that it embeds in the same CPE a CM but also Residential Gateway (RGW) functions that must meet with multicast distribution inside the Home. The last ones have to be prepared to accept and distribute the multicast data inside the home network. Thus, even when the D3.0 CM is used by the EROUTER could manage the multicast according to the specification, the internal RGW could not work properly with multicast. This point could be fixed with software upgrades, but if apart from that it is necessary to have strong WiFi capabilities to distribute the video, it is very probably that old EROUTERS, already deployed in the field, could not support the video service distribution as a managed service.

¹⁷ This paper (26) explains in the detail the usage of this attribute

Given this point and the aforementioned issue regarding the diversity of CM capabilities, the architecture decision was to deploy the IPTV services using the last EROUTERS generation and the next ones, with D3.0 CMs 24x4/8 already in deployment or next ones D3.1 CMs 32x8 + N x OFDM blocks¹⁸.

Another definition is that to provide IPTV services the SG size must be at least 24 D3.0 channels. However, in the future will be possible to extend the SG to 32 D3.0 channels, and to add D3.1 OFDM blocks. When that will happen, it will have to be careful with the already deployed CMs with 24 SC-QAM. Because the SG will have 32 SC-QAM, the CMs with 24 SC-QAM will have more than one set of QAM to register. In this scenario, to avoid having multiple multicast streams in the same SG, we MUST use static bonding group.

To maximize the multicast benefit (and minimize bandwidth) on the service group, the decision is to distribute the multicast on a set of 24 Single Carrier-QAM (SC-QAM) fixed in the spectrum but sharing this QAMs with the other services (HSI, telephony and Unicast Video). The way to define which are the 24 channels that will transport the multicast traffic is based in the Attribute Mask defined in DOCSIS (11).

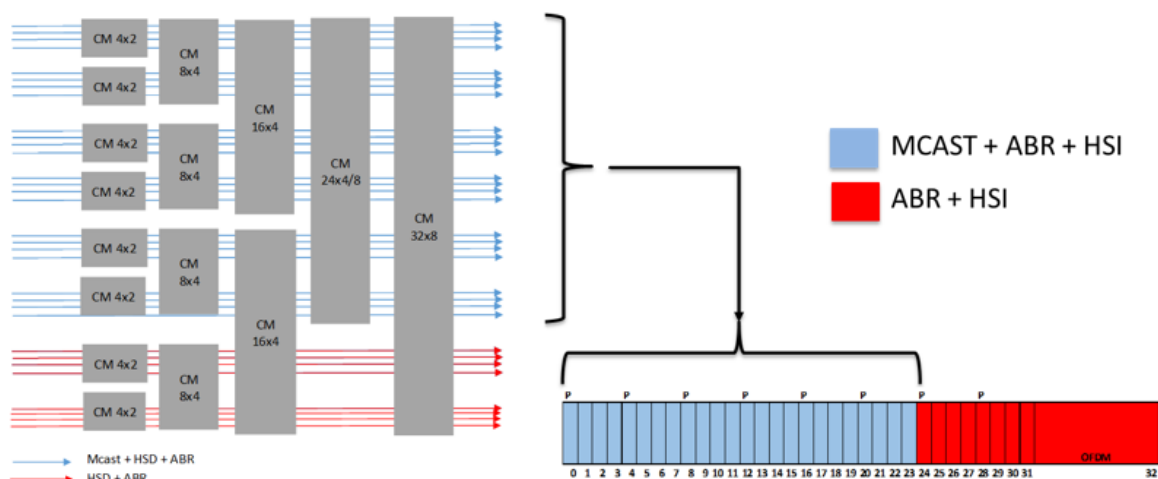


Figure 33 – Multicast distribution in the Service Group

As it was mentioned, when the SG is extended from 24 to 32 channels it is possible that 24x4/8 CMs could have different combinations of channel bonding (same happens with 4x2, 8x4, and 16x4). To assure which must be the group of channels in the SG, that must be used to connect CM with different capabilities, the bonding groups must be defined statically, and all the static bonding group with more than 24 QAMs MUST include the multicast bonding group.

Thus, defining where the CMs must bond the channels, with the static bonding configurations, and establishing which must be the channels that carriers' multicast using the Attribute Mask, the multicast traffic will be distributed across only one possible set of channels in the SG for the D3.0 24x4/8 or D3.1 CMs.

Figure 33 shows this configuration. It depicts the distinct set of bonding arrangements, where modems with different capabilities will use a diverse set of those BG. Particularly there is only one BG of 24 SC-QAM

¹⁸ In this deployment, CMs (and home gateways) that are used for the IP video service MUST be capable of receiving at least twenty-four SC-QAM downstream channels (independently of how many channels are used for multicast traffic).

where the D3.0 24x4/8 CMs can be registered. The Linear TV multicast is enforced to use the same channels that are based on the Attribute Mask but they are also used for HSI and Unicast ABR. Contrarily the additional capacity (the red one) must be used just for unicast: HSI, ABR video, telephony or other. The first 24 channels (the blue ones) are shared form Multicast and Unicast while the other 8 channels plus the OFDM block (the red ones) are used only for Unicast.

The picture shows channels with a “P” indicator, that is because those SC-QAM are primary channels, that are necessary because in the network there are CMs with 4x2 capabilities, so it must have one primary channel in each bonding group. Also, even today in our network there are some DOCSIS 2.0 CMs, and it needs a primary channel for those CMs. Using secondary channels instead of the primary could save some Mbps capacity per channel.

5.3.4. Uncorrected error and Partial Services - Multicast Resiliency in DOCSIS Networks

Figure 34 shows an SG with CMs that are consuming regular unicast data and joining to a multicast group. If, for example, there’s a problem in CM 2 with 2 QAM channels, it reports the problem to the CMTS and the CMTS declares Partial Service (PS) for this CM.

The CMTS sends the Unicast Traffic spreading the unicast data (that is different for each CM) through all QAM, except for the CM 2. In this case, the CMTS sends the data for this particular CM avoiding using those channels that for this CM have an issue.

However, for Multicast Traffic, the CMTS continues sending the data through all QAM channels (because the Multicast traffic for this Multicast group is the same for all CMs connected to this SG), so the portion of data carried by the QAMs, that CM2 is detecting errors, has uncorrectable errors and then the application that uses that data doesn't work (video is broken).

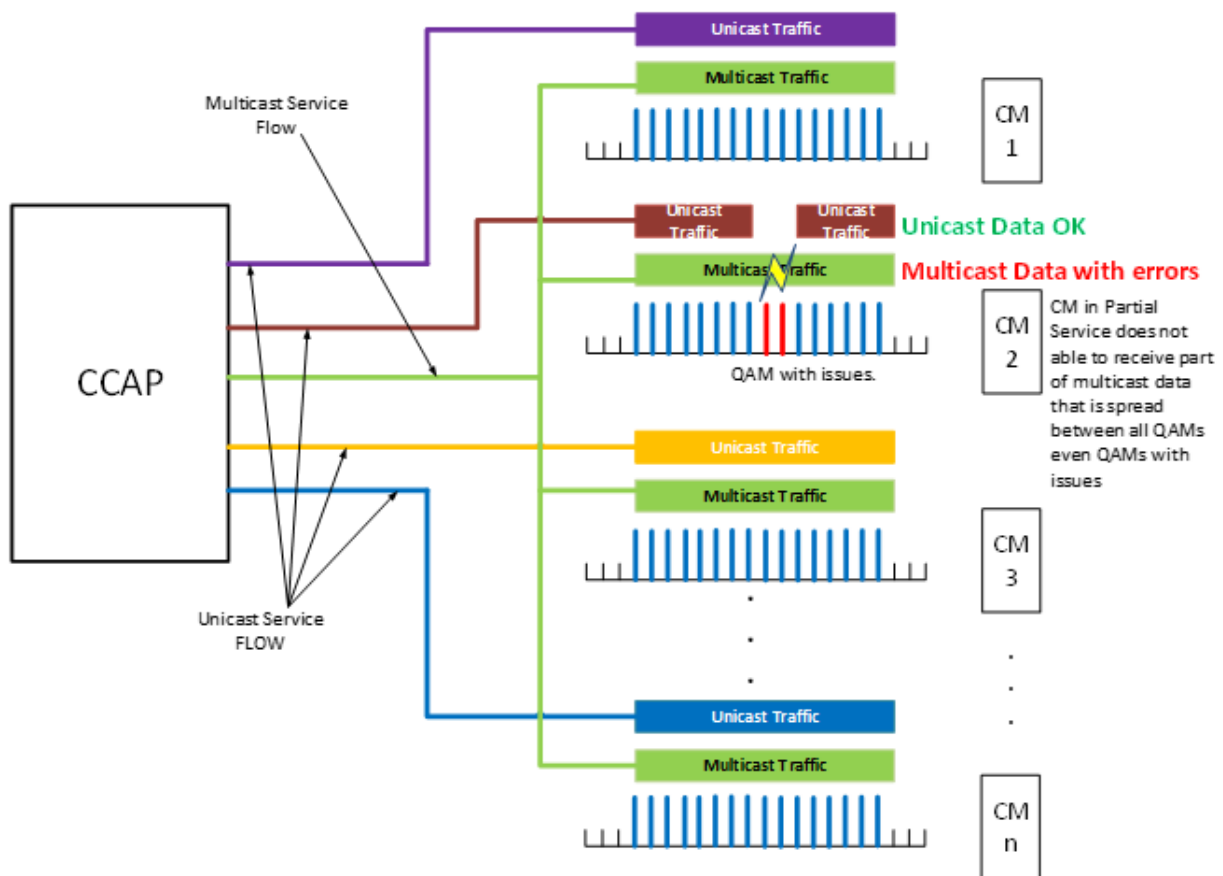


Figure 34 - Multicast Traffic issues due to Partial Service

Figure 35 is a diagram that shows the sum of CM's channels with a Partial Service mode declared in the CMTS for the whole network, a give CM could have Partial Service in more than one channel. There are around 17% of CM with PS, which means that all of them will have problems with video over multicast. The root cause is ingress noise that is generating in home networks but also in HFC network, for instance, frequencies in the range of 759-777 MHz come from mobile networks (LTE).




Figure 35 – CM’s channels in Partial Service Mode per frequency¹⁹.

If it avoids using the channels in the range 759-777 MHz, which are the most affected, then it could reduce the CM’s affected up to 6%. To do that it is necessary to reconfigure the arrange of channels used for Multicast used in the Figure 33 and reconfigure according to Figure 36.




Figure 36 – Reconfiguration of channels for Multicast to reduce Video Issues per PS.

However, even when the number of CMs with these problems are low in relative numbers, in absolute numbers means a big amount of client with video issues in the most popular channels distributed with multicast. It must be considered that in legacy digital TV, if some frequency has problems (noise, etc.), just the signals carried by this frequency are affected, but here in IP multicast, all the signals will have problems because the data is spread over several carriers. Again, those ingress noise must be fixed with

¹⁹ Note that these statistic numbers are the best case of the issue presented here, since there are many CMs that don’t an entry in partial service mode and still have the problem, that simply happens when there are uncorrected errors in some carriers.

maintenance, but even doing those operation tasks, the ingress noise could appear in any moment and in any part of the spectrum/network. This case should be considered not only from operational but also from a technological point of view, to do multicast in DOCSIS more resilient against these kinds of problems. The solution/s must be dynamic avoiding the necessity of operator's intervention and preventing the most as possible a complete blackout in the lineup for those clients behind a CM in PS mode. Following, four different mechanisms to reduce this effect.

Option 1: Multicast Resiliency by Capacity reduction of Multicast SG.

This method decreases the capacity of Multicast SG to avoid issues in CMs with Partial Service Mode. When a given number of CMs in a service group enter in partial service mode for a QAM channel, the CMTS stop sending multicast in that QAM channel. The number of CMs MUST be configurable starting at 1.

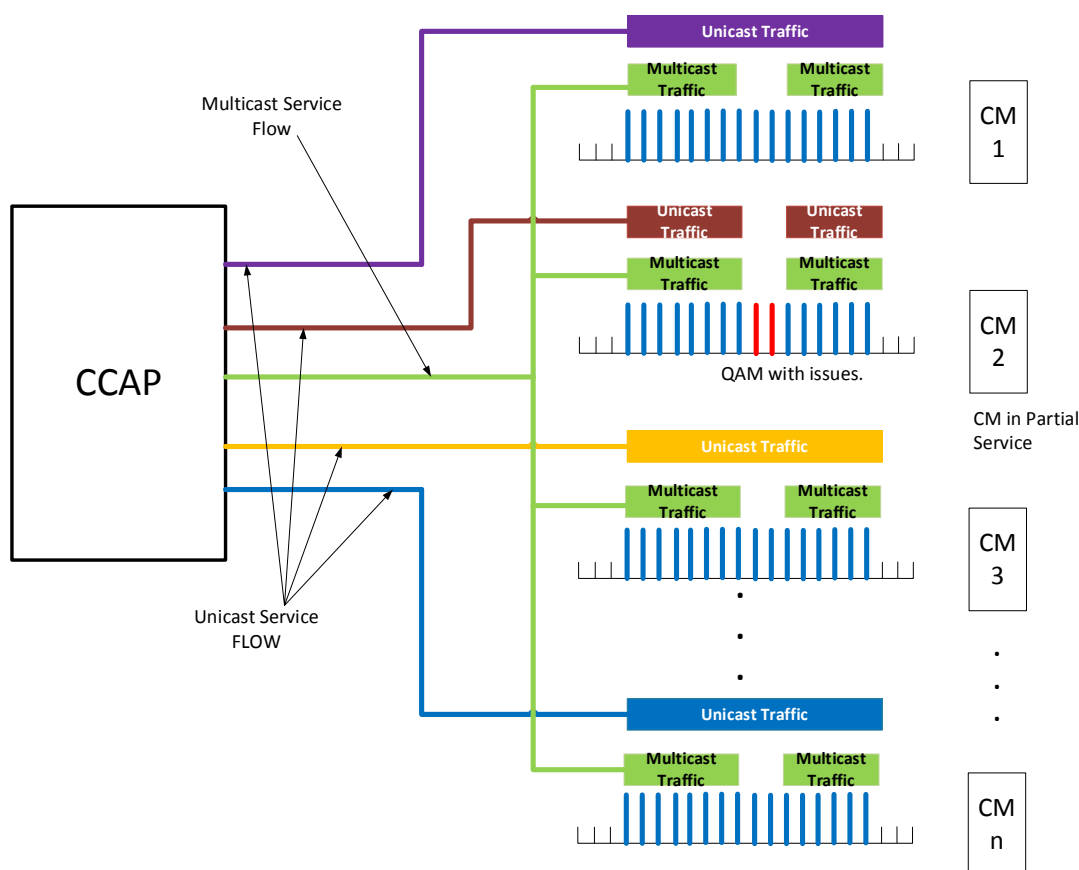


Figure 37- Multicast Resiliency by Capacity reduction of Multicast SG.

CMTS is who decides what QAM channel must be excluded or included dynamically in the Multicast SG based on the settings. The decision to include/exclude a given QAM channel in/from a Multicast SG depend on the amount of CM with Partial Service because of this particular QAM channel. The maximum number of QAM channels that can be eliminated from the “multicast bonding group” MUST also be configurable to avoid congestion problems in the SG because of the QAM reduction. There must be a configurable threshold level/timer before the decision to exclude a QAM channel from the Multicast SG. There must be a configurable threshold level/timer after the decision to include a QAM channel again in the Multicast SG.

Option 2: Multicast Resiliency by Multicast to Unicast in CM with Partial Service

This method does not modify the Multicast SG capacity for CM without problems, but for CMs that are in Partial Service, the CMTS will send the same data that should be carried using Multicast but using Unicast. This mechanism could affect the SG sizing.

When a given CM enter in partial service mode for any QAM channel, the CMTS stops sending multicast for that CM and convert the Multicast traffic that came from network side interface to Unicast in DOCSIS SG.

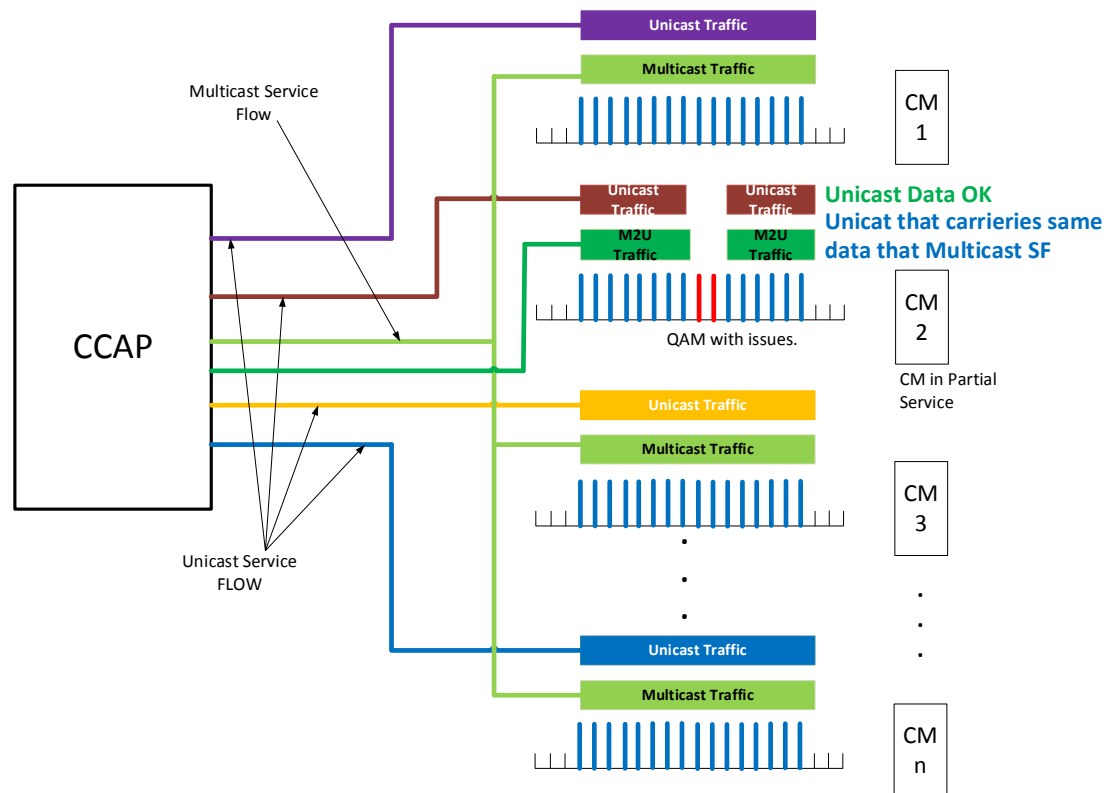


Figure 38 - Multicast Resiliency by Multicast to Unicast in CM with Partial Service.

CMTS is who decides what CM will receive a Multicast 2 Unicast (M2U) traffic instead of Multicast. There must be a configurable limit of CMs that could receive M2U traffic for a given SG. This is to avoid the congestion of an SG because of an excessive amount of CMs that are in PS and then a huge amount of M2U traffic that generates a congestion of the SG which sizing was done considering Multicast traffic.

It should be also possible to configure at the CMTS level that for a given Multicast Group, that, for instance, transports high audience video signals, the CMTS never generate the M2U mechanism.

The decision to change from Multicast to M2U is just the based on the Partial Service mode detection for a given CM. There MUST be a configurable timer before deciding the change to M2U.

The CMTS should come back from M2U to Multicast when for a specific CM that was working in Partial Service Mode it works again normally. Before deciding to go to Multicast it is necessary to wait sometime, this timer MUST be configurable in the CMTS.

Option 3: Multicast Resiliency Hybrid

To improve the amount of networks resources used we can combine Option 1 with Option 2. For a configurable number of CMs the CMTS uses M2U (option 2) and if that number is exceeded the CMTS starts forwarding the traffic to all CM in a reduced bonding using multicast as is explained in option 1.

Option 4: Multicast to Unicast based on STB and BO interaction

In this case the solution it is not based on the DOCSIS itself but is based on the interaction between the STB and the BO of the system. As was mentioned at the beginning of this section, Unicast traffic is not affected by the PS because DOCSIS manage that in a perfect way. So, it could be possible to generate a mechanism to detect the video impairments on the STB and based on that to request a unicast ABR video instead of multicast. Because the amount of the CMs with PS in an SG is small, it is possible to use Unicast ABR instead of Multicast for those CMs with Issues. This mechanic must be generated at control plane level between the STB and BO and could be a proprietary solution.

5.3.5. Video QoS for Unicast and Multicast.

DOCSIS® uses a shared medium but it provides tools to ensure QoS. To differentiate the QoS of video services over DOCSIS the ecosystem must be configured to carrier those data over specifics Service Flows (SFs). That configuration must consider the QoS not only for the media and for the control plane, and the last one includes the communication with the BO. The Figure 37 shows conceptually which are the different kind of traffic that must be managed. They are:

1. Unicast Video to connect Managed STB with CDN, over this it is delivered CoD and lowest popular Linear TV signals.
2. Multicast Video which carrier the most popular Linear programs to STBs.
3. BO control that is used to manage and signal the video services in any device.
4. HSI data access that is used to connect the general Home network devices with internet. This is also used to carry the video the Unmanaged Devices like Tablet, Smartphones, PC, etc., in this case even when the video come from the CDN, the product business case indicates that those devices must share the resources with HSI access.
5. IGMP message generated by the IP STBs (they are not depicted in Figure 39, to see Figure 40)

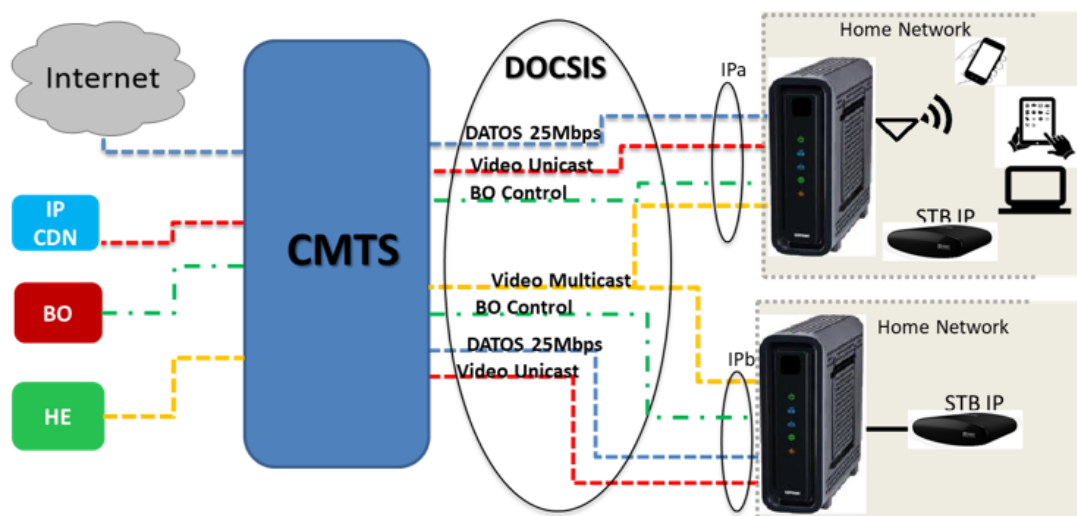


Figure 39 - Different kinds of traffic in DOCSIS

The Unicast Services Flows are based on DOCSIS Classifiers that are configured in the configuration file of CM. The Service Flow provide the QoS Parameters for treatment of those packets on DOCSIS. A Classifier is a set of criteria applied to each packet, which consists of some packet matching criteria, (destination IP address, for example) and a classifier priority. If a packet matches the specified packet matching criteria, it is then delivered on the referenced service flow. The classifiers could be based on Source: MAC Address, IP, Port; Destination: MAC Address, IP, Port; protocol, IP TOS/DSCP, 802.1Q VLAN ID.

It could use classifiers base on IP address, but if for instance the CDN is distributed and the IP address of those elements are not in the same IP segment, the configuration file must have several classifiers, and then it could be more complex in terms of operative and maintenance tasks, because every time that a new CDN node is added in the system the Classifiers must be modified. On the other hand, to use IPs and Port as filters in the classifiers do not allow to differentiate the Video Unicast traffic for Managed (STBs) and Unmanaged (OTT) devices, and the rule number 4 aforementioned is not fulfilled.

So instead of used IP:Port the classifier could be done based on DSCP, and then it is possible to use just one classifier per service flow and even when the operator will modify the CDN topology, or BO or HE configuration, the classifiers in the CM's configurations files are kept.

At the same time if the CDN could mark with different DSCP based on the User – Agents (of the devices) that generate the content request, then it should be possible to mark with a given DSCP the traffic from the CDN towards STBs and with other one towards OTT devices (with other or directly without DSCP mark). In this way, the traffic from CDN towards managed devices is carried on a specific service flow, but from the CDN to unmanaged devices is treated as a regular internet traffic and it goes in the HSI services flow.

It is important to note that the Service Flows are Unidirectional, and the classifiers must be configured in both ways, downstream and upstream. All of them have defined a set of QoS Parameter: all are Best-Effort Service Flows but they are differentiated between them with, Maximum Sustained Traffic Rate, Maximum Traffic Burst, Minimum Reserved Traffic Rate, and Traffic Priority, that are configured in each one. That can be appreciated in Figure 40 (the priority is showed as P: x in each service flow).

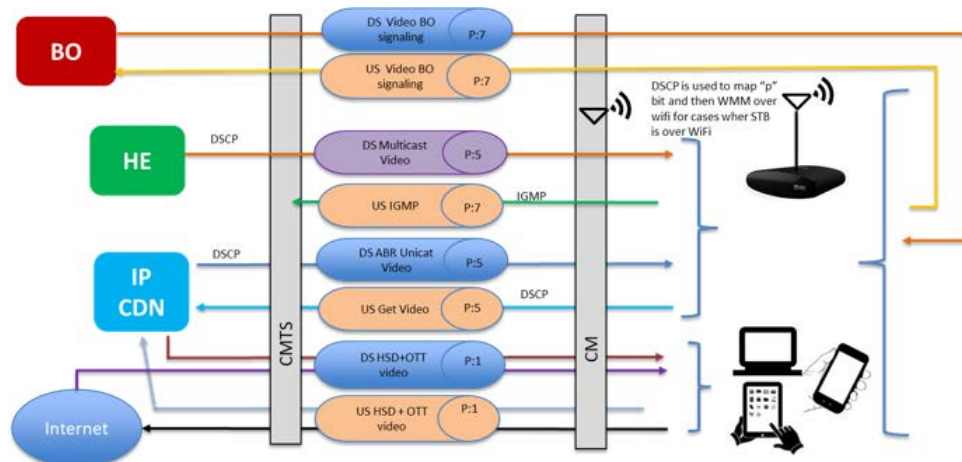


Figure 40 – Service Flow Configuration

It should be possible to combine some of those services flows, for instance in the Figure 40, US Get Video, US IGMP and US Video BO signaling, could be potentially managed in only one “Signaling” service flow, but the separation could be used to get statics on each service flow, for instance using IPDR, and because this is a starting deployment, it could be configured in this way and when the operator get experience then move to a simpler configuration.

Unlike Unicast, the Multicast traffic is not based on the CM’s configuration file but is based on Service Class Names (SCN) rules that are configured at CMTS level, that mandatory to work with Multicast as was required in this document. To set up QoS for multicast some definitions are required:

- Multicast Groups QoS: each multicast group segment address can be assigned at different Service Class Name, thus if it is required to apply different QoS parameters to a different kind of content, like SD, HD, UHD, (or other categorization), then it could be possible to represent those content types with diverse set of IP multicast segment address and then to apply different SCN. That is defined in a “container” that as is showed in Figure 41.
- Service Class Name: here is where the service flows are defined, once the Multicast Group QoS is defined, it is associated with SCN, also another set of QoS parameters are assigned:
 - Direction: Upstream (1) or Downstream (2) (for multicast is 2)
 - Attribute Mask: this is the value that is used to match with the DOCSIS channels defined in the BG which will carry the multicast traffic. The attribute mask can be: 1) required and in this case the CMTS will use BG with this bit marked it their associated channels attribute, or 2) forbidden which is the contrary.
 - Maximum Sustained Traffic Rate/ Maximum Traffic Burst / Minimum Reserved Traffic: this are the same parameters that must be configured in a service flows on CM’s configuration file, with those parameters is possible to configure the video’s bit rate profiles assigned at each type of video stream, for example 2Mbps for SD, 5Mbps for HD, 20Mbps for UHD (or other depending on the chosen criteria in the Multicast Groups QoS definition)
 - Priority: this the priority’s Service Flow parameter, that is used by the scheduler to give more (higher number) or less (lower number) priority to the traffic that is transported in this Service Flow when it compares with others.

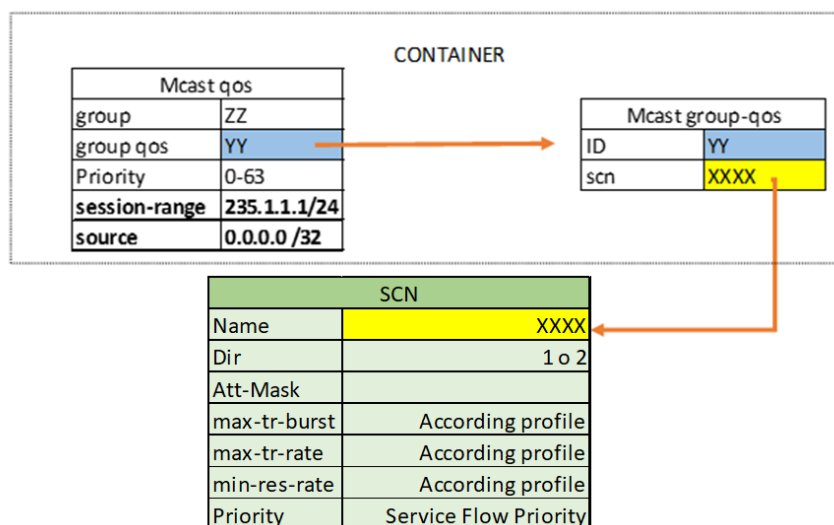


Figure 41 – SCN definition.

In order to protect the amount of multicast session that a given CM can subscribe it is recommendable to set maximum multicast sessions that a CM can join at the same time. The CMTS keeps the tracking of the multicast session for each client and limiting the number of multicast session behind a CM could prevent a denial of service attack. That is also part of the DOCSIS® specification (11).

Finally, in some CMTS it is possible to configure a Connection Admission Control (CAC) to limit the capacity used by multicast in the channels when this capability is configured just a percentage of the channel will be used for multicast. When the CMTS receives a JOIN message it is checked if there is enough capacity to allocate resources in the BG based on the Minimum Reserved Traffic, if there is no capacity then the new Service Flows it is not admitted.

5.4. Home Network

Finally, it is the Home Network (HN) domain, and this is maybe the hardest changes needed to move to Full IP Video, and not because of the technology itself, but because of the change of the paradigm for the operator and the end user.

Figure 42 compares the today model with the Hybrid HN (A) against the Full IP HN (B). The main difference is the change of the way that STB is connected. In the Hybrid HN, even when the UX experience could generate a disruptive change in the client's way to consume the service, because the new advance services, new UI, and others that were mentioned at the beginning of this document, the Hybrid STB is still installed and connected in the same way that a Legacy DTV STB. The technicians are still using the same tools and process, the connection inside of the HN is based on COAX. In the new Full IP HN the STB is connected in the same Ethernet network that is provided by the EROUTER's LAN, so the technician's "COAX's toolbox" does not work any longer. The IP STB is connected with Unshielded Twisted Pair (UTP) patch cords, WiFi, or even more through LAN device extensors, like Power Line Communications (PLC) or WiFi extenders.

So, the technician must add to their COAX's tools, Ethernet tools, UTP cable, with tools to check the WiFi connectivity and fix it in cases where the coverage it is not enough, the new devices as aforementioned PLC or WiFi extenders must be part of the devices in the technician's trucks just in case

of need in the installation. It could appear complications, for instance, the EROUTER has in general four Ethernet Ports, but all of them could be used at the installation moment, then the installation must be done using just WiFi, or connecting a local LAN Switch.

Those apparently simple facts imply new process in the installation, new skills for the technicians and several modifications in the MSO's Operations Support System and Business Support System. So, an IP STB installation is not a just a traditional TV services installation, it is really a Home Network Installation, wherein some cases it could be as simple as connect a STB to the WiFi HN, or it could be as complicated as whole Home Network installation to support the TV services.

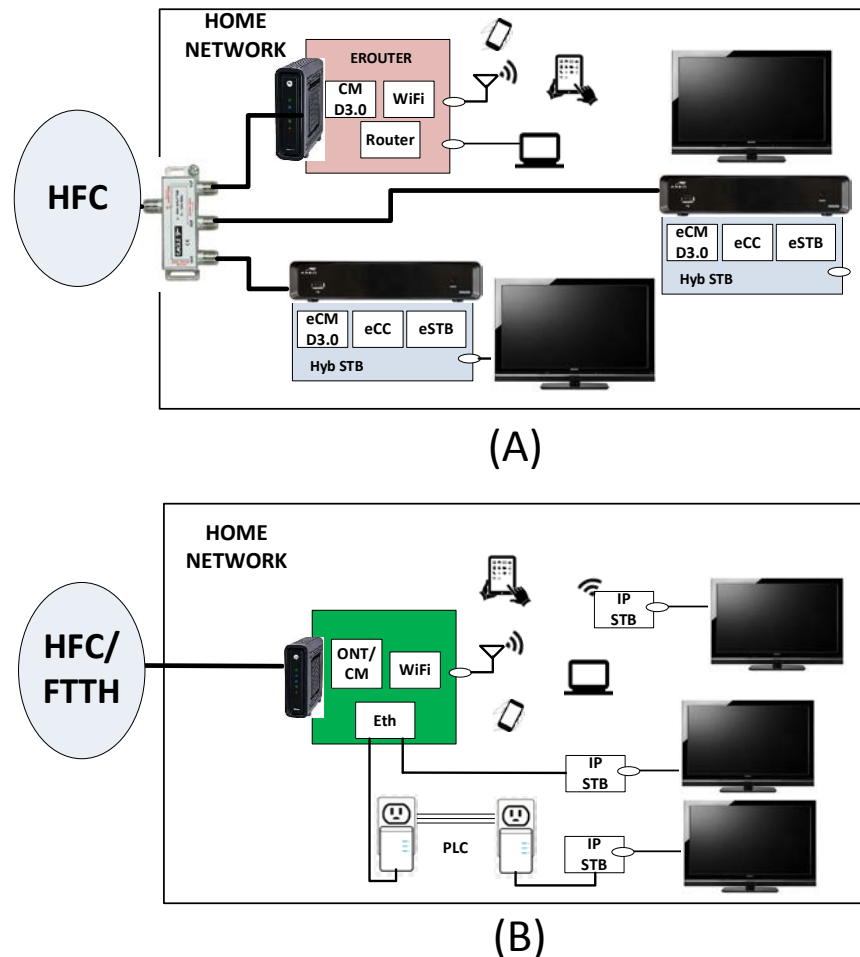


Figure 42 – Hybrid Home Network (A) vs Full IP Home Network (B).

Nevertheless, there is a lot of numbers of improvements and benefits, for the operator and the end client:

- The first one that it can be appreciated is that there is just one CM per home, then when in the Hybrid HN at least 2 CMs are required per home, one for the HSI and another one for at least 1 STB, so if every client in the DOCSIS SG has HSI and TV services with the Full IP home network the SG size is reduced in at least 50% of CM (in general there are 2.5 TV sets per home).
- At the same time there a cost reduction in terms of the STB's cost, the Hybrid STB has an embedded CM which means an extra cost per each TV.

- In the Hybrid HN, there are multiple points of the operator's HFC network entering to the client's house. In the Full IP HN, there is just one, and that is a demarcation point, it separates the operator's network. This demarcation isolates the ingress noise that could be produced inside the HN in the traditional COAX connections inside the client's house. That could generate a very good impact on the network operator in terms of maintenance and improvement in the QoS of the network.
- The Home Network model it is the same for DOCSIS or for xPON access networks, obviously the CPE will be deferent, EROUTERS vs ONT with embedded RWGs, however the RGW itself of both must have the same capabilities and functions, and the IP STB is agnostics to any of those access networks, even more, the TV service installation could be almost the same process, it is the same for booth access network.
- Even with the installation complications, that were aforementioned, the WiFi capabilities of the EROUTER/ONTs and IP STBs should be enough to cover most the homes, and with the help of an installation guide implemented in the STB, that could be shown during the first installation of the STB, it could mean the first steps towards an auto-installation, which means more cost savings.

From the LAN point of view, the capabilities must be the same for any CPE D3.0, D3.1 or xPON. For instance, the Figure 44 and Figure 45 are the minimum hardware reference architecture for the CM Residential Gateway for D3.0 and D3.1 respectively. The CPE is a device compliant with eDOCSIS specification (12), it includes an embedded D3.0 or D3.1 CM (eCM). Regarding the WiFi capabilities, the minimum requirements is a Dual Band Concurrent WiFi Radios in 2.4Ghz and other in 5Ghz, with the support of 802.11n (13) and 802.11ac (14) (wave 1/2). The antennas' arrangements are 3x3:3 for 2.4Ghz and 4x4:4 in 5Ghz. The CPE supports IPv4, IPv6 and Dual-Stack IPv4 and IPv6, but the video service up now only supports IPv4, thus the rest of the requirement explained here are based on IPv4.

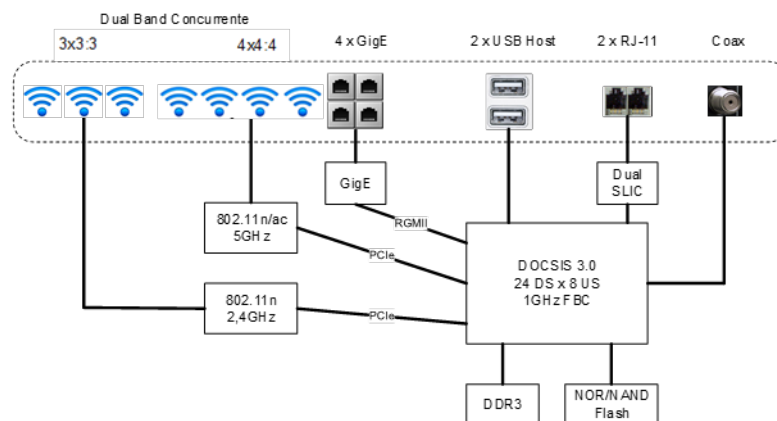


Figure 43- Cable Modem Residential Gateway WiFi D3.0 IPTV ready.

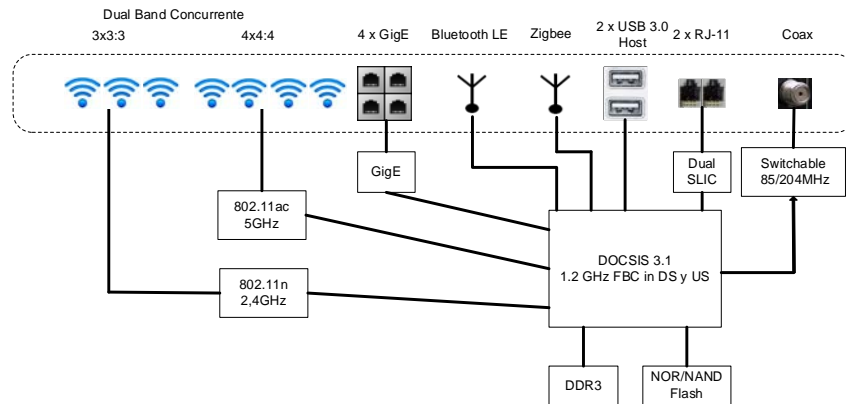


Figure 44 – Cable Modem Residential Gateway WiFi D3.1 IPTV ready.

Regarding the multicast frames forwarding the CPE must implement some specific functions (Figure 45):

- It must support IGMPv3 (10) in the Ethernet or WiFi interfaces.
- The internal LAN switch in the EROUTER must support IGMP Snooping (15) capabilities to forward the multicast frames only at the ports that have a connected device (receivers) that had requested a subscribe to a Multicast Group, and avoid sending multicast traffic to devices that are connected to the same LAN and have not request this data. In general, the tracking per-host membership status on an interface allows also to implement fast leaves, with this feature the switch can determinate when the last device (IP STB) sends an IGMP leave and it stops immediately to forward multicast traffic for this group.
- The LAN devices are not connected directly to the CMTS, who should receive the IGMP message transparently through the CM, in this case, the eCM. The IP STBs, same that other devices are connected to the access network using PAT, then the IGMP messages generated by STBs are not connected directly to the CM but to the internal router of the EROUTER. To manage that situation, the internal router must work as IGMP Proxy on behalf of IGMP messages generated form LAN devices. The internal router must keep the tracking of which device in the LAN requests IGMP messages and it resends those to the access network changing the host who generates the request by its WAN IP.
- When in same LAN there are devices working with IGMPv2 and IGMPv3, the router's local interface will receive IGMPv2 and IGMPv3 Membership reports, then when that happens, it will decide to do a fall back to IGMPv2, even when the router was configured to work with IGMPv3. If it is not possible to prevent IGMPv2 requests from devices other than an IP STB, then it could generate an issue when it is expected that the system works based on IGMPv3. To avoid that, the CPE must block the IGMPv2 message toward the internal router.
- Multicast over WiFi does not provide a good mechanism to guarantee the QoS of the all of IP STBs connected to the wireless LAN. When two or more devices are connected in a WiFi network they could place at different places in the HN, and then they could reach different bit rates based on the WiFi profiles that those devices can negotiate with the Access Point (AP). In such a case, the multicast packets, that must be just one stream for all the devices that have subscribed the same multicast group, and it will be forwarded with the lowest bit rate that can be reached by the device that has the worst situation in terms of WiFi. To avoid that the CPE must implement a Multicast to Unicast forwarding at the WiFi level. When multicast frames must be forwarded from the WAN interface to the WiFi segment, the MAC address Multicast Group it is replaced by the Unicast MAC address of the receiver that had subscribed to that Multicast group. Figure 45 shows how the CPE replace the multicast MAC address with host X's MAC address

when it forwards the multicast frame over the WiFi (which does not happen in the Ethernet segment).

- To avoid deny of service attacks the CPE must allow configuring a limit in the amount of the multicast session that could potentially be subscribed by devices behind the CPE. So, for instance, if the HN has 5 IP STB in each moment connected, it is not necessary to support more than 5 multicast groups. That could be configured with SNMP per device, and/or could be part of the CPE file configuration that is download during the registration, or it could be configured based on TR-069 (16) from the Automatic Configuration Server (ACS).
- An extra security mechanism that should be implemented is an IGMP messages threshold, after a given number of IGMP message per second the CPE should discard the IGMP messages, to avoid a waterfall of IGMP messages towards the access network.
- Another not mandatory but useful mechanism, it is to allow the operator to restrict which devices can subscribe to IPv4 specifics multicast groups used for the TV service. This is required to allow only STBs can subscribe to those multicast groups and it does not allow user devices to subscribe them to avoid attacks.

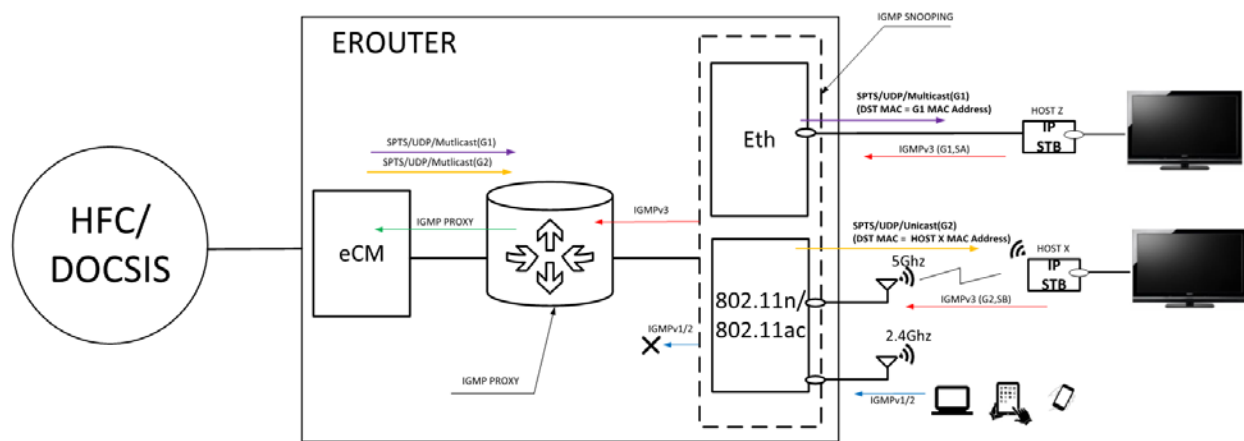


Figure 45 –Multicast Functions that CPE must implement.

The IP STB can be connected with Ethernet or WiFi, being Ethernet the first choice. But sometimes it is not possible to cabling the house because the technicians could find itself with objections from the client, installation time, etc.

The buildings are mainly located in urban areas with high population density and where the concurrence of Wi-Fi networks is consequently very high. If a network scan is made in any department within Buenos Aires city, it is difficult to find less than 20 networks within the 2.4GHz band. Added to this is the fact that these conditions are not constant but tend to vary over time.

The band of 5GHz first, it is a much less congested frequency, the number of networks that emit on this frequency is quite small, so the interference is much lower. Also, the coverage of this frequency is lower in distance than 2.4GHz, thus the potential interference problems are reduced.

In situations where cabling is not possible, and the connection is made through Wi-Fi, the STB IP must be connected only to the 5GHz network considering the signals levels. The levels reported by Received Signal Strength Indicator (RSSI) of the IP STB must be better than -60dBm. That level is verified in the IP STB's connection process and if the level is not reached then the STB will show on the screen a message indicating that WiFi conditions are not enough for the video service.

To guarantee the QoS in the WiFi it is used Wireless Multimedia (WMM) what is based in 802.11e (16). With this, the WiFi is prioritized according to dissimilar categories of data which is determined based on the “P” bit of 802.1p. Table 5 shows the mapping between the P bit and the data types. The P bit in the multicast packet is getting from the DSCP mark as can be appreciated in the Figure 40. The Figure 46 shows a capture of the WiFi traffic where the DSCP and then “P” bit is mapped correctly to the WMM category.

Table 5- WMM “P” bit mapping.

802.1p Priority	WMM Access Category
1	Background
2	
0	Best effort
3	
4	Video
5	
6	Voice
7	

43349	2018-04-20	15:49:22,220064	10.254.245.200	239.200.10.10	UDP	1434	58105 → 4000	Len=1328
43542	2018-04-20	15:49:22,298823	10.254.245.200	239.200.10.10	UDP	1434	58105 → 4000	Len=1328
43732	2018-04-20	15:49:22,378462	10.254.245.200	239.200.10.10	UDP	1434	58105 → 4000	Len=1328
43951	2018-04-20	15:49:22,456201	10.254.245.200	239.200.10.10	UDP	1434	58105 → 4000	Len=1328
44139	2018-04-20	15:49:22,539592	10.254.245.200	239.200.10.10	UDP	1434	58105 → 4000	Len=1328
44384	2018-04-20	15:49:22,620971	10.254.245.200	239.200.10.10	UDP	1434	58105 → 4000	Len=1328

✓ IEEE 802.11 QoS Data, Flags:F.C
Type/Subtype: QoS Data (0x0028)
Frame Control Field: 0x8802
.000 0000 0011 0000 = Duration: 48 microseconds
Receiver address: IntelCor_22:4a:a8 (24:77:03:22:4a:a8)
Destination address: IntelCor_22:4a:a8 (24:77:03:22:4a:a8)
Transmitter address: Sagemcom_5a:8d:67 (58:90:43:5a:8d:67)
Source address: Cadant_81:64:46 (00:01:5c:81:64:46)
BSS Id: Sagemcom_5a:8d:67 (58:90:43:5a:8d:67)
STA address: IntelCor_22:4a:a8 (24:77:03:22:4a:a8)
.... 0000 = Fragment number: 0
1011 0100 1111 = Sequence number: 2895
Frame check sequence: 0xddbdf9fd [correct]
[FCS Status: Good]
QoS Control: 0x0005
.... 0101 = TID: 5
[.... 101 = Priority: Video (Video) (5)]
.... 0000 = EOSP: Service period
.... 0000 = Ack Policy: Normal Ack (0x0)

Figure 46 – WiFi Traffic capture with WMM.

Conclusion

There are several drivers to move towards Full IP video services, it is possible to get this objective in incremental steps. Move from legacy DTV system directly to Full IP system is very disruptive not only for the clients but also for the operators. The evolution of the TV services should be done in an agile way

allowing the coexistence of legacy technologies with the new ones, and at the same time protecting the investment, giving the agility to speed up the deployment.

From the cable operators point of view the first step is to move to a Hybrid ecosystem, that allows providing a new UX to the end clients but keeping the robustness of the already deployed DTV services, and simultaneously generating new advance ways to consume TV, like better UI, TSTV functionalities, recommendations, and introducing new unmanaged devices like second screens.

There are two layers to evolve the ecosystem, the control plane, and the data plane; the first one is the main to introduce the changes in the UX and here is the first place where IP is needed, the second one is for the video distribution. For the Hybrid TV services, a new BO, new HE components, and a CDN are deployed. The control plane and media distribution for CoD and TSTV services as well as the second screen devices are based on IP. But most of the Live TV services are kept in DTV, which allows reusing the TV signals that are already in the HE, the transport of those signals through the networks does not change, in the access networks there are no changes, and even more important, the installation process of the TV services at the Home Network remains mostly like in the legacy DTV.

As it was mentioned at the beginning of this document, there are more necessities which must be covered, for instance: to deploy video services not only in HFC networks but also in other IP based access networks like xDSL, FTTH, Mobiles and in general other unmanaged networks; there are also better cost-effective STBs in the IP world and there are more vendors diversity. Another good point is that with IP STBs it is possible to have a more agnostic HN that could be connected in different access networks. Those drivers and the improvements in the technologies, like new video encodings, are facilitators in order to finally move from the Hybrid ecosystem to a Full IP Video.

It also was analyzed the linear TV audience behavior and the relation with the media transport mechanism. In small service areas with no more than 64 HHPP using unicast is it doable, but for a bigger size of HHPP to use only unicast for Linear TV could generate issues with the network capacity. Here is where Multicast helps, mainly for the most popular channels meanwhile unicast only for less popular ones (Figure 23) and CoD.

Finally, it was explained which would be the transformation to move from the Hybrid TV to the Full IP Video, starting in the HE, and then it goes through the transport and access networks where the main difference is the implementations of the multicast, and finally the HN where the transformation is visible not only for operator but also for the end client. It has been realized that Multicast over DOCSIS has some issues due to the ingress noise in the HFC plant or in the HN, and some improvements on DOCSIS side or in the video control plane are needed to provide a better resilient service.

For the future there are more improvements that it should be implemented, for instance, some of them are:

- It was mentioned that less popular Linear TV channels and the complete lineup for unmanaged devices, could be transported with ABR in Unicast delivery, however, there could be some constraints. There is a delay between the DTV and ABR signals which is around 20 to 40 secs (depending on the ABR's segments size, encodings, player buffer, etc.). This delay is not acceptable for the client point of view, especially for the specific type of content like live events or news. There are new technologies of Low Latency ABR that could improve that, for instance, Common Media Application Format (CMAF) (6).
- Another issue with ABR technologies are the subtitles for linear content, most of the content provider (at least in LATAM) provides Digital Video Broadcasting Subtitles (DVB-SUB) (17) which are Bitmap images that are not directly supported in ABR technologies, they support text

formats like Web Video Text Tracks (WebVTT) and others similar. There are mechanisms that use real-time Optical Character Recognition (OCR) that allows converting the subtitles images into text and then to generate the WebVTT. Also, the last versions of HLS and DASH allow also to use SMPTE-TT (18) embedding the images subtitles in base64 encoding, but for that new mechanisms must be implemented in OS at HE side and video Players in devices.

- Multicast is well supported in managed devices, but unmanaged devices do not support in general the Multicast, they use just Unicast ABR. That is not a big deal for small second screens at the moment, but it could be, again, an issue for the network sizing/planning when the operator starts to deploy the video services over smart TVs or Console Games in a massive manner because the audience behavior is like STBs (big screens). To address that there are new mechanisms that allow the delivery of ABR over Multicast, for instance, there are CableLabs specifications (19) and a DVB draft (20).

As it could see in this document, the challenge to reach the objective is huge, not only for the changes in the technology but also for the changes in the skills of the operator which must deal with a new process in order to deliver the video in this new way, however this a new way that CSP must start to transit to address the clients and networks necessities, and finally to get the benefits of a real convergence.

Abbreviations

ABR	Adaptive Bit Rate
ACS	Automatic Configuration Server
AP	Access Point
ATS	Adaptive Transport Stream
BBIP	Backbone IP
BO	Video Back Office
CA	Conditional Access
CAC	Connection Admission Control
CBR	Constant Bit Rate
CDN	Content Delivery Network
cDVR	Cloud Based DVR
CE	Consumer Electronic
CM	Cablemodem
CMAF	Common Media Application Format
CMS	Content Management System
COAX	Coaxial Cable
CoD	Content on Demand
COTS	Commercial Off the Shelf
CPE	Customer Premises Equipment
CSP	Communication Service Provider
CuTV	Catchup TV
D3.0	DOCSIS 3.0

D3.1	DOCSIS 3.1
dB	Decibel
DBC-REQ	Downstream Bonding Change Request
DBG	Downstream Bonding Group
dBm	Decibel milliwatt
DCII	Motorola Digicipher II
DCS	Downstream Channel Set
DOCSIS	Data-Over-Cable Service Interface Specifications
DOCSIS	Data Over Cable Service Interface Specification
DOM	Document Object Model
DRM	Digital Rights Management
DSG	DOCSIS STB GATEWAY
DSID	Downstream Services ID
DSL	Digital Subscriber Line
DVB	Digital Video Broadcasting
DVB-SUB	DVB Subtitle
DVR	Digital Video Recorder
eCM	embedded CM
EPG	Electronic Program Guide
EQAM	Edge QAM
eSTB	Embedded Set Top Box
FHD	Full High Definition
FPS	Frames per second
GigE	Gigabit Ethernet
HD	High Definition
HE	Head End
HEVC	High Efficiency Video Coding
HHPP	Homes passed
HTTP	Hyper Text Transfer Protocol
HTTPS	Secure Hyper Text Transfer Protocol
IANA	Internet Assigned Number Authority
IC	Intermediate Cache
IEEE	Institute of Electrical and Electronics Engineers
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IPTV	Internet Protocol Television
ISMV	IIS Smooth Streaming Media Video
LAN	Local Area Network
LTE	Long Term Evolution
Mbps	Megabits per second

MDF	Multicast DSID Forwarding
MHP	Multimedia Home Platform
MP4	Format created by the Moving Picture Experts Group (MPEG) as a multimedia container
MPEG	Moving Picture Experts Group
MSO	Multiservice Operator
NAPT	Network Address Port Translation
NAS	Network-attached storage
NAT	Network Address Translation
nDVR	Network Based DVR
nPLTV	Network Based Pause Live TV
OCAP	Openable Application Platform
OCR	Optical character recognition
OFDM	Orthogonal Frequency Division Multiplexing
OLT	Optical Line Termination
OOB	Out of Band
OTT	Over the Top
PAT	Port Address Translation
PC	Personal Computer
PHY	Physical
PIFF	Protected Interoperable File Format
PLC	Power Line Communications
QAM	Quadrature amplitude modulation
REPG	Reverse EPG
RGW	Residential Gateway
SCN	Service Class Names
SC-QAM	Single Carrier-QAM
SCR	Video Scrambler which integrated with a CA encrypts the video
SD	Standard Definition
SDH	Synchronous Digital Hierarchy
SDI	Serial Digital Interface
SF	Service Flows
SG	Service Group
SPTS	Single Program Transport Stream
STB	Set Top Box
SVOD	Subscription VoD
TS	Transport Stream
TV	Television
TVOD	Transactional VoD
UHD	Ultra-High Definition
UI	User Interface

URL	Uniform Resource Locator
USB	Universal Serial Bus
UTP	Unshielded Twisted Pair
UX	User experience
VBR	Variable Bit Rate
VoD	Video on Demand
WebVTT	Web Video Text Tracks Format
WiFi	Wireless Fidelity

Bibliography & References

1. **SCTE, Steven R. Harris Senior Director Advanced Network Technologies Program Development.** *Home Networks – It's Not That Simple Anymore.* s.l. : SCTE Digital Home Symposium, 2012.
2. **Christopher Mueller.** bitmovin. [Online] March 29, 2015. <https://bitmovin.com/mpeg-dash-vs-apple-hls-vs-microsoft-smooth-streaming-vs-adobe-hds/>.
3. **CableLabs.** Adaptive Transport Stream Specification. *www.cablelabs.com.* [Online] 02 14, 2014. [Cited: 05 1, 2018.] <https://apps.cablelabs.com/specification/5404>.
4. **IETF.** HTTP Live Streaming (RFC 8216). *RFC 8216.* [Online] 08 2017. [Cited: 05 01, 2018.] <https://tools.ietf.org/html/rfc8216>.
5. **MPEG-DASH.** Guidelines for Implementation :DASH-IF Interoperability Points. *dashif.org.* [Online] 04 2015. [Cited: 05 01, 2018.] <https://dashif.org/w/2015/04/DASH-IF-IOP-v3.0.pdf>.
6. **Weil, Nicolas and Bouqueau, Romain.** Ultra Low Latency with CMAF. [Online] 07 2017. [Cited: 05 01, 2018.] https://parisvideotech.com/wp-content/uploads/2017/07/Bouqueau-Weil_UltraLowLatencyWithCMAF.pdf.
7. *Sustained Throughput Requirements for Future.* **Jeroen Wellen, Prudence Kapauan, Amit Mukhopadhyay.** s.l. : SCTE-ISBE, 2017.
8. **Reuss, Ron.** *IP Unicast v. Multicast Modeling Overview.* s.l. : CableLabs, 2012.
9. **IETF.** Host Extensions for IP Multicasting(RFC 1112). *tools.ietf.org.* [Online] August 1989. <https://tools.ietf.org/html/rfc1112>.
10. —. Internet Group Management Protocol, Version 3 (RFC3376). *tools.ietf.org.* [Online] October 2002. <https://tools.ietf.org/html/rfc3376>.
11. **CableLabs.** DOCSIS 3.0 MAC and Upper Layer Protocols Interface Specification. *www.cablelabs.com.* [Online] 12 07, 2017. <https://apps.cablelabs.com/specification/CM-SP-MULPIv3.0>.
12. —. CM-SP-eDOCSIS. *www.cablelabs.com.* [Online] 09 06, 2017. <https://apps.cablelabs.com/specification/CM-SP-eDOCSIS>.

13. **802.11n** - *Local and metropolitan area networks—Specific requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput*. **IEEE**. 2009.
14. **802.11ac** - *Telecommunications and information exchange between systems—Local and metropolitan area networks-- Specific requirements--Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications--Amendment 4: Enhancements for V*. **IEEE**. 2013.
15. **IEEE**. Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches (RFC 4541). *tools.ietf.org*. [Online] 05 2006. <https://tools.ietf.org/html/rfc4541>.
16. **Forum, BroadBand**. TR-069 CPE WAN Management Protocol. *www.broadband-forum.org*. [Online] 03 2018. <https://www.broadband-forum.org/technical/download/TR-069.pdf>.
17. **ETSI**. Digital Video Broadcasting (DVB) Subtitling systems - ETSI EN 300 743 V1.3.1 (2006-11). *www.etsi.org*. [Online] 11 2016. https://www.etsi.org/deliver/etsi_en/300700_300799/300743/01.03.01_60/en_300743v010301p.pdf.
18. **SMPTE**. Timed Text Format - SMPTE-TT - SMPTE ST 2052-1:2010. *www.smpste.org*. [Online] <https://www.smpste.org/sites/default/files/st2052-1-2010.pdf>.
19. **CableLabs**. IP Multicast Adaptive Bit Rate Architecture Technical Report. *www.cablelabs.com*. [Online] 10 26, 2016. <https://apps.cablelabs.com/specification/ip-multicast-adaptive-bit-rate-architecture-technical-report/>.
20. **DVB.org**. Digital Draft - Video Broadcasting (DVB); Adaptive media streaming over IP multicast - DVB Document A176. *https://www.dvb.org*. [Online] 3 2018. https://www.dvb.org/resources/public/standards/a176_adaptive_media_streaming_over_ip_multicast_2018-02-16_draft_bluebook.pdf.
21. *TRANSPORT, CONTENT, AND SERVICE IMPLICATIONS ON VOD NETWORK*. **George Kajos, Conrad Clemson**. 2005.
22. **Juniper**. Understanding PIM Source-Specific Mod. *www.juniper.net/*. [Online] https://www.juniper.net/documentation/en_US/junos/topics/concept/multicast-pim-ssm.html.
23. **Cisco**. IP Multicast Technology Overview. *www.cisco.com*. [Online] April 4, 2002. https://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/ip_multicast/White_papers/mcst_ovr.pdf.
24. **IEEE**. ANSI/IEEE 802.11e-2005 - IEEE Standard for Information technology--Local and metropolitan area networks--Specific requirements--Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 8: Medium Access Control . *http://standards.ieee.org*. [Online] 2006. <http://standards.ieee.org/findstds/standard/802.11e-2005.html>.
25. **John Horrobin, Gitesh Shah**. *Pioneering IPTV in Cable Networks*. SCTE Expo, Atlanta : Cisco, 2013.
26. **Hanks, William T**. *Configuration Recommendations for DOCSIS Transport of Managed IP Video Service* . SCTE Expo, Philadelphia : Arris, 2016.

Designing Video Services for Low-Latency Distributions in IPTV Cable Systems

A Technical Paper prepared for SCTE•ISBE by

Yasser Syed

Comcast Distinguished Engineer
TPX/VIDEO/VAST Dept., Comcast Cable
1701 JFK Boulevard Philadelphia, PA 19103 USA
303-246-8413
yasser_syed@comcast.com

Ali C. Begen, Professor, Ozyegin University, Turkey

Alex Giladi, Comcast Distinguished Engineer, TPX/VIDEO/VAST Dept., Comcast Cable

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
1. Types of Latency	3
2. The Player/Decoder Model	4
3. Types of Delivery.....	6
3.1. Pre-Packaged Delivery	6
3.2. Live Event Delivery	7
4. Handling Low-Latency Live Events	7
4.1. Segment Creation	7
4.2. Segment-Based Delivery	8
Conclusion.....	10
Abbreviations	10
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1: Types of latency in the end-to-end system.....	4
Figure 2: Player/decoder model.	5
Figure 3 - Tradeoffs between low latency, scalability, and QoS (Quality of Service).	5
Figure 4: CMAF segment, fragment and chunk structures.	8
Figure 5: Server distribution using HTTP over TCP and QUIC.	9

Introduction

Service providers are moving towards IPTV (Internet Protocol Television) technologies using fragmented delivery as a way distribute video services over bandwidth-varying environments. This has already been successful with on-demand or pre-packaged video services (television shows, movies) and has been adapted, to some degree, for live event-based services (sports, news) -- but fragmented delivery works best for content that is already pre-packaged into a final format before delivery. Distributing live events, by contrast, presents additional latency-related challenges due to the real-time handling of requests, storage and delivery. For instance, customers expect to be able to record live broadcasts for later viewing. However, creating low-latency, live streaming by altering existing technology tools and structures simultaneously creates new and often prohibitive complexities. This paper examines some of the existing technologies in use today to optimize for network latency, and introduces possible implementations for achieving low-latency streaming.

1. Types of Latency

Latency can refer to one or more of latency causes in multimedia streaming that really depend on the application context. In live streams, latency is mostly discussed as an end-to-end latency, which describes the time delay between the action occurring in front of the camera and the same action being observed on the display. This is also known as glass-to-glass or handwaving latency. This cumulative latency can be broken up into three sub-parts that comprise of latency due to content preparation, latency for distributing it, and latency on the player.

For live events, content preparation latency is mostly associated with how the content is captured and encoded (including conditioning) through real-time processing using codecs such as AVC (Advanced Video Coding)/H.264 or HEVC (High Efficiency Video Coding)/H.265. Latency is inherently linked to the encoding process (e.g. look ahead) and the encoding structure through the use of temporal compression techniques that alter picture reordering in the coded stream, in order to increase picture quality while allowing for random access of the stream. Another part of content preparation latency is the packaging latency, which happens as the segment is received at all (or most) representations and parsed for indexing, so as to make it fetchable (manifests/indexing byte offsets) or streamable (multicast/uncast).

Once the content is prepared, it needs to be distributed. In file-based streams, distribution latency becomes one of the major factors discussed in this context. It is comprised of segment distribution (or the lack of it) to the CDN (Content Delivery Network) latency (CDNL), which does include delay for the I/O (input/output) storage of the segment. Another major factor is server latency, which describes the delay from the segment's availability on the server and the actual request for the content, to the receipt of it by the client player. This is often described as the delay from "the live edge". This type of latency is important for any type of file-based streaming, including pre-packaged content, because it prevents lag in player controls (e.g., seek, skip, rewind, fast forward). For live events, reducing this latency reduces the overall "glass-to-glass" latency.

The last sub-part to end-to-end latency is the player latency, which is a combination of the decode latency, buffer management and stream delay. Decode latency occurs relative to the coding structures used by the codec. Buffer management uses latency to avoid buffer underflow by extending the buffer to accommodate for "chunky" loading, due to the segment length, which may slow down the fill rate into the buffer. Lastly, stream delay is partly caused by player startup delay, which provides time to initially grow a buffer's capacity before playout occurs. This delay can avoid buffer underflow situations attributable to network

congestion that can slow down the buffer's fill rate. The startup delay is different from latency, which could take 5 seconds from pressing play, but the stream is only behind 3 seconds from the “live edge” (and the transmission time of the segment from the server.) The other part of stream delay is related to seeking behavior on the stream, and the time to initiate the action and see the result, which is a measure of delay when the end user time-shifts in the buffer.

Yet another type of latency happens at the system control plane level, where the number of requests to the server necessarily increases as a function of the number of players it handles. Above a certain capacity, this slows down the HTTP response time to send out segments to specific players, ultimately the filling rate into the client buffer. An increase in the number of players that a server handles would start making this type of delay more noticeable.

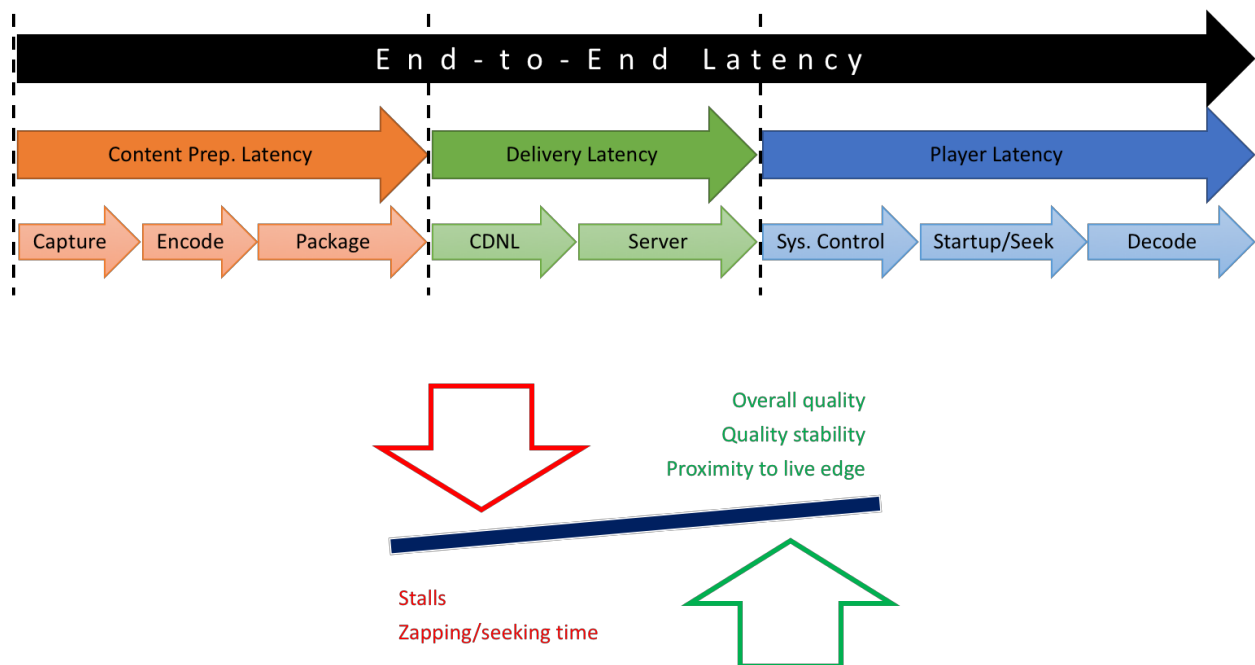


Figure 1: Types of latency in the end-to-end system.

A certain amount of latency is useful in network media streaming applications to ensure proper playback of the stream. It is a hedge against network jitter, keeps the player buffer filled, and allows for increased quality in low bitrate coding. Providing for low-latency is always a tradeoff against robustness of playback of the stream captured, and yet having low-latency streaming is important for live streaming. Additionally, developing low-latency techniques will also be helpful in emerging and latency-averse applications including, but not limited to, augmented and virtual reality (among many others.)

2. The Player/Decoder Model

In both live event and pre-packaged delivery, the player/decoder model is common. It is important to understand the mechanics of the player/decoder model and how it behaves, with respect to latency. Figure 2 illustrates a typical player/decoder model. The client makes requests for content segments from the HTTP server, which responds with the specified content segment for each request. This process starts filling the client buffer at a fast pace until the designated startup delay is achieved. Then, the client buffer starts to offload segments into the Coded Picture Buffer (CPB) at a rate to maintain the HRD (Hypothetical

Reference Decoder) model (an equivalent is done on the audio decoder side for audio segments), while simultaneously requesting new segments from the server. The goal is to maintain a constant client buffer level, which is much larger than the buffer maintained for the HRD model. When it encounters network congestion, the client player adjusts its requests according to a player algorithm to attempt to maintain a sufficient client buffer level. If the client buffer drains, then a rebuffering instance (observable as a stall in stream playout) happens. Other approaches may also attempt to adjust the emptying rate of the client buffer.




Figure 2: Player/decoder model.

The client buffer feeds the elementary stream's coded picture buffer (CPB) for both audio and video with the respective media segments. Depending on the algorithm, the loading of the CPB may be “chunky,” and if not smooth enough, it may affect the HRD model's playout of the stream, which may also cause a rebuffering state. For video, the HRD model for a particular GoP (Group of Pictures) structure is maintained, from the ordering of the coded pictures, which improves video quality by retaining these pictures as a reference picture in the decoded picture buffer. A decoding latency occurs throughout this decoding process to typically accommodate a 2-4+ picture delay. Typically, gaining higher quality at the intended bitrate employs techniques at the cost of increasing latency. In fact, increasing quality of the stream or increasing scalability of the number of players happens with a tradeoff in latency (see Figure 3).




Figure 3 - Tradeoffs between low latency, scalability, and QoS (Quality of Service).

Additionally, shortening the segment size can reduce latency by filling in the CPB buffer more “smoothly” and adjusting the segment requests coming from the client. For example, if DASH (Dynamic Adaptive Streaming over HTTP) is used with 2-second segments, this results in latency of about 8-10 seconds (cumulative), while HLS (HTTP Live Streaming) using 10-second segments can result in latency as high as 30 seconds (assuming at least three segments are loaded into the buffer). With 6-second segments, which were allowable since mid-2016 for HLS, the latency for HLS streams reduced to 18-20 seconds. Thus, smaller segment durations can assist in reducing latency -- until the segments become too small and increase network traffic as a result of the increased requests to the server.

3. Types of Delivery

In the current fragment delivery models, the services all use different types of segmented, file-based delivery processes. Initially, on-demand applications were popular for ABR (Adaptive Bitrate)-enabled video services, because the creation of the content segment could be done well before the delivery. Furthermore, it could be prepositioned on the servers in the CDN, which shortened the delivery pathway from the server to the players.

Live events are another type of segmented file-based delivery process that differ from pre-packaged delivery because content segment creation is done serially with the delivery of those segments. In the live case, the viewing experience is sensitive to the latency from creation to delivery, while in the pre-packaged delivery case, it is more crucial to make sure the viewing quality is good relative to the latency considerations.

3.1. Pre-Packaged Delivery

For pre-packaged delivery, the encoded content already exists as a single file or set of segments. The goal for file-based delivery is for the client/player buffer to be filled such that it always allows the next frame to be decoded and played out. Delay in delivery of the next segment would cause the buffer to underflow, thus pausing the playout of the content -- to the obvious detriment of the QoE (Quality of Experience). Factors that can cause disruption in delivery include insufficient bandwidth between the server and client; congested bandwidth that causes delay of segments; unequal segment arrival times due to jitter; or simply missed segments traveling over a best-effort system that never arrive due to reaching the TTL (Time-to-Live) limit. Possible solutions to these issues include the following strategies:

- Create a longer startup time to allow for the CPB to grow larger before the playout process starts draining the buffer. This allows for more time to allow the next segment to reach the buffer (or to request an older, missing segment).
- Constrain the size of the segments being sent to avoid the risk of rapidly depleting the buffer. This can be done on a segment-by-segment basis, using adaptive streaming technologies.
- Monitor and manage bandwidth to the client by reducing the size of the segments. This can also be done by adaptive streaming technologies, where the player can monitor the buffer and request alternative and smaller-sized addressable content segments, of lower quality, in order to deal with best effort connections that at times experience bandwidth congestion. Other ways to avoid bandwidth congestion on the connection involve managing the capacity of the connection to mitigate bandwidth fluctuations.
- Create segments that are self-decodable, so as to avoid a loss because of segment corruption and the consequent need to retransmit it.

Pre-position segments at multiple servers situated physically closer to the client players, which can be done through a CDN network that can reduce the path latency of the packets. Having multiple servers can also reduce delays by reducing the number of segment requests at a single server during flash demand periods.

Each of these strategies can help to ensure smooth robust playback of the streams, but many of these plans create the need for a lengthened buffer that is more smoothly filled relative to both the client buffer and CPB. This adds to delay of the initial start of stream playback, and increases latency on the stream on the order of seconds or more.

3.2. Live Event Delivery

For live event delivery, the segments are created in the current moment by the camera capture and encoding processes. The segments, once created, are then transmitted over the network and delivered to the client player. It is meaningful, in live services, to reduce the glass-to-glass latency to reduce the amount of latency that may be more noticeable if alternative viewing distribution formats are available (e.g., MPEG2-TS (Moving Pictures Experts Group, Transport Stream) QAM (Quadrature Amplitude Modulation) delivery or attending the live event). There are several methods to reduce latency, which can even be cumulative, but each approach has tradeoffs in either latency, quality or robustness of playout. Methods to create lower latencies can be grouped in the following two categories:

- Modifying the segment creation process and providing better integration of this to the delivery process,
- Modifying the segment delivery process to bring this closer to the live edge. A lot of these techniques can be applied to file-based streaming as well.

In both of these categories, there are requirements wherein live streaming should still enable a recordable service (e.g., cDVR [Cloud Digital Video Recorder]) that can be used for later viewing of the content. Also, it is important to know that not all pre-packaged streaming features make sense for real-time streaming events.

4. Handling Low-Latency Live Events

The current adaptive streaming technologies have been able to handle live events, but at the cost of adding significant end-to-end delay. For next-generation technologies, a focus on reducing transmission latency is paramount. Such a latency reduction is well-served to focus on two areas: i) the content creation process for producing segments, and ii) the delivery of these segments to one or more players. In turn, some of these techniques, described below, can aid in pre-packaged segment delivery of content, too.

4.1. Segment Creation

To reduce latency in the area of segment creation, temporal compression techniques in the encoding process should be reduced. This will result in an alignment of the coded elementary stream with the presentation order of the frames, which can be done through the use of known spatial compression techniques, or very simplified FPP (Forward Predicted Picture) encoding processes. Both result in a loss of quality, so the existing and traditional techniques to improve quality (e.g., look ahead or pre-processing methods) need to be used sparingly to avoid adding latency to the stream. The resulting reduction in compression efficiency either results in a loss of quality or is compensated for by an increase in bitrate, which would increase the average connection bandwidth.

In terms of delivery of the in-process created segment, the packager does not have to wait for the entire segment to be processed, because it can release chunks (which do not need to be all self-decodable and could resolve to even a single frame) of the segment, which are then subdivided into decodable fragments [see Figure 4] as soon as they processed. This is a tool employed by the CMAF (Common Media Application Format) specification for ISO-BMFF media segments. This way, the application client can release the received chunks earlier to smoothly fill the player buffer as the chunks are being received. These chunks can also be playable as they are received, so long as the prior chunks for that segment have arrived. This approach can aid in filling the client buffer quickly, without waiting for the entire segment to be available for delivery.

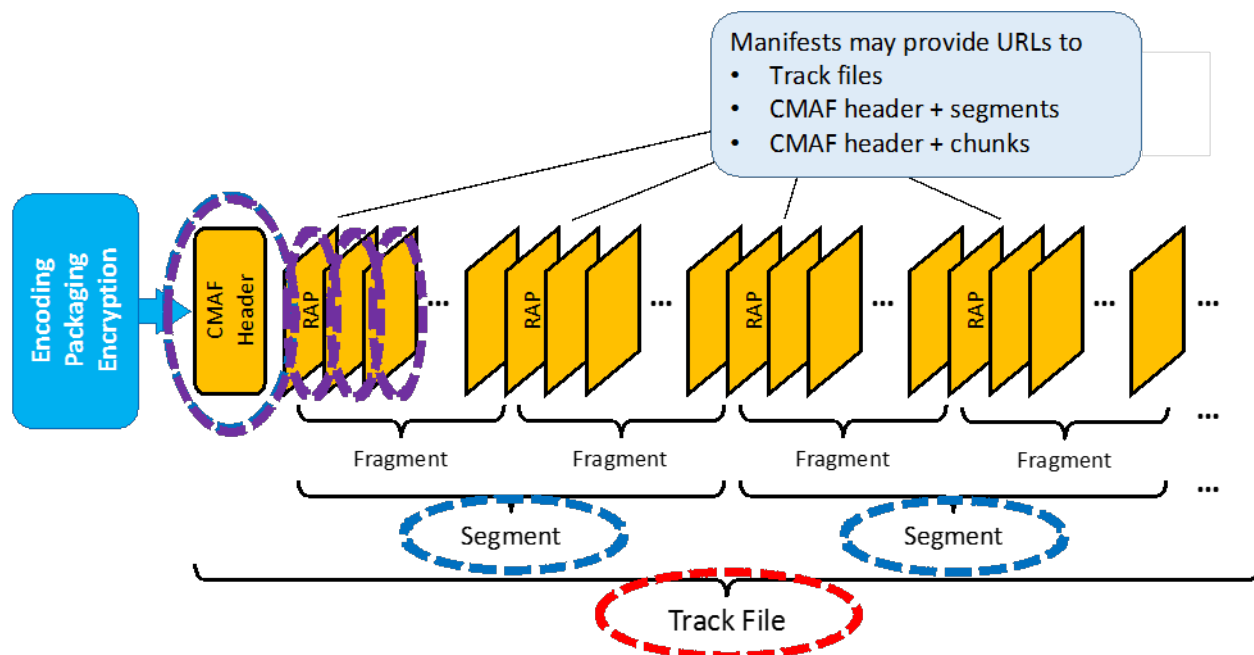


Figure 4: CMAF segment, fragment and chunk structures.

4.2. Segment-Based Delivery

Other opportunities to reduce latency are more applicable to how the segments are being delivered. In pre-packaged file delivery, prepositioning the segments beforehand, in the CDN system, may reduce delivery latency once the segments are requested. For live streaming this might not be as efficient. CDNs need to be cognizant of when they are handling a low-latency stream. Pre-positioning segments into more edge servers can cause start-up delays of a newly requested live and low-latency stream and can cause multiple players to not be synchronized in their playout. This can be attributable to delays with manifest creation/segment availability, or even just the delay of copying and distributing segments on multiple cascaded servers in the delivery path. Moving more towards a centralized server distribution or minimal hierarchy layered distribution can reduce this type of delay but would extend path latency issues.

An alternative to using a full CDN distribution mechanism to reduce live stream latency is to stream from a central server, but in this case, scalability would become an issue. If a low-latency stream is highly requested, then the central server/server(s) would need to be scaled to respond to a high traffic volume of request/responses from thousands or even hundreds of thousands of clients. An alternative approach would be to use a type of multicast/broadcast transmission to a player or an edge server. Using an approach that multicasts segment transmission, such as ATSC (Advanced Television Systems Committee) 3.0, could handle the traffic volume, but potentially at a latency cost, where the latency sacrifices come from delays in manifest creation and segment availability. Another alternative approach could be a hybrid of this, where content is multicast to local servers and then distributed from that point. The multicast component does not have to be a multicast of segments but could instead be a multicast of a marked up MPEG2-TS stream using AF, or Adaptation Field Descriptors, as described in the MPEG-2 Systems document and using the manifest described in SCTE 214-4, which is sent to a local server or even at a home gateway and packaged at that point. These types of streams can be automatically outputted from the encoder and traditionally multicast to a local server without the need for a manifest. A linear packager at the local server can then create the segments and the manifest locally. As one can see, there are several different approaches that could reduce the distribution latency to the servers.

Another area to look for possible latency reduction is the internet transport protocols for handling information exchange. At the application layer, HTTP/1.1 is commonly used for server- to-client exchanges of media. The HTTP/1.1 protocol may use a free and open-source web server implementation. Using the HTTP protocol allows for adaptive streaming across generalized IP (Internet Protocol) hardware (using port 80) without designing specifically for each piece of hardware; this is widely known as HTTP adaptive streaming (HAS). Modifying HTTP can solve for specific problems of adaptive streaming in HTTP, but the cost is deviating from an open-source implementation and needing both the transmission and receiving points to download additional code to handle this modified connection -- which is hard to implement consistently in an open general IP environment.

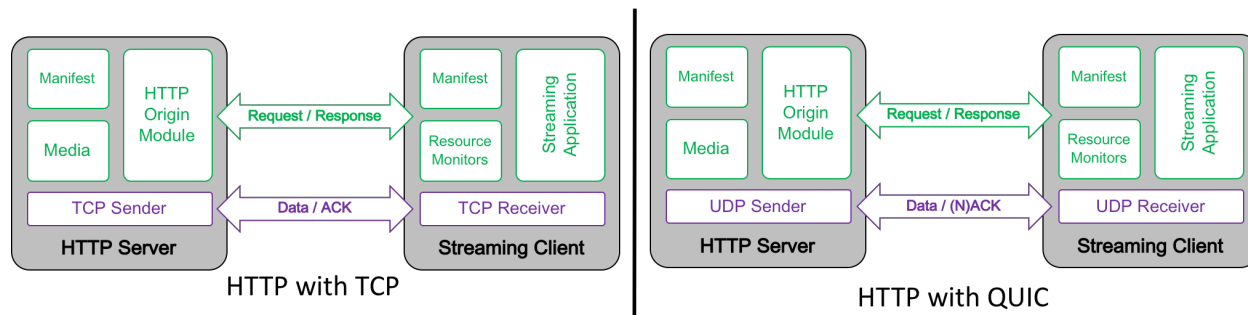


Figure 5: Server distribution using HTTP over TCP and QUIC.

A lower layer area that could be focused on would be the transport layer protocol. The protocol typically used with HTTP is TCP (Transmission Control Protocol) (see Figure 5), which is a protocol between two nodes to set up a connection and then send data packets over the internet. This type of protocol ensures that all data packets are received, in order, with a method to detect missing packets and request that they be sent again. For low-latency streaming, two features of TCP are highly sensitive to low-latency streaming. The first factor is that the re-request of a data packet must come in time to be useable by the buffer. With a low buffer for low latency, there are greater chances that the packet may not arrive in time, which means that the playability of the stream would get interrupted -- so this operation would fail more in low-latency situations anyways.

The second factor is that the packets must be received, in order, in the application buffer, before they are released to the player buffer (head-of-line blocking). These factors increase as best effort systems encounter more bandwidth congestion. There are several approaches to deal with this, but require a change in the protocol. An immediate approach would be to open up several TCP connections for each exchange, such that head-of-line blocking could be avoided. Another approach is to use UDP (User Datagram Protocol) instead of the TCP protocol, where UDP benefits because it does not attempt to even request a missing data packet.

An alternate but increasingly popular protocol that takes advantage of this is QUIC (Quick UDP Internet Connections), where the initial round-trip exchanges to set up are reduced (see Figure 5). Other approaches would be to send multiple data packets upon a single response with only negative acknowledgements sent. These types of modifications are being developed in new specifications such as the QUIC protocol and in the future HTTP/2 versions. Both versions require moving away from an already accepted open-source web server approach, and for that reason, acceptance of the improved protocol, and transitioning to them, may take time. Another way to mitigate these types of factors is to avoid bandwidth congestion or reconnections where possible. This can be done by adapting the bandwidth environment with respect to the particular service of the stream, which can then be handled in managed bandwidth environments more easily than in best-effort systems.

For low-latency live streaming, other alterations that would help would be to move closer to the live edge of delivery. For pre-packaged delivery, the priority is to provide robust streams to players to create an uninterrupted viewing experience, which involves requesting segments that are three segments behind the live edge. This ensures there is always a segment sent for each request, which reduces the risk of rebuffering. In live low-latency streaming, the priority may be to receive packets as quickly as possible, and the uninterrupted experience may be secondary to that. Using this premise, the use of manifests may change. For instance, the three-segment delay (to ensure availability of the segment) for manifests in adaptive streaming could be shortened. The use of predictive templates may be useable where the client may request future segments that have not been published yet, but will be. This, combined with error handling techniques at the player, may provide a useable live low-latency stream. Some of these player techniques to avoid rebuffering situations could be to use to repeat frames, skip frames, or apply temporary slowing of the frame rate, where the low latency of the stream is treated as a priority.

Conclusion

Reducing latency will improve the viewing experience in live event streaming. In current generation adaptive streaming situations, which were more designed to enable pre-packaged on-demand streaming, the priority was to ensure the robustness of the played stream -- which may add latency to the system, but improves the overall QoE. With a renewed focus on live streams, a balance needs to exist between robustness of the stream and a reduction of end-to-end latency.

There are three areas to reduce latency: 1) segment creation, 2) distribution, and 3) playout; it's a tradeoff between reducing latency, minding delivery quality, and accommodating scalability (especially for live events, which often generate flash crowds). Reducing latency cannot be done with just one change: It is a cumulative sum of incremental changes in all these areas. For content creation, it's about how to encode/decode the ES stream faster and yet still maintain quality. That requires integrating segment creation into delivery and distributing it quickly into the network, through a series of chunks that can be re-assembled into a segment or fragment. In distribution, it is a matter of how to get it to the edge servers quickly, while shortening the request/response communication between server/client, through the use of multiple TCP connections, or by moving to UDP-types of connections. Lastly, on the player side, the player should recognize that this is a low-latency stream where it is okay to anticipate segment requests, adjust the buffer for shorter durations, and allow for new playout behavior to keep the stream playable in a low-latency mode.

Working on reducing latency in the system will help in live event streaming, but will also generally aid in all segment-based deliveries (including pre-packaged delivery), by moving distribution closer to the live edge.

It is to the benefit of next-generation IPTV systems to integrate low-latency streaming into part of their design. This will improve performance of the next-generation IPTV systems.

Abbreviations

ABR	Adaptive Bitrate
AF	Adaptation Field

ATSC	Advanced Television Systems Committee
AVC	Advanced Video Coding
CDN	Content Delivery Network
cDVR	Cloud Digital Video Recorder
CMAF	Common Media Application Format
CPB	Coded Picture Buffer
DASH	Dynamic Adaptive Streaming over HTTP
DPB	Decoded Picture Buffer
ES	Elementary Stream
FFWD	Fast Forward
FPP	Forward Predicted Picture
GoP	Group of Pictures
HAS	HTTP Adaptive Streaming
HEVC	High Efficiency Video Coding
HLS	HTTP Live Streaming
HRD	Hypothetical Reference Decoder
I/O	Input/Output
IP	Internet Protocol
IPTV	Internet Protocol Television
ISBE	International Society of Broadband Experts
ISO/BMFF	ISO Base Media File Format
MPEG	Moving Pictures Experts Group
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
QUIC	Quick UDP Internet Connections
RWD	Rewind
TCP	Transmission Control Protocol
TS	Transport Stream
TTL	Time To Live
UDP	User Datagram Protocol
VoD	Video on Demand
SCTE	Society of Cable Telecommunications Engineers
NCTA	National Cable & Telecommunications Association

Bibliography & References

ISO/IEC 23009-1 Information Technology- Dynamic adaptive Streaming over HTTP (DASH) Part 1: Media presentation description and segment formats – Second Edition

ISO/IEC 23000-19, Information Technology- Multimedia application format (MPEG-A)- Part 19: Common media application format (CMAF) for segmented media- First Edition

ISO/IEC 13818-1:2018, Information Technology – Generic coding of moving pictures and associated audio – Part 1: Systems

SCTE 214-4 2018, MPEG DASH for IP-Based Cable Services Part 4: SCTE Common Intermediate Format (CIF/TS) Manifest for ATS Streams

SCTE 223 2017, Adaptive Transport Stream

Category-aware Hierarchical Caching for Video-on-Demand Content on YouTube. Christian Koch, Johannes Pfannmuller, Amr Rizk, David Hausheer, Ralf Steinmetz, MMSYS'18, June 12-15, 2018, Amsterdam, Netherlands

E. Thomas, R. Koenen, A. Begen, J. Boyce, “Streaming-First design for the MPEG-I project”, M43753-MPEG Meeting Documents, July 2018, Ljubljana, SI

A. Langley, A. Riddoch, A. Wilk, A. Vicente, C. Krasic, D. Zhang, F. Yang, F. Kouranov, I. Swett, J. Iyengar, J. Bailey, J. Dorfman, J. Roskind, J. Kulik, P. Westin, R. Tenneti, R. Shade, R. Hamilton, V. Vasilev, W. Chang, Z. Shi, “The QUIC Transport Protocol: Design and Internet-Scale Deployment”, SIGCOMM '17 August 21-25, 2017 Los Angeles, CA, USA

A. Begen, C. Timmerer, L. Ma, “Delivering Traditional and Omnidirectional Media”, ICME Tutorial, ICME 2018, San Diego, CA

T. Stockhammer, “Low-Latency DASH”, MHV 2018, Denver, CO, July 31st-Aug. 1st, 2018

Y. Shen, “Low-Latency Live Streaming at Scale”, MHV 2018, Denver, CO, July 31st-Aug. 1st, 2018

Digitizing the Customer Experience: Win Loyalty and Sell More with Last Mile Service Trackers

What to Do Now That You've Captured Subscriber Attention During the Last Mile of Service

An Operational Practice prepared for SCTE•ISBE by

Chris Ruff

CEO and President

Glympse

1424 11th Ave #300, Seattle Washington

(206) 237-1010

sales@glympse.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Repurposing Customer Engagement Dead Zones for Profitability	3
1. Learning Outcomes	3
2. From Last Mile Profitability to Last Mile Revenue	3
3. The Right Department for the Job	4
4. Required Equipment and Technology	5
4.1. Customer-facing field technician tracker	5
4.2. Embedded Web Analytics	6
4.3. Method for summarizing, analyzing and/or visualizing customer engagement data	6
4.4. APIs and webhooks to pass customer and task data	6
4.5. Content management platform.....	6
4.6. Ad Server	6
4.7. Digital ordering or ecommerce capabilities (for advanced projects)	6
5. Key Performance and Engagement Metrics	6
5.1. Success Metrics	7
5.2. Baseline Metrics Required	7
6. Calibration and Equipment Preparation	8
7. Detailed Procedure.....	8
7.1. Setting up basic integration.....	8
7.1.1. Collect data	8
7.1.2. Categorize and segment.....	8
7.1.3. Set Up Integrations	8
7.1.4. A/B Testing.....	9
7.1.5. Continuous Measurement and Trend Spotting	9
7.1.6. Define Time-Based and Seasonal Initiatives	9
7.2. Advanced options for consideration	9
7.2.1. Rich Communication Services	9
7.2.2. Gamification	9
7.2.3. Live Chat or Chatbots	9
7.2.4. Artificial Intelligence	10
7.2.5. Machine Learning	10
8. Recording Results	10
Summary of Lessons Learned	10
Abbreviations	11
Bibliography & References.....	11

Introduction

The cable industry has spent years solving for how to give customers waiting for an installation or service appointment better visibility and accurate information about when an appointment will start and end. This critical moment makes or breaks customer satisfaction. Today, operators are equipped to solve this last mile customer frustration problem with digital insight and accurate estimated time of arrival (ETA) predictions provided by customer-facing technician trackers.

Problem solved, but that doesn't mean operators can afford to stop innovating. So, what's next?

Operators that provide last mile customer portals and technician trackers with the goal of reducing inbound calls and customer no shows have *already* built the foundation for deeper customer engagement and revenue generation. In the quest for transparency, they've created a captive audience that's dialed in and paying close attention to a centralized web experience monitoring their technician's location and ETA, much like ride sharing apps have spawned crowds of rapt travelers staring at their screens as they wait in the airport pickup zone.

This is an (often missed) opportunity to introduce new experiences into the customer journey. Not only that, but digital marketing spend is predicted to reach nearly \$120 million annually in the U.S. by 2021. With a large part of that spend made up of digital display and video advertising, web-based technician trackers are an untapped channel that could reap the rewards of the digital investment trend.

This paper presents ideas for transforming this customer engagement dead zone into a highly engaging and interactive customer touchpoint.

Repurposing Customer Engagement Dead Zones for Profitability

1. Learning Outcomes

After reading this paper, the reader should possess a clear understanding of the four key elements of this operational best practice, and how to deploy those elements in their own operations.

Learning outcomes include:

- How to identify the right moment or moments during the critical last mile with the most potential for innovation
- Strategies for leveraging customer attention to deepen engagement, grow loyalty and create natural connections that guide customers to related digital branded experiences
- How to create and embed interactive experiences that make sense to customers and improve satisfaction
- Strategies for integrating sales and marketing programs that drive add-ons and additional revenue to take advantage of customer down time during the last mile

2. From Last Mile Profitability to Last Mile Revenue

Jobs per day. First time fix. Wasted truck rolls. Rerolls. Then there's mileage, optimized routing and service parts management. These are the metrics cable operators strive to minimize every day, every hour

for field services. In an industry facing rising costs and customers expecting more amidst unprecedented competition from OTT services, it's no wonder executives are constantly seeking new and creative ways to reduce field service operations expenses. More recently, best-in-class organizations have focused on driving better customer service around field service as a path to better margins. The connection to efficiencies: happier, more informed and less frazzled field service customers are less likely to cancel their contract or service, less likely to miss an appointment and less likely inundate customer support channels looking for answers or apologies. A step further: numbers show that informed customers who keep their appointments begin generating revenue immediately. Empowered customers with the flexibility to change a field service well in advance are less likely to churn.

Many best-in-class organizations also connect customer-focused field service to customer loyalty and/or lifetime customer value (LCV) and have focused transformation strategies on this approach. If field service appointments are the only in-person interaction a customer has with the business, this last mile of service delivery is what they'll most likely judge the entire brand on. When it is time to reach out again – whether to seek support, buy more or extend/renew a service – that field service interaction is the one they'll remember.

Data supports this approach. A recent NewVoiceMedia study that found that 49 percent of consumers switched vendors due to poor customer service. This number increases for 24-34-year-olds, where 62 percent of consumers have switched to another business because of subpar customer service.

While reducing costs and improving customer satisfaction both contribute to healthier margins and long-term revenue growth, field service arms will succeed faster if they can play a role as an active revenue generator for the business *today*. Strategies that focus solely on continually cutting costs and improving service levels as a means to improve profitability will only create expectations from ever-cheaper, ever more “delightful” service from management and customers alike. It perpetuates the vicious cycle of pressure to do more with less. There are more direct paths to generating field service revenue, and penny-pinching will do little to solve the larger industry threat.

3. The Right Department for the Job

There are plenty of reasons field service leaders should actively pursue revenue-generating initiatives. There's the obvious:

- To reach personal and departmental goals
- To establish safety, security and even notoriety for your department within the broader organization despite mounting competitive pressure
- To grow the size of your department and influence by demonstrating success

Perhaps more importantly, leaders that present demonstrable results are more likely to win funding for their own priority projects faster. For example, showing the connection between how securing the latest technology and training for service technicians can result in a new revenue stream during the last mile of home services and installations is a creative way to bundle strategies for holistic modernization. Well-intentioned projects often get stuck behind major system upgrades and tech debt that can take months, or worse, years to complete. Good luck prioritizing employee retention programs in this environment.

Even when there is a strong business case for an innovative project that can also save the business money, that project is often “funded” by cutting the budget of the team that is leading it. Build a case for an important and exciting new project around new *revenue generation* rather than *cost savings* instead.

There is one reason that stands out among all others as to why the same team that owns and executes the last mile the customer's journey – whether for a new installation or a service – should be driving new proactive engagement, marketing and sales initiatives.

Field service teams already intimately understand the nuances of the critical last mile of the customer journey; field service departments have engagement data and they own the direct, two-way customer communication and interaction that happens around the most emotional touchpoint in a customer's lifetime journey.

A field service interaction is the first, and often the only, face-to-face interaction many operators will have with their customers. It's a deeply personal interaction because it is in a person's home or workplace, and it's often happening at a critical moment when customers form strong perceptions of a brand – at the start of the relationship or when something is broken or wrong.

Considered in context of the holistic customer journey, field service owns what happens at a pivotal moment in the customer relationship. It's a prime opportunity to stop thinking of field service as an operational necessity with costs to be minimized and instead, to start positioning it as a lucrative engagement channel. This becomes even more powerful for field service leaders who can transform not just the on-site field service experience, but every point of customer engagement leading up to and immediately after that field service meeting. It's a chance to own and transform a major gap in the typical customer journey for the better.

The following sections explore the required technology, processes, performance and engagement metrics required for foundational planning. Then, we present an outline for the detailed procedure and best practices for collaborating with marketing and sales to go from last mile revenue vision to last mile revenue reality. It all starts with customer-facing technician trackers.

4. Required Equipment and Technology

This operational best practice requires one primary platform and subsequent data as a launching point, as well as several supporting technology platforms for integrating sales and marketing into customer-facing last mile engagement.

4.1. Customer-facing field technician tracker

As the introduction of this paper clearly articulates:

Operators that provide last mile customer portals and technician trackers with the goal of reducing inbound calls and customer no shows have already built the foundation for deeper customer engagement and revenue generation. In the quest for transparency, they've created a captive audience that's dialed in and paying close attention to a centralized web experience monitoring their technician's location and ETA.

Thus, the operator must have a customer-facing technician tracking experience in place. Ideally, this solution should be web-based (though an app-based experience can work with additional development), it should be interactive, and it should be flexibly architected for easy communication with other customer experience platforms and systems of record. Ideally this experience will be multi-phased in nature, with the ability to update/change content and engagement options served to the customer, corresponding with various stages of the last mile of service delivery.

4.2. Embedded Web Analytics

Google Analytics is a common and popular tool. Regardless of which tool is used, this approach requires web page analytics. Data about customer engagement and activity on each web page of the journey – ranging from clicks to user flows and heat map data – is what will form the foundation for building personalized, targeted marketing programs. Tools like Piwik go so far as to connect web page activity to marketing campaigns. Others like CrazyEgg, Mouseflow and Clicktake incorporate eye tracking heat map data and to even visualize in-page behavior.

While these tools will all be familiar to marketers, the innovation is in understanding how these same tools can be applied to collect engagement data during service delivery experiences, whereas they are traditionally only used to predict and measure shopping and initial buying behavior.

4.3. Method for summarizing, analyzing and/or visualizing customer engagement data

Whether it's an integrated console or raw data in excel format, users will need to set up a process for collecting and organizing the data collected from web page analytics and behavior for their last mile technician tracker. The ability to segment and compare this data by individual steps and stages in the customer-facing journey is ideal.

4.4. APIs and webhooks to pass customer and task data

These are required in order to connect individual customer preferences, previous behavior, collective trends and new offers to one another, as well as to make actions available directly in the technician tracker and/or customer platform.

4.5. Content management platform

Regardless of final decisions for segmenting customers for campaigns and A/B testing, a content management, content marketing or marketing automation platform will play a vital role in uniting customer data with appropriately timed and themed offers.

4.6. Ad Server

This may be the same platform or a platform that's connected to your content management or marketing automation tool. At the most advanced level, this will be a predictive tool that defines which offer to serve to which customer or segment at an ideal time or based on certain engagement triggers

4.7. Digital ordering or ecommerce capabilities (for advanced projects)

Related to ensuring the appropriate RESTful APIs and web hooks are in place, if the strategy includes an offer for an on-site upgrade, a method for completing the order and collecting payment within the same screen as the technician tracker is necessary.

5. Key Performance and Engagement Metrics

Two primary metrics will determine overall success – revenue and customer engagement. However, multiple last mile customer engagement metrics will be critical both to establish a baseline and in order to effectively categorize customers into segments for personalized ads and offers.

5.1. Success Metrics

As this best practice is focused on generating new revenue streams via customer-facing technician trackers and appointment tracking portals, revenue is the primary success metric. Some variations to consider are:

- change in average revenue per customer (ARPU)
- installation revenue
- service cycle revenue
- incremental revenue per customer

Secondarily, engagement metrics can be useful to determine which offers and new revenue initiatives need revising. This is where web page analytics will be crucial to track which offers are least effective, the ideal timing and on-page placement/location for offers and which offers miss the mark on final conversion. For example, which offers do customers often click on but don't confirm purchase? Is that because they don't perceive the offer to be valuable, or because the process for completing the order and payment are too cumbersome? Or, was the offer surfaced too early or too late in the journey?

5.2. Baseline Metrics Required

Prior to measuring success, revenue programs must be devised and implemented. This will require a clear understanding of baseline engagement for each step or stage in the web-based, customer-facing journey. Data about which elements of this journey receive the highest engagement and when will also help inform decisions on sales, marketing and eCommerce programs.

Start with understanding views per stage of the journey. That is, how many views does each stage in the technician tracking experience receive. Which stages receive the most views, and the least? Does this data vary by season, job type or customer type? How long does each customer – as well as the average customer – spend engaged with each step of the customer-facing experience? Some may be extraneous, and some may be more critical than originally anticipated.

Next, document and analyze engagement with various interactive elements for each journey stage. Common elements include:

- Add to calendar widget
- Technician or field engineer details
- Buttons to chat, call or reschedule
- Real-time live map tracker
- ETA countdown
- Feedback capture
- Downloadable PDF installation or services summaries
- Options to share feedback or comments to social channels

Understanding how customers interact with these elements and when each element is most popular will provide clues and context as to what types of offers can be most beneficial at various phases of the last mile service journey. Even deeper, they may provide inspiration for entirely new lines of service that existing customers may buy into. Finally, consider contextual language analysis of comments provided via feedback forms and surveys to help identify opportunities to change the experience – for example, adding a new page to the live journey that would entertain them as they track their technician.

6. Calibration and Equipment Preparation

No calibration of equipment is required. Some ongoing iteration of engagement metric parameters will be useful, as well as some adjustments to proactive ad services. This is covered below in detailed procedure.

7. Detailed Procedure

7.1. Setting up basic integration

7.1.1. *Collect data*

Set up the recommended web page analytics tools to collect engagement data from your customer journey. Customer-facing technician tracker experiences have resulted in customer engagement of five to more than ten minutes for each appointment. The goal is to leverage web page analytics tools to gather more nuanced data about which steps in the journey customers are most interacting most and least with, which elements of the web experience they are interacting with versus ignoring, and how these measurements change as the scheduled appointment gets closer.

7.1.2. *Categorize and segment*

Best practice for this step is to collaborate with marketing and/or sales teams. Field service teams should start by parsing and analyzing data for engagement at each progressive step or new screen in the technician tracking experience. However, marketing can provide insight to help further segment engagement data based on demographics, customer history, etc.

7.1.3. *Set Up Integrations*

Next, consider which platforms will need to be integrated with the tech tracking experience. The best practice is to start with something simple and get progressively more complex as processes are refined. For example, a common first step is to integrate a marketing platform or ad server with the tracking experience, essentially enabling the brand to surface a targeted advertisement at various steps in the appointment tracking lifecycle. Some test a static ad throughout the entire journey, while others test different advertisements at different steps of the journey, and even immediately after a service is complete during a feedback capture phase.

Examples of more advanced integrations include interstitials, links to my loyalty or my account programs that encourage customers to try new services, etc. eCommerce solutions can be embedded directly as a widget in a technician tracking portal or web view so that customers can “add-on” and even complete payment and verification for incremental products. The simplest example is the offer for a customer to add a sports or premium channel package for a discounted rate, with that rate only valid while the technician is on site. This type of intelligent and sophisticated upsell solution requires seamless bi-modal communication between the technician tracking tool and the platform managing transactions, for example.

Organizations with advanced marketing automation and digital personalization tools in place will also need to decide whether to define offers based off average customer data or rather, to surface predictive offers based on a unique customer profile.

7.1.4. A/B Testing

All that said, consider testing multiple methods of segmentation and personalization. A/B testing can be as broad as offering two entirely different services to different customer segments or as granular as testing ad copy performance with those segments. It may be difficult to conduct A/B testing at the individual level only because the frequency of home service visits is low. This means, testing should be robust and modeling based on similar customers and similar engagement behaviors will be key to capitalizing on each individual's engagement during the last mile tracking experience.

7.1.5. Continuous Measurement and Trend Spotting

As typically follows from A/B testing, ongoing analysis of results will be critical. As previously stated, each customer may have only one in-person engagement with a cable operator every one to two years. Relying on aggregate data of previous successes and failures, identifying trends and testing potential motivation for behavior will mean the difference between success and failure. With this method, unlike email marketing or digital advertising, there won't be multiple opportunities to get the offer right. Data will be the foundation of successful programs.

7.1.6. Define Time-Based and Seasonal Initiatives

Don't overlook the opportunity to incorporate seasonal offers along with personalized ones. Seasonal offers are also a good, simple place to start for those not equipped to support targeted or segmented offers to customers directly within a technician tracking tool – whether that's because they lack the marketing tools or the buy-in and necessary collaboration from other parts of the organization. Seasonal offers can be a good testing ground to demonstrate results and drive for more advanced programs.

7.2. Advanced options for consideration

7.2.1. Rich Communication Services

Rich communication services (RCS) are becoming increasingly popular and mainstream. RCS offers the ability to visualize not only various steps and stages of a last mile tracking service directly in message, but to add interactive elements. This is a prime channel and flow to test interactive ads and new offers.

7.2.2. Gamification

Those that take the care to analyze tracker engagement by stage or step in the appointment process may come to realize that customer engagement ebbs and flows in relation to the scheduled time for the pending appointment. For example, a customer will likely spend less time viewing a "pre" type phase that establishes a broad ETA and allows a customer to add an appointment to their calendar, and more time actively watching a live map view of a technician's truck on a map. This may mean that ads and offers at earlier stages of the appointment lifecycle will be less frequently viewed and thus less effective. This is where advanced teams may get creative with gamification or other interactive elements they can insert into the web experience in earlier stages in order to drive more time spent on these web pages.

7.2.3. Live Chat or Chatbots

Similar to gamification, the chance for customers to chat with an agent or even with a chat bot service directly within the tracking experience can be highlight valuable. Beyond answering questions and offering support about the upcoming home service, this is a massive opportunity to engage with customers who are considering complementary services.

7.2.4. Artificial Intelligence

Taking it a step further, artificial intelligence can add a predictive layer to these experiences and increase the probability of converting new sales. The key, going back to the required technology component of this best practice, is to ensure that the technician tracking experience you are building around is open, flexible and can easily embed or share data with partner solutions via the appropriate APIs, software development kits (SDKs) and web hooks.

7.2.5. Machine Learning

Further, applying machine learning algorithms with in web page or unique view of a customer journey can offer a path for much more unique and sophisticated offers at the individual level so much that it's possible to offer different choices, offers, advertisements or engagement paths depending on actions a customer takes to engage with distinct elements of the technician tracker tool.

8. Recording Results

Recording of results should be simple and straightforward. Following the guidelines established during the detailed procedure to establish parameters, track results in a central dashboard, console or if less advanced, in raw data format. Clearly define engagement phases and clearly label engagement metrics by phase. Data should then be mapped back to customer records, which will be associated with individual tasks corresponding to field service actions. This framework will provide the foundation you will use to map marketing content management and personalized ad services to each customer, during the appropriate phase or status of each task. Finally, track success over time in terms of revenue and engagement as defined for the business (using the defined success metrics) – as well as recording the impact of each change on all results.

Summary of Lessons Learned

The reader should now possess a clear understanding of how to apply web analytics tools and data analysis to identify the right moments for new engagements. This means the reader should be equipped to start planning how to actually leverage a subscriber's attention as a technician travels to their home for a service, in order to create a measurable new revenue stream. Identifying the right moment or moments during the last mile with the most potential for innovation requires thoughtful collection and segmentation of data, and a creative approach for viewing customer behavior. Most importantly, success with this operational best practice requires tight alignment and close collaboration with sales and marketing departments. IT will also play a crucial role in connecting the required systems to offer a seamless experience.

It's possible to create and embed interactive experiences that make sense to customers and improve satisfaction by first understanding them, and then learning and responding to how they react to new offers and interactive components of what used to be a dead zone for customer communication and engagement.

Abbreviations

API	application program interface
ARPU	average revenue per customer
ETA	estimated time of arrival
LCV	lifetime customer value
RCS	rich communication services
SDK	software development kit

Bibliography & References

<https://www.forbes.com/sites/shephyken/2016/08/27/bad-customer-service-costs-businesses-billions-of-dollars/#54bb86165152>

<http://www.iacquire.com/blog/15-google-analytics-alternatives>

<https://www.forbes.com/sites/forrester/2017/01/26/us-digital-marketing-spend-will-near-120-billion-by-2021/#4a09cdd7278b>

Edge Compute and Software Life-Cycle Management

Creating Consumer Value and Flexibility

A Technical paper prepared for SCTE•ISBE by

Patrick Goemaere

Chief Architect Cloud Services Connected Home
Technicolor
942 Birmingham Rd 91504 Burbank US
+1 (818) 442 7183
Patrick.goemaere@technicolor.com

Rajat Ghai

VP Wireless & Open Networking
Technicolor
rajat.Ghai@technicolor.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	6
1. Compute Technology Cycles	6
1.1. From the Cloud to the Edge	7
1.2. Agility at the Edge, Learning in the Cloud	7
1.3. Edge Compute Financial Model	8
1.4. The Edge : A New Frontier of Next Generation Services & Experiences	9
1.4.1. MSO Edge Computing use cases: Residential	9
1.4.2. MSO offered Edge Compute Enterprise – B2B use cases	10
2. Edge Compute and decentralization	12
2.1. IoT platforms	12
2.2. Edge compute	14
2.2.1. Latency	14
2.2.2. Bandwidth	15
2.2.3. Security	15
2.2.4. Agile developer friendly IoT Edge devices.....	16
2.2.5. Autonomous operation on the Edge devices	17
2.2.6. Privacy	18
2.2.7. From cloud to Fog and Edge endpoints	18
2.3. Examples of IoT Edge Platforms.....	19
2.3.1. Azure IoT Edge	20
2.3.2. AWS Greengrass	21
2.3.3. Resin OS.....	22
2.4. Life Cycle Management of IoT Services, Development And Deployment	22
2.4.1. DevOps Practices to the Rescue	23
2.4.2. Securing IoT Devices the Agile Way	23
3. Lightweight Virtualization – Containers	24
3.1. Container Technologies	25
3.1.1. Machine/System Containers (LXC/LXD)	27
3.1.2. Process Containers.....	28
3.1.3. Application/Desktop containers.....	31
3.1.4. Serverless containers	32
3.2. Container standarization	33
4. MSO CPE landscape & CPE device characteristics.....	34
5. The New Embedded Stack	35
Conclusion.....	38
Abbreviations	39
Bibliography & References.....	40

List of Figures

Title	Page Number
Figure 1 - The Sinusoid of Processing	6
Figure 2 - Analytics Functions And Characteristics	8

Figure 3 - Financial Model Edge Compute	9
Figure 4 - Technology Convergence.....	13
Figure 5 - E2E IoT Functional Architecture.....	13
Figure 6 - Security vulnerabilities for IoT devices	16
Figure 7 - Developer Communities	17
Figure 8 - IoT GW for autonomous operation	18
Figure 9 - Edge/Fog computing.....	19
Figure 10 - Azure IoT Edge.....	21
Figure 11 - AWS Greengrass.....	21
Figure 12 - Resin OS	22
Figure 13 - Developer friendly workflows.....	24
Figure 14 - Virtual Machines Versus Containers	25
Figure 15 Container Building Blocks.....	26
Figure 16 - Container Categories.....	26
Figure 17 - LXC Containers	27
Figure 18 - Docker Building Blocks.....	29
Figure 19 - Docker packaging and image format.....	29
Figure 20 - Docker Containers and images	30
Figure 21 - Docker ecosystem	31
Figure 22 - Edge compute use-case domains	35
Figure 23 - Modern LCM layer for CPE	38

Introduction

Over the years, we have witnessed a fundamental shift in how software gets developed and deployed in the cloud. Born-in-the-cloud web companies have moved to a much more agile way of working where continuous deployment and integration became the norm, and where these companies are able to introduce new functionality and features on a daily basis, with fine grained control to limit the exposure of these feature first to a subset of their customer base using deployment strategies like canary releases or a derivate like blue/green and other variants.

This transformation has helped these companies to create a competitive advantage so that they can rapidly check the viability of new functionality, collect user feedback and respond quickly to changing requirements. The introduction of full automation and DevOps methodology, has furthermore improved the complex social interaction that typically has plagued many complex product and service development in large scale organizations by breaking up traditional boundaries between product and business owners, R&D, IT and operations so that products can quickly move from business requirements to deployment and has allowed more flexibility in trying out new ideas in a fail fast way of working.

The innovation around Internet of Things (IoT) has built further on the evolutions started in the Web and mobile domain on providing the foundational technology building blocks on how to communicate with embedded devices at web scale, and in order to keep solutions viable, has pushed the boundaries of these technologies further and further to the edge of the network.

This evolution of decentralization has blurred the traditional demarcation between operational embedded technologies and IT technology. On top of these newly cloud native implementations of ultra-web scale communication/messaging infrastructure services, a new generation of services like software life cycle management (LCM), configuration management and orchestration are being developed, that borrows the same concepts as their well-established cloud DevOps counterparts but addressing better the enormous scale of IoT deployments, and their very heterogeneous and fragmented nature.

Our current way of working in the embedded software device space differs today fundamentally from this agile cloud methodology, due to the very nature of these devices themselves. Because of their very constrained nature in terms of CPU power, memory and storage, hard real time requirements and cost structure, these solutions tend to be optimized for very specific functions leading to monolithic firmware development running in specialized execution environments, typically stripped-down Linux distributions. Embedded devices lack the uniform compatibility environment that has been addressed in the cloud by virtualization of compute - either virtual machines (VM's) or containers, and networking. Typically, these development cycles are complex, lengthy, and require lots of validation around industry specific certifications, interoperability and standardization. Furthermore, the reluctance of introducing operational impact on the installed base hinders the deployment frequency.

This paper will focus on a subset of device types used today in the Cable industry as Customer Premise Equipment (CPE) equipment, more notably internet gateway's and set-top boxes and their very specific nature, and tries to identify the current container technology landscape and their applicability for these types of embedded CPE devices.

It proposes a dual track approach with respect to separating out critical firmware functions and more agile customer facing functionality, avoiding a big bang approach and creating an architectural evolution path for the current existing legacy deployments.

It will further elaborate on the evolution of Internet of Things technology of the last years and it's evolution from a centralized approach, towards a decentralization of this architecture towards the Edge of the network and the profound effects it has on improving user experience and customer intimacy for the next generation of IoT services, and the disruption it causes on the entire value chain.

Content

1. Compute Technology Cycles

Computing architectures have always been optimized for most efficient workload placement economics




Figure 1 - The Sinusoid of Processing

First Wave

At the dawn of computing, a centralized structure in the 'mainframe' age was most economical due to very expensive compute and transport cost structure. This was the pre-microprocessor era of mainframes and mini computers built using Solid State semiconductor Devices leveraging the invention of the Transistor in late 1940s.

Second Wave

Computing architecture changed in early 70s with the invention of the Microprocessor and the arrival of Personal Computers; Moore's law kicked in and the cost efficiency economics moved to a distributed 'client-server' model. In this golden age of personal computing, which lasted for the two decades (1980s and 1990s), consumers and enterprises adopted local premise computing, primarily because the cost of computing was significantly lower than the cost of Internet transport. Internet speeds in those decades were slow and expensive. This is also considered the time associated with the dawn of public Internet. Internet Service Providers (ISP's), notably Multiple System Operators (MSO's), made large public/private investments which led to faster and more affordable Internet connectivity.

Third Wave

At the turn of the 21st century, with the arrival of innovations and investments in affordable high-speed wireline and wireless Internet access, computing architectures evolved once again. With both computing and Internet access being now affordable, cloud computing was born. Cloud computing (though similar to a centralized design of Mainframe) was a centralized co-location of compute, memory and storage clusters offered as Infrastructure as a Service (IaaS) by public cloud providers like Amazon, IBM, Microsoft and Google, followed by a long tail of smaller providers. Cloud computing for enterprise workloads created immense cost efficiencies in enterprise IT. Cloud computing has given rise to a radical transformation of traditional enterprise technology over the last 10 years. Due to the fully distributed and virtualized nature of cloud, a new generation of cloud native solutions have arrived, architected to cope with resilience in a world where failure is the norm and moved to horizontal scalability using commodity hardware instead of scaling vertical using expensive hardware.

Fourth Wave

Now, in 2018, computing is at the doorstep of another epoch. Computing has become ubiquitous and miniaturized to a point that it can be embedded in virtually every physical object. At the same time wireline and wireless Internet connectivity continues to follow economies of scale. This has led to a new universe of ‘smart things’ (*things that can compute and can network*). This radical shift is referred to as the Internet of Things (IoT). The world has started utilizing IoT for the digital transformation of the businesses and society as a whole. Every known business process is being digitized, automobiles are becoming autonomous vehicles, Internet connected smart sensors are helping automate the world around us and the society as a whole, is becoming digital as a result. It is expected that by 2025, there will be up to 25 Billion Internet connected Smart Things, far exceeding the number of humans on this planet.

As sensor technologies and networked services pervade every industry, the gravity of *so much data generated by so many diverse endpoints* renders centralized computing topologies inadequate.

1.1. From the Cloud to the Edge

In vital industrial sectors like agriculture, aerospace, mining, healthcare, or energy, reliable connectivity can be sparse; devices and equipment generate massive amounts of data chatter placing unnecessary demand on networks and battery life; failure can cost significant revenues, time, even lives. Because the cost of bandwidth is not falling as fast as the costs of compute and storage, the ability to aggregate and filter data *locally on the device* also serves an economic justification: edge compute decreases reliance on bandwidth. Furthermore, when critical real-time decision-making is required (think aircraft in-flight or self-driving cars), latency to the cloud is simply unacceptable. In consumer markets, current IoT implementations have fallen short of expectations—unmet needs for interoperability, inadequate monetization for providers and poor user experience (UX) for end users. Moreover, privacy concerns stifle consumer adoption and related, if highly variable, regulatory compliance requirements mean some data are simply better left local on the device.

1.2. Agility at the Edge, Learning in the Cloud

Instead of relying purely on cloud infrastructure to process data, local compute ‘at the edge’ enables greater speed, flexibility, security, privacy, economy, scale, and most importantly *learning*. To enable this, organizations need the tools to aggregate, define, and filter data in order to process lower-value data locally, while uploading high-value data back to the cloud for analytics, innovation, and more sustainable storage and data management. They need partners to enable intelligence in the body (edge), while taking

greater advantage of intelligence of the brain (cloud). The following picture shows the balancing of different analytic functions amongst Edge, Fog, and Cloud.

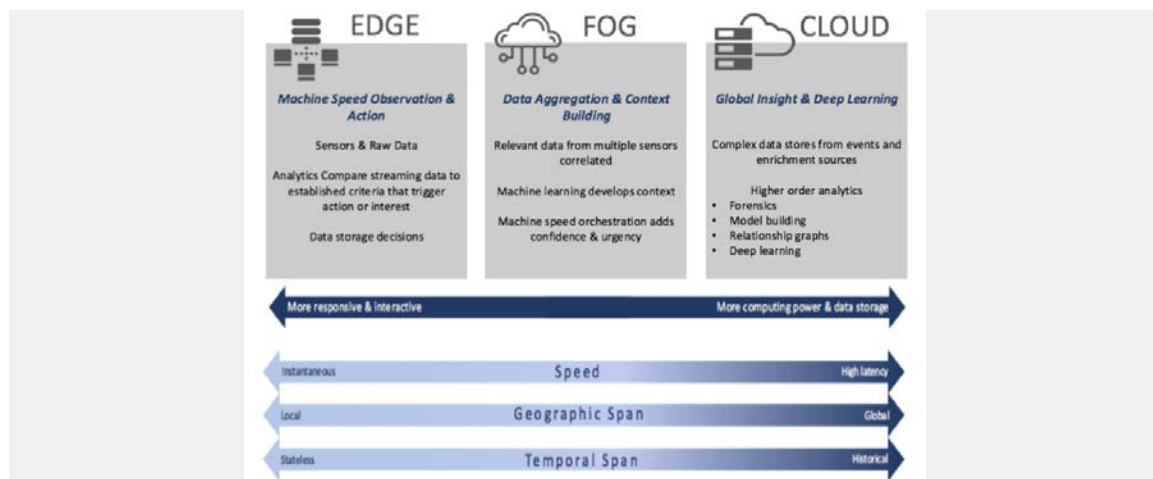


Figure 2 - Analytics Functions And Characteristics

1.3. Edge Compute Financial Model

Shown below is a case study (Financial model Edge compute, n.d.) done by a research firm (Floyer) for a wind farm site equipped with 100 sensors and two video streams. It compares the total cost of managing and processing on three different kinds of architectures:

1. cloud-only processing with a dedicated network,
2. A Mobile Network Operator (MNO) cellular network with hardware and cloud processing, and
3. An edge-and-cloud processing with a dedicated network.

In the first scenario, the three-year Total Cost of Ownership (TCO) of the solution including cost of transmitting the data, plus cloud costs and equipment costs, worked out at **\$254,552 /yr.**

In the second case the three-year TCO for a MNO hosted solution was less than 50% of the first scenario came to **\$113,884/yr.**

However, with a combined edge-computing and cloud approach the 3-year TCO dropped to 15% of the first scenario and 33% of the second scenario, to just **\$37,628/yr.**

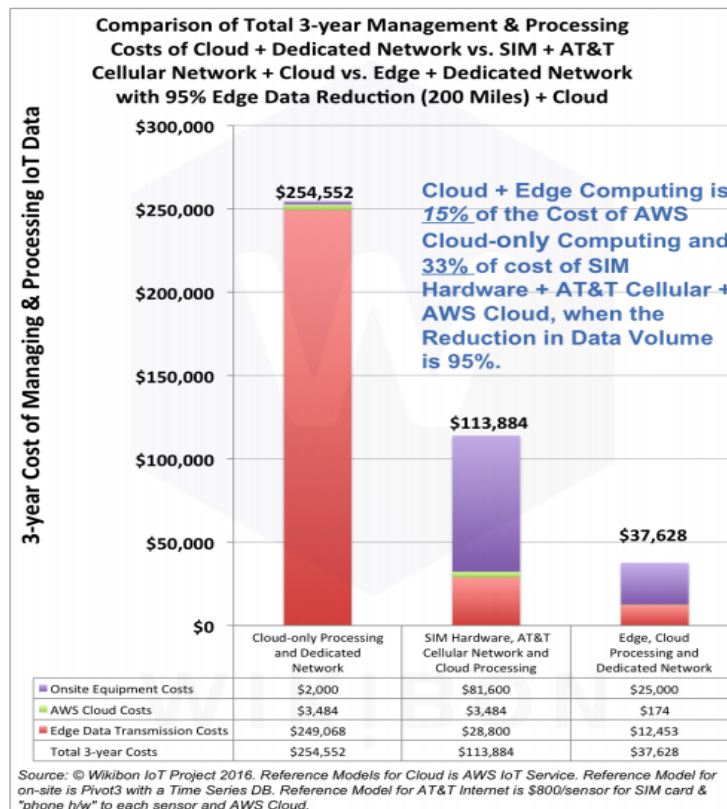


Figure 3 - Financial Model Edge Compute

As shown in the case study, the economics work best when Edge computing is not used to replace centralized cloud-based infrastructure in its entirety, but rather to complement it by increasing the computing and storage resources available on the edge by adopting platforms that provide intermediate layers of computation, networking, and storage and offloading light weight computational workloads locally at the edge and executing computationally intensive workloads (e.g. machine learning workloads that require Graphics Processing Unit (GPU) etc.) in the centralized cloud.

1.4. The Edge : A New Frontier of Next Generation Services & Experiences

While the move to the edge was triggered by the practical realities of the physical world and the tactical needs for real-time data processing, the broader implication of this shift is its role as a force multiplier of new services and experiences. New use cases abound for both end user and enterprise adopters.

1.4.1. MSO Edge Computing use cases: Residential

Seamless Plug & Play: Activating new devices and connected services is finally easy, really easy. When gateways, for example, are voice-supported and equipped with automated detection of nearby devices, users' onboarding experience is a matter of simply 'powering-on' and speaking. Behind the scenes, the edge-enabled gateway handles the rest, automatically provisioning devices, finding Wi-Fi, identifying and executing updates, and activating new services.

Autonomous Systems Administration: Instead of forcing users to become IT systems admins, edge compute-enabled devices handle the cumbersome but critical burdens of performance and network monitoring, anomaly detection, identity authentication, permissions, and encryption, even facilitating better connectivity and quality of service by automating channel changes, roaming decisions, and Wi-Fi optimization.

User-Centric Security & Privacy: Finally, a technology that allows users to have agency and peace of mind around their most sensitive data. When devices are equipped with access management, encryption, and security capabilities right out of the box, data is never exchanged between products without authenticating identity. Moreover, businesses and users can easily configure data filtering and transmission rules, so that highly sensitive data can remain private, local to the gateway, and not sent to the cloud.

Single-Stop Service Hub: When a gateway is built with onboard compute capabilities powerful enough to coordinate, monitor, configure, and control multiple devices, new modes of self-service and support open up. Through voice-interaction, users can simply speak to control and configure their own network services, and solicit support requests through a convenient single-stop self-service portal.

‘Getting Better All the Time’: The opportunity for innovation enabled through edge-level intelligence means products gain in value over time. For end users, this translates to a single hardware investment which continuously delivers new features and services, constantly learns and adapts to user speech patterns and preferences and evolves to support new use cases through open interoperability with any other device... on the body, in the home, or on the go.

1.4.2. MSO offered Edge Compute Enterprise – B2B use cases

Automatic Implementation: When in-field devices are outfitted for more robust local processing, enterprises enjoy a streamlined device installation process. Devices can automatically provision, certify, pair, and register themselves into Managed Service Provider (MSP) and Customer Relationship Management (CRM) systems.

Automatic Service Activation: These same out-of-the-box configurations mean enterprise user registration, certification, and customer service activation are ‘baked in’ to the onboarding process. Such a gateway, for example, can automatically connect and troubleshoot to onboard nearby devices, and offer voice-enabled pairing for additional devices.

Managed IoT Cloud Services: Edge-level compute opens up new ways service providers can add value to their partners and customers through managed services. Real-time monitoring of multi-device performance, application performance, network activity, security features (like vulnerability and intrusion detection), and device compliance, eases the burden for enterprises and allows them to focus on their core competencies. Service providers can also use such a gateway to manage new feature releases, API improvements, firmware, and other updates so innovations are securely folded into the broader device ecosystem. These capabilities render it economically feasible to create true end-to-end, personalized, and secure contextual experiences with end-users.

Integrated Support Services: Enterprises’ product and support programs also benefit from edge devices’ ability to run preventative maintenance such as patch management, bug-fixing, anti-virus updates, and better manage incidents through remediation, customized help desks, and integrated workflows and analytics via Customer Relationship Management (CRM) integration. Management in areas like compliance and refurbishment also help with end-of-life support.

Unified Analytics: When devices are enabled to interact in real-time, enterprises can parse and batch data to prioritize agility, while still gathering insights across data sets to learn over time. Integration with Enterprise Resource Planning (ERP), CRM, Finance, and other systems doesn't just enable enterprises to unify analysis and machine learning with a single dashboard, but also facilitates new use cases like predictive maintenance, automatic hardware replacements, and invisible improvements to performance efficiency.

Baked-In-Innovation: In the construct of balanced intelligence, where edge devices handle agile processing and new features, and deeper learning and data mining occurs in the cloud, enterprises are poised to support improvement across every level of the product (and therefore *customer*) lifecycle. Onboarding is improved over time as machine learning helps deliver better, more proactive support and troubleshooting to offload call center interactions. Product and service performance gain in efficiency as analytics and machine learning reveal areas to optimize such as connectivity, data transmission, and compute resources. Most importantly, such a construct embeds innovation into service architecture, wherein open development environments enable new and diverse services to flourish

2. Edge Compute and decentralization

2.1. IoT platforms

Designing, implementing, securely operating, managing and maintaining IoT projects are complex. In fact, there are entire organizations whose sole mission is solving a specific problem within an IoT architecture. The problems that can be found within such architectures can range from connectivity to figuring out where apps live.

Here are just some examples:

- Solving the daunting challenge of how to connect various very fragmented sensor protocols to the Internet.
- Handling intermittent connectivity, let alone trying to update an application over the air or patch a remote edge system
- Predicting when something is going to break before it happens
- Correlating telemetry with contextual data to generate better models
- Storing different information in different databases (NoSQL, time series, in memory database)
- Making sure your architecture is scalable, flexible, and can be deployed anywhere
- How to abstract your developers from where their applications will be running in this chaotic, highly distributed world

But these are only pieces to the bigger puzzle of designing an end-to-end IoT system where connected devices, business processes, DevOps practices, and people collaborate together across four different, complementary, but most-of-the-time disconnected, worlds:

- Operational Technology (OT)
- Information Technology (IT)
- Analytics and Machine Learning (ML)
- Traditional and modern application development

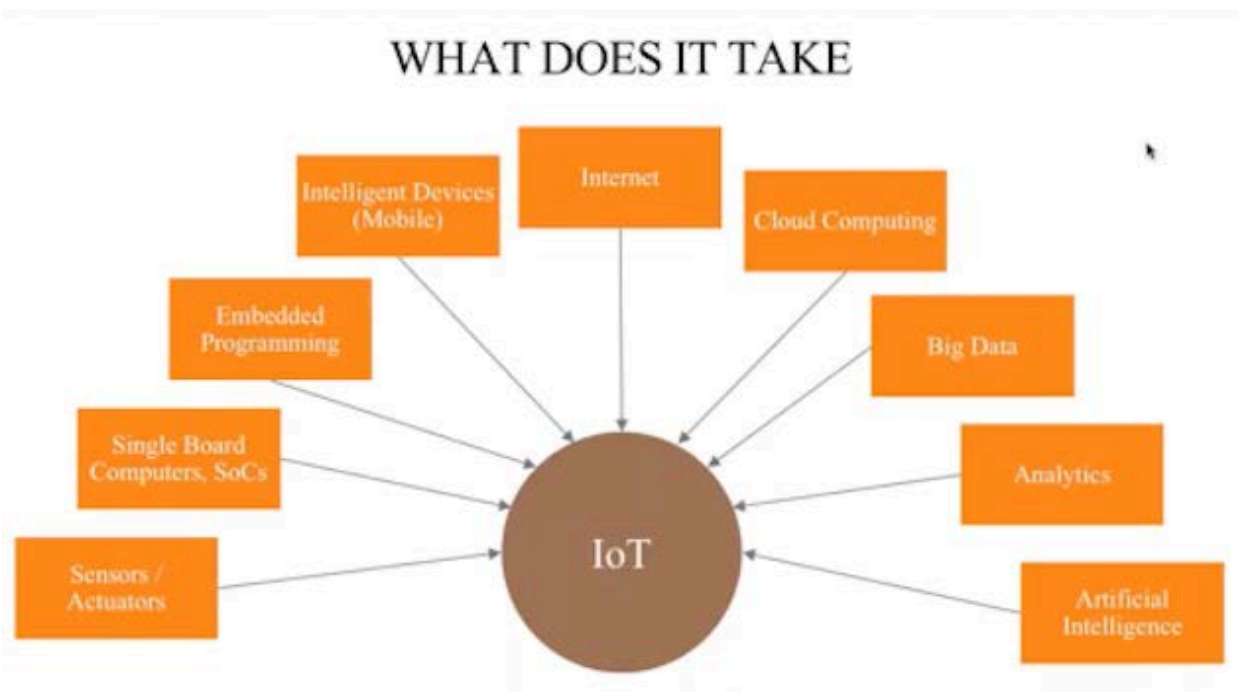


Figure 4 - Technology Convergence

Connecting devices; ensuring communication protocols are translatable; verifying accuracy and security features are running across the entire network; capturing, managing, analyzing, and using data to create better business outcomes; integrating operational data with existing informational technology systems and applications; and gaining intelligence at the edge are all parts of the functionality of an IoT architecture as depicted below:

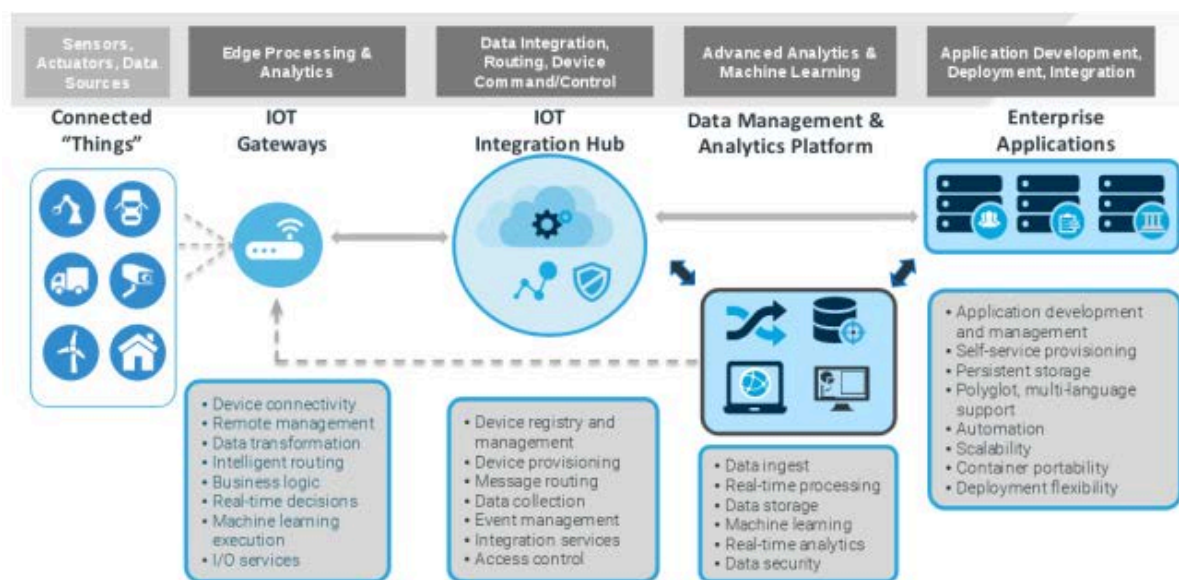


Figure 5 - E2E IoT Functional Architecture

Deploying IoT solutions in real production scenarios need more than a single provider to successfully accomplish the functionality in the above diagram.

Successful implementations have used a combination of technology providers, network providers, security specialists, developers, and system integrators, all operating as a cooperative, open, modular, and flexible team.

Many companies have addressed this complex End to End (E2E) architecture by launching IoT cloud platforms that typically focus on a centralized cloud IoT integration Hub. These platforms typically handle:

1. Device registration, management, security, and monitoring
2. Aggregation of data from multiple IoT devices
3. Integration with Big data storage and analytics
4. Sharing of data with other systems and applications
5. Handling and responding to device events

What is common among these IoT cloud platforms is that they offer the service on a low-cost, per-device subscription model, making it accessible even to hobbyists and the vendor community.

2.2. Edge compute

In the previous chapter on the technology cycles, we have described the evolution towards more decentralization as a mandatory step to make current IoT solutions more economically viable.

The underlying reasons for moving more processing to the Edge are better understood and accepted today.

The word edge in this context means literally its geographic distribution. Edge computing is computing that's done at or near the source of the data, instead of relying on the cloud at one of a dozen data centers to do all the work. It doesn't mean the cloud will disappear. It means the cloud is coming to you.

The main drivers to make current IoT more economically viable in this context are:

- Bandwidth preservation
- Latency improvements
- Security
- Privacy
- Autonomous operation
- Agile developer friendly IoT edge devices

2.2.1. Latency

One of the drivers for edge computing is the speed of light. If Computer A needs to have a response from Computer B, half a globe away, before it can do anything, the user of Computer A perceives this delay as latency. The brief moments after you click a link before your web browser starts to actually show anything is in large part due to the speed of light.

Moving compute closer to the origin of the data reduces the latency involved in the round trip to the cloud. Some of the evolving use-cases such as Augmented Reality (AR) and Internet of Things (IoT)

benefit from edge computing. End users of these applications enjoy immersive experiences delivered by the edge.

The goal of edge computing is to minimize the latency by bringing the cloud compute capabilities to the edge.

2.2.2. Bandwidth

IoT incorporates devices, sensors and data sources to send data from the field to the cloud to help businesses make real-time decisions. Edge computing uses connected and non-internet connected devices to process data near where it is collected. So, instead of sending all of the data to the cloud, edge computing can process data and send a smaller selected set, accelerating the delivery of data.

Edge requires less bandwidth by processing and analyzing data where it's gathered. A [McKinsey](#) article about the potential of IoT mentions an oilrig that has 30,000 sensors, but only one percent of the data are examined. Consider how much more efficiently a system could function if only the necessary data points were sent to the cloud, instead of all the information which needed analysis. In a network with millions of devices and data sources, this is an invaluable benefit. With the addition of edge computing, businesses can process data near the collection site and send selected data to the cloud when desired; reducing bandwidth and limiting security risks.

For data intensive applications, cost improvements can easily be calculated as expressed in the Edge compute financial model calculated in previous chapter.

2.2.3. Security

Securing these new edge devices is vital and should not be forgotten. You will need to enforce data encryption, both in transit and at rest, and protect communication with master clouds. Only by establishing security by design and embedding security mechanisms in all the components/layers involved will your edge workload be on the right track. IoT devices often collect sensitive information, and there's little agreement on how to protect this data. A study by HP found that 70 percent of IoT devices are vulnerable to attack.

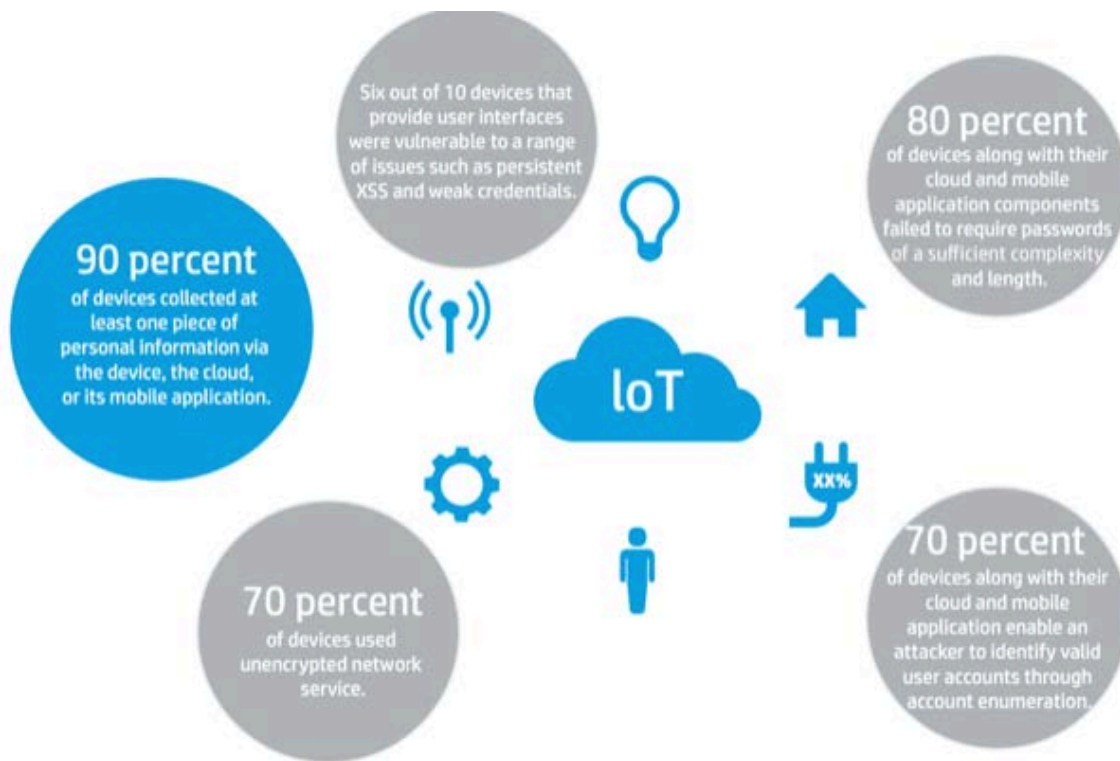


Figure 6 - Security vulnerabilities for IoT devices

2.2.4. Agile developer friendly IoT Edge devices

The benefit of the mobile app economy is that independent developers can create apps with a modest investment. And the ease of distribution and monetization provided by app stores creates a strong incentive for innovators. The IoT economy also needs independent developers, but the barriers to entry are higher due to the fragmented world of embedded developers and the complex nature of the development cycle, with its low level programming languages like C/C++, which require a steep learning curve before becoming proficient.

In these respect, the same kind of democratization which propelled the mobile and cloud world needs to happen on the embedded side as well, giving the developers a choice of selecting any modern programming language that makes them productive, and will open this traditional embedded world to a much larger community of developers.

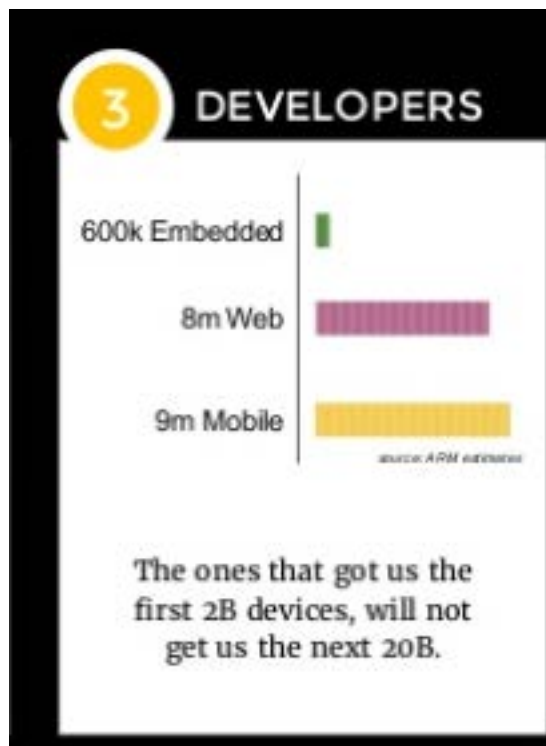


Figure 7 - Developer Communities

On top of this multi-programming language (polyglot) execution environment, the same frictionless experience needs to be given to the developer's with respect to the cloud's self-service model, friendly to use API's and easy to consume Software as a Service (SaaS) solutions.

2.2.5. Autonomous operation on the Edge devices

Edge compute can also maintain service operation during a network failure. On-device, service-specific processing that is enabled by edge compute can act as a buffer during a network failure and then synchronize data and state with cloud-based processes when Internet connectivity is restored. Some sensor applications have additional redundancy requirements that may include using a battery backup to operate during a power outage. Edge compute, coupled with a battery backup, creates a robust platform for offline data processing.

A cloud based IoT Hub or messaging service typically provides a very lightweight Agent that can be easily integrated in very constrained devices (20K – 50K SDRAM) and enables these devices to communicate with the IoT cloud platform directly.

When more functionality and processing is required a miniaturized version of the IoT core platform is offered that can act as an execution environment for local computation.

Vendors of this solution typically provide this as compatible solution with their cloud services so that from the sensor device point of view, there is no distinction between the local processing and the cloud processing.

The Edge IoT Gateway (GW) equipped with this IoT local environment then acts as a kind of bridge or proxy between the sensors and the cloud infrastructure.

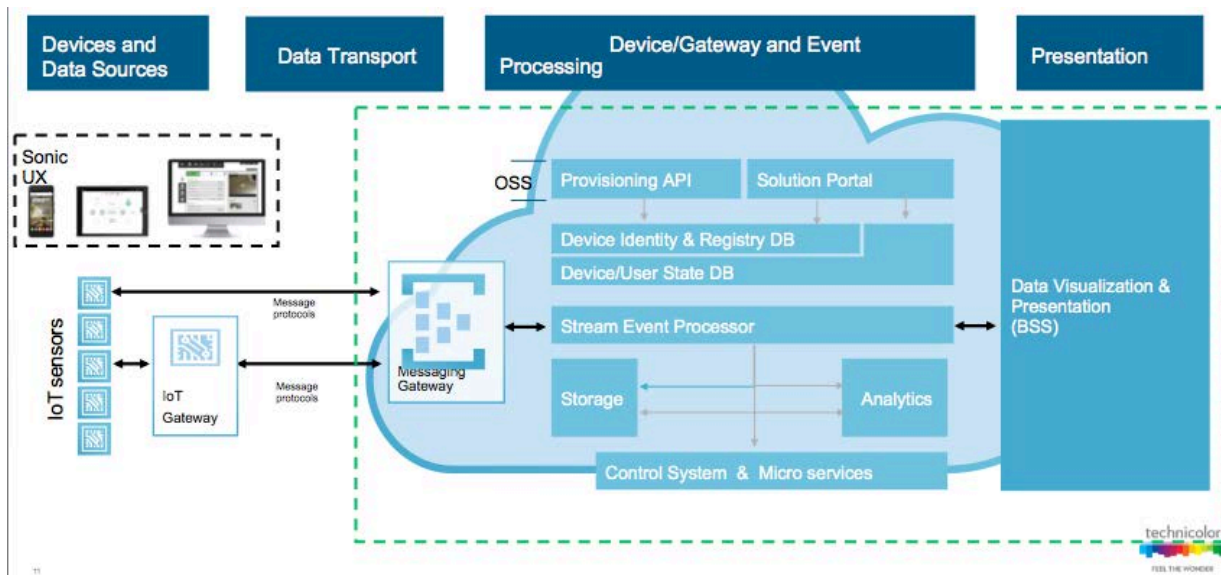


Figure 8 - IoT GW for autonomous operation

2.2.6. Privacy

Privacy has become a major concern for consumers in IoT markets like US and Europe, and is only one of the concerns reduced by successfully leveraging edge compute. Regionally, European markets have instituted greater regulatory protections over privacy than the US. This is best exemplified by the implementation of the General Data Protection Regulation (GDPR) which took effect May 2018. The GDPR significantly changes how companies handle EU citizen data privacy. GDPR places requirements for managing EU citizen data, such as a 72-hour notice to citizens after first having become aware of a data breach. Other aspects of the regulation require data erasure, data portability, and reporting. In general, personal data about identity, habits, speech and video will become a growing concern. Edge compute can enable new service architectures where personal data is processed locally to minimize the exposure of personal data.

Edge computing will undoubtedly provide the next level of efficiency for IoT solutions. For enterprises looking to build a powerful, scalable and secure IoT solution, it only makes sense to incorporate edge computing.

2.2.7. From cloud to Fog and Edge endpoints

There are actually two related concepts at play: edge computing and fog computing. Both models push data processing capabilities closer to where the data originates but differ in their emphasis. Crudely, fog computing locates the intelligence in the core operator's network while edge computing puts it inside the devices themselves.

Exactly how computing at the edge will finally be implemented will come out of a competition between various vendors, consortiums, and standards. The result will be a complex hybrid system that puts computing power and decision making at whatever location is optimal.

SOLUTION – EDGE / FOG COMPUTING

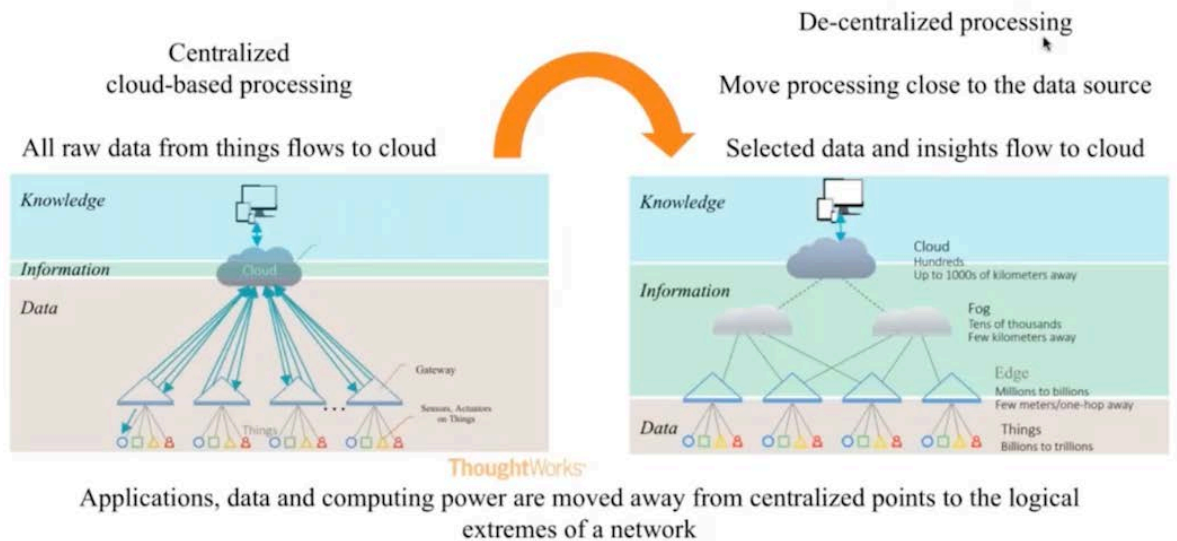


Figure 9 - Edge/Fog computing

2.3. Examples of IoT Edge Platforms

“Containers are a natural way to allow not only portability but also reusability so that you can assemble just the right application. Most modern Edge platforms are using containers under the hood, either by providing direct support for Docker containers, or by providing a serverless execution environment, compatible with their cloud offerings (AWS Lambda, Azure Functions) Edge apps or services are assembled from one or more containers and the containers perform inter-container communications using a local message broker in the Edge.” The compatibility of containers provides deployment flexibility and allows efficient upgrade scenarios due to the capability of doing incremental container updates.

Typically, the Edge architecture contains similar functional blocks as their cloud counterparts but are miniaturized to be able to run on embedded devices.

This section focuses on exploring three current solutions in the market, 2 commercial ones from Amazon (AWS Greengrass) and Microsoft (Azure IoT Edge), and one open source implementation called Resin OS, which offers a stripped-down version of the Docker engine leveraging the open source project Balena.

These Edge platforms are currently mainly addressing the Industrial IoT space. The minimal hardware requirements typically start from a quad core central processing Unit (CPU) similar to a Raspberry Pi and memory requirements of around 1G Byte of SDRAM, and a couple of GBs for storage to be able to run a couple of containerized applications, and base images of a client Operating System (OS).

Both AWS and Azure solutions are tightly integrated with their cloud IoT core functionality, so that they can be centrally controlled for security reasons and administrated with respect to the Life cycle management of their containers and runtime engines. They naturally extend their cloud IoT platform towards the Edge.

Although the code is written in Lambda, Azure Functions or containers are technically capable of integrating with other cloud infrastructure or SaaS solutions. However, the risk of vendor lock-in is high due to the frictionless, tight integration of their own services.

Besides a container and/or server-less runtime execution environment both AWS and Azure offer local implementations of IoT messaging (IoT Hub) components which daisy chain to their cloud counterpart as a kind of federated broker infrastructure.

They also both provide a local copy of the device shadows/twins registry that can be optionally configured to synchronize their device/application state with the cloud IoT registry so that in case of connectivity problems cloud applications call still query the cloud registry, and the device can continue to operate on the local state of the device/application.

Both solutions are relatively young. AWS released their first version at Reinvent 2016 and Microsoft one year later, but they are actively adding new components like machine learning frameworks, and storage solutions to their offerings.

Amazon only offers a server-less compute environment, while Microsoft offers both a server and server-less environment as a Docker compatible container framework.

The advantage of the latter is that you can use existing software packed in Docker images, while the server-less environment needs newly designed software components in line with the server-less methodology.

None of them provides multi tenancy support to allow different vendors to control their own container infrastructure.

2.3.1. Azure IoT Edge

The Azure IoT Edge consists of:

- Docker compatible container runtime
- Azure Functions
- Message broker
- Offline/synchronized device twins
- Cloud management and deployment
- Cloud development/test support
- Azure stream analytics
- Azure machine learning

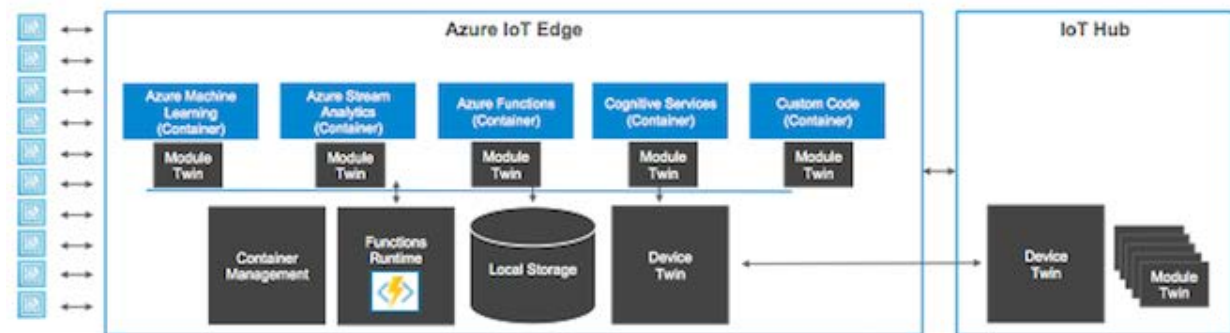


Figure 10 - Azure IoT Edge

2.3.2. AWS Greengrass

The AWS Greengrass Core software consists of:

- A message manager that routes messages between devices, Lambda functions, and AWS IoT.
- A Lambda runtime that runs user-defined Lambda functions.
- An implementation of the Device Shadow service that provides a local copy of shadows, which represent your devices. Shadows can be configured to sync with the cloud.
- A deployment agent that is notified of new or updated AWS Greengrass group configuration. When new or updated configuration is detected, the deployment agent downloads the configuration data and restarts the AWS Greengrass core.
- Analytics Engine
- Discovery
- Free RTOS Greengrass compatible version without container support for low footprint devices.



Figure 11 - AWS Greengrass

2.3.3. Resin OS

Resin OS shares a lot with cloud operating systems that support containers. They share the focus on minimalism, getting out of the user's way and letting their container do the heavy lifting, and using Docker, which is the standard way of running containers, and well understood by a large developer community. Resin OS applies the same principles to a different domain, that of embedded Linux devices, sometimes called "connected devices", "Internet of Things" or "Industrial Internet", depending on the use case.

At the core, it relies on a more optimized version of the Docker runtime engine, called Balena, which is an open source project, and is based on the Docker's Moby open source engine, but incorporates a number of optimizations to make the engine more usable for constraint embedded devices.

Some of these changes include:

- 3–4 times smaller footprint by using one single binary and de-duplicating shared GoLang libraries.
- Multiple CPU architectures support, for a heterogeneous set of different embedded platforms.
- Bandwidth efficient delta updates with binary diffs, which result in dramatic size optimization for updating the container images.
- More robust support for flash updates during devices failure cases, like reboots or power failures, and minimized disk writes by on-the-fly extraction from images on the registry.
- Removing non-relevant features not applicable for IoT embedded devices like SWARM support, plugin support and removing cloud logging drivers.




Figure 12 - Resin OS

2.4. Life Cycle Management of IoT Services, Development And Deployment

The Internet of Things means smart devices connecting every aspect of our lives, which creates a lot of challenges. And there's more to IoT than hardware and connectivity. To run and connect all these machines, we need software, lots of software. And we must be able to develop and update it quickly while making sure it's performing securely and efficiently.

So here's the problem: how do you deploy code to millions of devices in people's homes, offices, cars and beyond? How do you monitor these devices? What are the security implications? How do we design, develop and test all this stuff, and how do we configure all these devices? Fortunately, the tech industry has been addressing these challenges, albeit at a comparatively small scale, for over a decade.

2.4.1. DevOps Practices to the Rescue

The growth of DevOps is crucial here because it has led organizations to create the kind of automated systems that will be required for IoT software. Continuous delivery and continuous deployment at scale are only feasible through automation. Scripting or—better still—containerizing these processes with appropriate automated tests is the answer.

Mixing agile practices into DevOps accelerates things even further. Incremental builds with frequent releases of small batches of code makes for faster and safer development. Close customer contact to solicit feedback and usage analysis coupled with continuous improvement helps agile organizations respond quickly to emerging requirements.

2.4.2. Securing IoT Devices the Agile Way

We have barely touched upon the most significant of the concerns that organizations have about IoT: security. With all those connected devices out there in the wild, the attack surface is almost inconceivably huge. More than ever, the common development error of treating security as an afterthought just won't cut it.

Luckily, DevOps and agile practitioners don't think this way. Security is not silos—it's part of the development cycle, in the same way that product, quality and performance are. With this most significant of concerns addressed, it should be clear that organizations that use modern development practices are well set up to cope with the challenges of IoT.

The end result will be solutions that balance their functionality between the edge and cloud and creating a flexible approach to shift and lift these components back and forth when required.

Fortunately, most existing cloud IoT platforms include services around device management and SW Life cycle management as part of their core offering, allowing these solutions to easily scale to millions of devices. These foundations can be used to build modern, frictionless workflows and integrate them into existing developer's SaaS solutions, giving the developers the tools, they are used to and love like GitHub, a Container registry etc. After all, these ecosystems already play a crucial role in mobile and cloud development.

"It's just Git push and forget about it. It's that easy." – Sam Levy, Pact Coffee

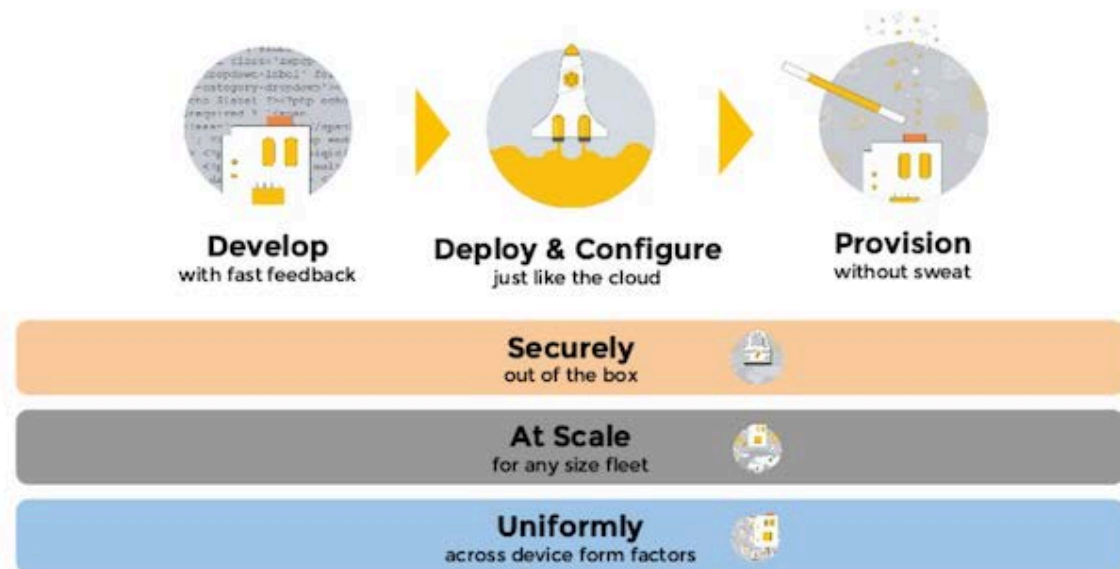


Figure 13 - Developer friendly workflows

Today, startups are already working on solutions that deliver this seamless, frictionless experience for developers working with IoT devices in the field.

They use containers as a means to develop and deploy their functionality and provide a workflow and tools, which are familiar for developers, and which were born in the cloud.

This brings IoT services development almost to the same level as their mobile counterparts, which have already established this mature practice for years with ecosystems that provide similar solutions and tools.

Today these solutions still have modest scale in terms of the amounts of devices they can handle, but thanks to the advances made on the basic core IoT foundation building blocks this type of approach will shortly come to the realm of the operators' large fleet deployments.

3. Lightweight Virtualization – Containers

At the turn of the century when cloud computing became popular, public cloud infrastructure providers used Virtual Machines (a software abstraction of a physical computer/server) as a virtualization technique to create abstraction of the physical hardware, create large aggregated pools of logical resources consisting of CPUs, memory, disks, file storage, applications, networking, and offer those resources to users or customers in the form of agile, scalable, consolidated virtual machines. Virtual Machine, while providing complete isolation of resources from one user to another, did come at a price. The price paid was in terms of overhead related to CPU virtualization, virtualizing the disk (block) and network I/O. For example, a typical state of art physical server could only host a few dozen VMs at most.

While a VM-based virtualization is a mandatory requirement in some business models, many business models exist where a lightweight virtualization mechanism might suffice, as long as it provided adequate application isolation.

In the past few years Containers have emerged as such a key application isolation, packaging and delivery technology, combining lightweight application isolation with the flexibility of image-based deployment methods. Containers use core technologies such as Control Groups (Cgroups) for Resource Management, Namespaces for Process Isolation, thus enabling secure multi-tenancy for applications without the overhead experienced by Virtual Machines.

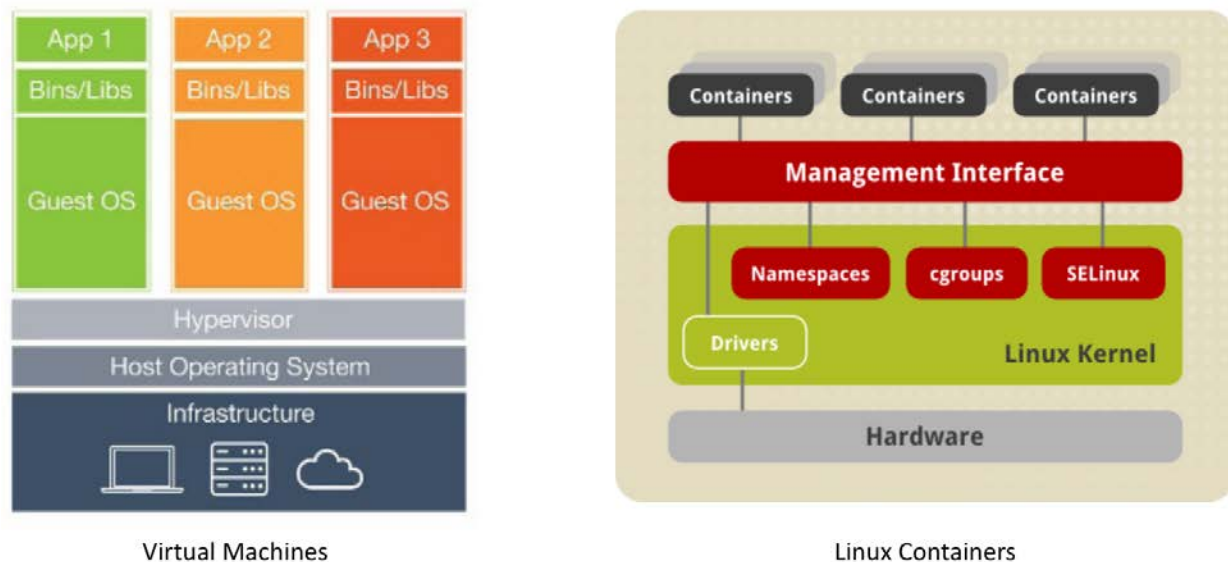


Figure 14 - Virtual Machines Versus Containers

Containers have since further evolved into various kinds that have varying degrees of abstraction based on application runtime requirements.

For Edge Compute Nodes, Virtual Machines are generally not viable due to the limited compute capabilities in the Edge device, hence a lot of virtualization for Edge applications is focused on Container based technologies.

3.1. Container Technologies

Compared to hypervisors, container-based virtualization provides different abstraction levels regarding virtualization and isolation. Hypervisors virtualize hardware and device drivers, generating a higher overhead. In contrast, containers avoid such overhead by implementing process isolation at the operating system level. A single container instance combines the application with all its dependencies, and runs as an isolated process in *user space* on the host operating system (OS), while the OS kernel is shared among all the containers as shown in the figure below. The lack of needing hardware/driver virtualization, together with the *shared kernel* feature, provide the ability to achieve a higher virtualized instance density because the resulting disk images are smaller.




Figure 15 Container Building Blocks

In the past few years several types of containers have come to prominence, each type most optimized for a particular application class and business logic. Containers themselves are not a new technology. Solaris platform has offered the concept of Solaris Zones for many years, and many Linux administrators have experimented with FreeBSD jails as a lightweight alternative to virtual machines. In order to understand the current container landscape, this chapter attempts to categorize the most fundamental approaches and technologies used.

At the core, Docker, LXC and other container technologies depend on the key Linux kernel features of Cgroups and Namespaces. Depending on use cases or requirements these container engines can make different trade-offs with respect to architectural choices and implementation. In order to understand these choices and validate the impact they have on performance, security, portability and footprint in the context of constrained device types, the categorization is depicted in the figure below:

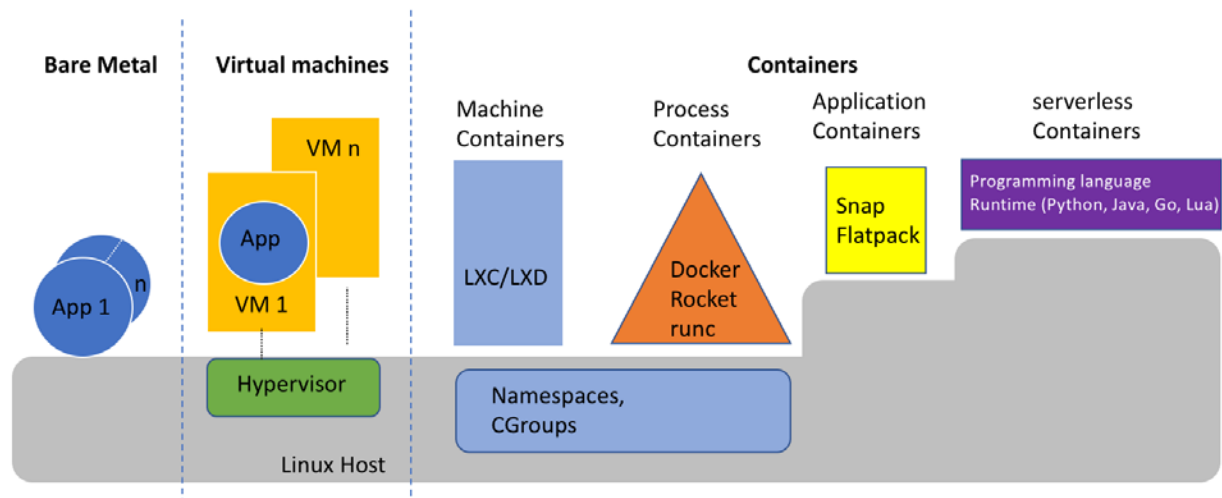


Figure 16 - Container Categories

3.1.1. Machine/System Containers (LXC/LXD)

Linux Containers (LXC) main focus is system containers. That is, containers which offer an environment as close as possible to the one you'd get from a VM but without the overhead that comes with running a separate kernel and simulating all the hardware. This is achieved through a combination of kernel security features such as Namespaces, mandatory access control and control groups. Just like Solaris Zones and BSD Jails, it tries to provide a lightweight VM experience for system administrators.



Figure 17 - LXC Containers

LXC storage management is rather simple. It supports a variety of storage back ends like btrfs, lvm, overlayfs, and zfs. However, by default (if no storage backend is defined), LXC simply stores the root file system under `/var/lib/lxc/[container-name]/rootfs`. For databases and other data-heavy applications, you can load data on the Rootfs directly or mount separate external shared storage volumes for both the data and Rootfs. Creating an image out of an LXC container just requires tar'ing up the Rootfs directory.

LXC is focused on IT Operators with the goal of providing a lightweight virtualization solution. This means, for a system administrator to transition from hypervisor-based virtualization to LXC is rather painless. Everything from building container templates, to deploying containers, to configuring the OS, networking, mounting storage, deploying applications, etc. all remain the same. In fact, LXC gives you direct SSH access. This means all the scripts and automation workflows written for VMs and physical servers, apply to LXC containers too. LXC also supports a template notion, which essentially is a shell script that installs required packages and creates required configuration files.

One can create either privileged or unprivileged LXC containers. The default today is unprivileged which allows for configurable access to the Linux kernel system functions.

LXD is an extension of LXC. While LXC is used under the hood, and provides a user space CLI interface, LXD is a system daemon that extends the LXC library with REST API's to provide container

management as a service and is written in GoLang. It simplifies the management of multiple LXC containers, takes advantage of host-level security measures (which in the case of LXC is more problematic) and simplifies resource sharing like networking and data storage.

In contrast to Docker it is only available for Linux platforms, which is not a fundamental problem for our type of Linux based platforms, but limits the development experience a bit.

3.1.2. Process Containers

While the term *process containers* was initially coined by Google in 2006 and led to the creation of Cgroups, this type of containerization focuses on limiting, accounting and isolating resource usage for a limited collection of processes. The new generation of container runtimes like Docker and Rocket (RKT) take this concept further by restricting their containers to only carry one service. A typical consequence of that is that even secure shell (ssh) access is not supported inside these containers, however this delivers a better separation of concerns.

This is very different from the traditional system containers like the ones created by LXC where you can have multiple services running on the same OS.

3.1.2.1. Docker

Containers became truly mainstream with the Docker project. Like most great technical innovations, Docker stood on the shoulders of older giants but with a new focus that addressed a market need at exactly the right time. In Docker's case it was a focus on developer experience and ease of (re)use that took containers from a fairly arcane infrastructure technology into something truly transformational.

While Docker was originally designed on top of LXC, it created its own version of that with the introduction of Runc. Runc resembled the basic functionality as LXC, with the addition that it is standards compliant with an OCI format. Docker is a first-class citizen in the design of Runc.

Figure 18 depicts the current building blocks of Docker.

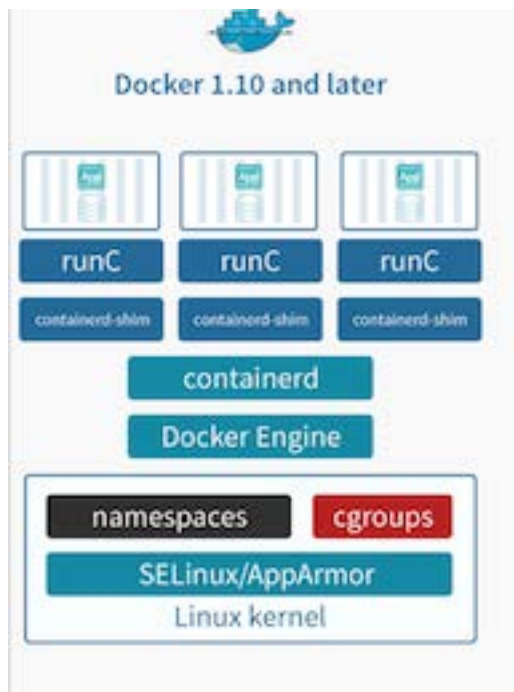


Figure 18 - Docker Building Blocks

The biggest innovation that Docker added compared to the previous mentioned system containers is that it added the container management images and made them stackable to achieve portability and flexibility with respect to reuse of these images for the developer.

Each Docker image references a list of read-only layers that represent file system differences. Layers are stacked on top of each other to form a base for a container's root file system. The Docker storage driver is responsible for stacking these layers and providing a single unified view.




Figure 19 - Docker packaging and image format

The major difference between a container and an image is the top writable layer. All writes to the container that add new or modify existing data are stored in this writable layer. When the container is deleted the writable layer is also deleted. The underlying image remains unchanged. Because each container has its own thin writable container layer, and all changes are stored in this container layer, this means that multiple containers can share access to the same underlying image and yet have their own data state.

Layering helps Docker to reduce duplication and increases re-use. This is very helpful when one wants to create different containers for various components. One can start with a base image that is common for all the components and then just add layers that are specific to your component. Layering also helps when one wants to rollback changes as you can simply switch to the old layers, and there is almost no overhead involved in doing so.

In order to achieve this a lightweight Union File System is used, that allows sharing files that are not changed, reducing the storage and memory footprint and enhancing the startup time.




Figure 20 - Docker Containers and images

Copy-on-write is a similar strategy of sharing and copying, in which the system processes that need access to the same data share the same instance of that data rather than having their own copy. At some point, if any one process wants to modify or write to the data, only then does the operating system make a copy of the data for that process to use. Docker makes use of copy-on-write technology with both images and containers.

Currently the Docker engine is stripped out in different components, where the current Docker engine focuses on image management, container orchestration and the interaction with the Docker container repository, and Containerd is managing the complete container lifecycle of its host system: image transfer and storage, container execution and supervision, low-level storage and network attachments, etc, similar to what LXD does for LXC.

Docker is more than an image format and a daemon, though. The complete Docker architecture comprises the following components:

- Docker daemon: runs on a host and is the combination of Containerd and the Docker engine.
- Client: connects to the daemon, and is the primary user interface
- Images: read-only template used to create containers
- Containers: runnable instance of a Docker image
- Registry: private or public registry of Docker images

- Services: a scheduling service called Swarm which enables multi-host, multi-container deployment. Swarm was introduced in version 1.12

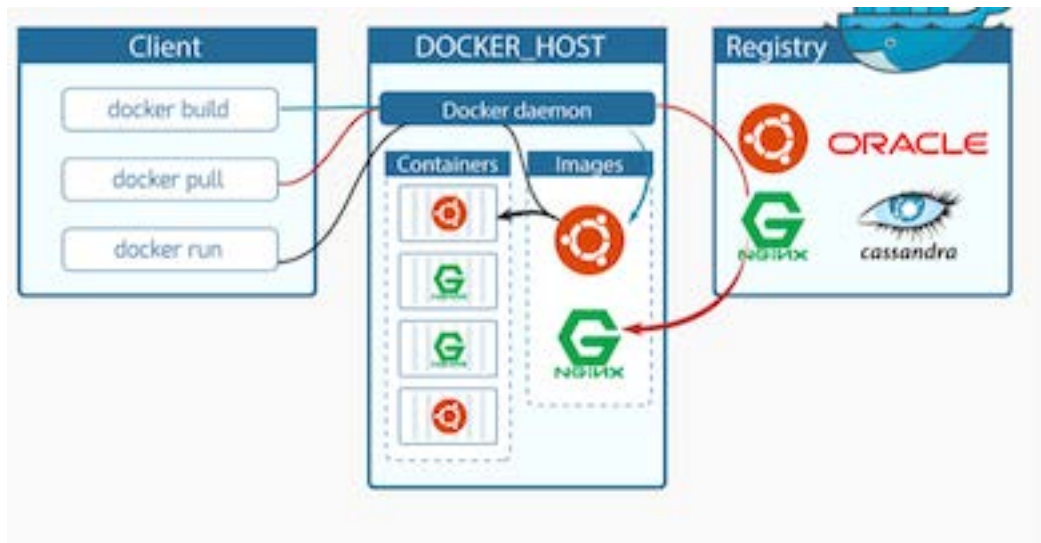


Figure 21 - Docker ecosystem

3.1.2.2. Microservices

The application space where technologies like Docker and RKT are a perfect fit for, can be roughly categorized as – modern, Microservices-based, and traditional enterprise applications.

MSOs can adopt a Microservices architecture which has gained popularity amongst new web-scale companies like Netflix, Google, Twitter, etc. Applications with a Microservice architecture consist of a set of narrowly focused, independently deployable services. This would give the MSOs the advantage of increased agility and resilience to roll out new value added services to their subscribers. Agility since individual services can be updated and redeployed in isolation. Given the distributed nature of Microservices, they can be deployed across different platforms and infrastructures, and the developers are forced to think about resilience from the ground up instead of as an afterthought.

To support ephemeral and elasticity workloads a lot of focus and functionality has been added to make these solutions a perfect match for container orchestration frameworks like Kubernetes, DCOS, SWARM and the like, but as a side effect of some of their choices, their runtime footprint is heavy, and the impact on performance by adding a specific storage backend and copy-on-write functionality can be a bottleneck for IO intensive applications.

3.1.3. Application/Desktop containers

Machine containers like Docker or RKT are almost as versatile and adept to orchestrate a complex system of Microservices; however machine containers are not ideal for a significant number of Edge-Compute use cases like IoT where the physical H/W is an embedded platform with very limited compute, memory and storage like a Cable Modem/Gateway or a Wi-Fi Router. Such platforms typically have dual core ARM processors (Quad core at best) and have just enough resources to run lightweight processes in an isolated sandbox.

An isolated sandbox is a mechanism to separate running programs with the aspect of keeping the operating system and other applications isolated and safe when running programs. The sandbox does so by isolating or virtualizing an environment where the program can be executed in, limiting the possible access to the underlying operating system and other applications. Monitoring in a sandbox can occur as well as trying to check whether or not the executing program is malicious.

These techniques have been heavily used in the past by web browsers and other graphical applications, to allow running untrusted code in these environments, without compromising the security and robustness of the overall system.

Over the recent years, we have seen these techniques being reused in the Linux Desktop Application space to overcome the fragmentation that exists with respect to the numerous flavors of Linux distributions out there. The idea here is to have a minimum base OS and run everything in distribution-agnostic containers.

Due to the nature of this use case, which does not require running hundred of containers, there is also less need to optimize for size and reuse of different image layers and using techniques as copy-on-write. It allows keeping the storage model simple, and avoiding the negative impact on performance that these union file systems tend to have.

Most notable players in this field are AppImage, Ubuntu Snap and Flatpack.

While originally developed for Desktop applications, Snap for example is really trying to broaden their applicability to other domains like IoT edge computing, and recently also added OpenWRT support. The disadvantage of the current Snap approach is that it is very vendor specific, and the current base image only supports a heavy 250MB Ubuntu distribution.

On the flip side, these tools allow secure and controlled access to a shared system bus like Databus (DBUS) or eventually Microbus (UBUS) for OpenWRT, which allows interaction between different containers in an efficient way.

3.1.4. Serverless containers

Serverless computing is a cloud-computing execution model in which the cloud provider dynamically manages the allocation of machine resources. The name "serverless computing" is used because the server management and capacity planning decisions are completely hidden from the developer or operator.

AWS Lambda introduced by Amazon in 2014,^[5] was the first public cloud vendor with an abstract serverless computing offering, soon followed by Azure and Google, with their variant Cloud Function. Today, both Amazon and Azure also offer these runtimes for their Edge computing offering, notably AWS Greengrass and Azure IoT Edge.

Although it is out of the scope of this paper to discuss the specific programming model for these new types of solutions, they are using the same kind of containerization and isolation techniques as their former counterparts. The way they achieve isolation is by using a programming language runtime that acts as a barrier between the host and untrusted code inside the containers.

As such they provide a higher level of abstraction than just the operating system by making the language runtime accessible from within the containers, and leaving package management up to specific language

specific solutions or simple tar or zip approaches. The resulting implementation as such tends to be less complicated in terms of functionality and footprint.

As of today these solutions offer runtimes for different languages like Python, Java, C# and Go.

Serverless is sometimes mistakenly considered as more secure than traditional architectures. While this is true to some extent because OS vulnerabilities are taken care of by the underlying environment, the total attack surface is significantly larger. Serverless approach is not mutually exclusive of the other types of containerization techniques but can be used to complement each other.

Today containers technologies are supporting the major range of Linux distributions, even Android, and serverless computing frameworks can run on top of Dockerized environments or application container frameworks like Snap. As an example Greengrass today is available both for Docker and for Snap, and Azure IoT is built on the foundation of Docker infrastructure.

3.2. Container standarization

The **Open Container Initiative (OCI)** is a Linux Foundation project to design open standards for operating-system-level virtualization, most importantly Linux Containers. There are currently two specifications in development and in use: Runtime Specification (runtime-spec) and the Image Specification (image-spec), and the consortium recently also released a Distribution Specification for Registries heavily based on the Docker protocol.

The OCI was launched on June 22nd 2015 by Docker, CoreOS and other leaders in the container industry, and is today supported by the major public cloud players Amazon, Microsoft, Google, IBM, chip manufacturers like Intel, and some notable Telco players like AT&T and Huawei.

Their goal is to address the fragmentation in the container landscape by defining clear specifications to align on interoperability and compatibility and avoiding vendor lock-in to accelerate container adoption.

Besides specifications they also released an open source initiative in collaboration with Docker to provide a basic runtime called Runc to drive innovation and allow building higher-level tools on top of it. It is used today under the hood for all Docker types of frameworks and on GreenGrass and Azure Edge.

Also all major cloud providers are supporting the OCI image specifications currently in their container registries, and will soon also support the upcoming API, which is defined in the Distribution specification for registries. The distribution-spec is the latest in a series of OCI initiatives highlighting the rapid rate at which container technologies are maturing. In fact, rather than focusing mainly on building pipelines to construct single applications, more attention is being focused on how to manage supply chains that ultimately will consist of multiple registries and pipelines. Constructing supply chains at scale assumes a base level of interoperability between registries enabled by a Docker Registry v2 protocol, which can now be deployed anywhere to access images in public or private clouds.

App Container (appc) is an open specification that defines several aspects of how to run applications in containers: an image format, runtime environment, and discovery protocol.

The image format defined by appc and used in RKT is the Application Container Image or ACI.

As part of the success of OCI, alternative technologies like LXC and RKT are starting to offer more support and integration of OCI specifications, by adapting their roadmap, however this evolution is still ongoing.

4. MSO CPE landscape & CPE device characteristics

For the last 20 years or so fixed broadband connectivity has become a de-facto choice of Internet connectivity for most users. MSO-led cable broadband has led the way of providing broadband Internet connectivity to the users. During the early years of broadband era, the desktop/PC used to be the primary end user device at home. MSOs used to provide a DOCSIS (Data Over Cable Service Interface Specification) modem that served as the NTU (Network Termination Unit) at the home premise. PCs in the home typically connected to the modem via Ethernet cables. Once MSOs started offering home voice/telephony service, MSOs combined the modem and voice telephony gateway into one integrated broadband device called an eMTA. eMTAs are typically a purpose-built network device with very little compute capabilities outside of their core function of providing voice and internet termination function.

In the last 10 years, however, the connectivity landscape inside a typical residential home or enterprise has significantly changed compared to early days of broadband described above. Today, wireless (Wi-Fi) is the primary way users access Internet in their homes. Also, the number of end user devices have proliferated from a few desktops to Wi-Fi capable smart phones, Tablets, laptops etc. Further, users are also bringing home internet connect smart devices like smart TVs, Coffee Pots, garage door sensors, door locks, home assistant / speakers etc. These smart things are typically Wi-Fi, Bluetooth or Zigbee connected and an average US household has many as 13 such connected devices at home ^[2] and the number is expected to go up to 50 by 2025. Proliferation of IOT end user devices and extensive Wi-Fi usage in the home, gives MSOs a chance to move upstream and provide a new category of value added services to their subscribers. MSOs are thus integrating many wireless connectivity technologies in the current generation DOCSIS gateways e.g. Wi-Fi, Bluetooth, Zigbee/ZWave, LoRA etc. to provide a wide spectrum of wireless connectivity options in the home. Further Moore's law has been bringing the cost of computing down to a fraction of what it used to be 20 years ago. Multi-core, general-purpose computing costs for embedded systems of the likes of Broadband Gateways are now common. (see [Raspberry Pi project](#)).

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 GPU (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. All such integrated Gateways are capable of performing computing at the edge / customer premise.

Edge Compute capable broadband gateways provide many tangible benefits to the MSOs. Edge Compute gateways are actually platforms that let the MSO disaggregate the applications from underlying gateway firmware. In the traditional workflow, the MSO have to do intense planning for 5-7 years to do a CPE refresh. Since the apps and the firmware were tightly integrated, once the CPEs were deployed, there was little to no new feature development on these gateways. This made the service rollout cycles for MSOs very slow and risk of failure very high, as MSOs usually had to plan 5-7 years in advance. On the other hand, Edge-Compute Gateways are capable of disaggregating the 'Gateway OS' from the network applications (e.g. IoT apps etc). Just like in Desktop computing where a desktop OS is disaggregated from the desktop applications, Edge-Compute capable broadband Gateways let MSOs deploy the Gateways and the applications independently. Typically, Gateway OS will evolve at a slow cadence while applications will be rolled out in a much more agile fashion. This helps the MSO operate at the Web scale, giving them much needed agility and service velocity, which is needed to compete in the marketplace.

5. The New Embedded Stack

The telecom industry is currently going through a radical shift towards softwarization of the current hardware base networking functionality. Originally, they had started mapping their networking functions on a VM based infrastructure. In the mobile world with the introduction of 5G, they more and more started to apply a cloud native approach towards the implementation of network function virtualization (NFV) functions, with more prominent usage of container technology and orchestration. ETSI has been standardizing the Multiaccess Edge Computing environment also called MEC (Mobile Edge Computing). AT&T recently merged their efforts around the Open Network Automation Platform (ONAP) with the China Mobile Open Initiative, by contributing their code to the Linux Foundation as an open source initiative, which broadens the application domain towards wireline and enterprise.

From an operator perspective, the Edge means the base station and their core network infrastructure, but extending their reach towards the physical edge. Applying these lightweight virtualization techniques on their customer premises equipment 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.

Operators' close proximity to their users due to their large installed base of CPE's could make them more competitive by exploiting edge computing on these types of devices.

Specifically, in the IoT applications domain, the physical edge will play a dominant role in creating value-add services, especially for those novel use cases that are at the intersection of different IoT domains and different verticals, where latency and autonomous operations are crucial. Edge compute will be a key enabler that will foster and promote new compelling use cases. Figure 22 shows a subset of these use cases.




Figure 22 - Edge compute use-case domains

By extending the edge, CPE devices can play a crucial role in bridging the current fragmented and heterogeneous landscape of IoT cloud platforms, that are using different messaging/communication protocols and API's, with their own vendor specific semantics, the limitless amount of low power

connectivity protocols and standards that are currently dominating most of the smart home use cases of today, and allowing third parties to fully exploit the hardware capabilities present in those devices.

To fully exploit this evolution, the key elements that need to be addressed are creating the proper abstractions, the programmability of these devices and the interoperability between various ecosystem solutions. Lightweight virtualization technologies like containerization appear to be a perfect fit to address these challenges, since they provide a good computational abstracted environment and allow independence for various technology choices to be used inside this isolated environment, as well as independence and full end-to-end ownership with respect to the integration of these components with their own cloud infrastructure and services.

Given the fragmented IoT landscape and importance of cross domain use cases that benefit from a CPE edge solution, this edge platform should support a multi-vendor environment, and hence this platform needs multi-tenancy support as a first class citizen to host different solutions, that is controlled and operationally managed by the operator, so that issues around security and privacy can be fully covered.

To improve integration and create more flexibility in the choice of where to place these containerized application or service components, either on the CPE device itself or close by in the network, the CPE edge solution should have a common execution environment that enables cross platform deployment. We see this evolution today in commercial approaches from the public cloud providers like Amazon and Microsoft that make their edge environments compatible with their cloud counterparts so that containers or serverless functions can run practically anywhere.

Taking into account all previously mentioned arguments, we recommend taking an evolutionary approach by segregating the current monolithic CPE firmware into a dual track approach, where a clean separation of concern is obtained between traditional development and deployment strategies of the firmware versus a more open and dynamic environment for containerized components. As such the years of investment of this network-centric firmware and integration with the operators' OSS/BSS (Operations and business Support Systems) and remote management infrastructure can be safeguarded and can be used to upgrade both the firmware as the containerized components.

As an example of that, the current TR69 standard has already been enhanced with support for LCM of software components, more notably by the TR-157 data model, which can support different execution environments.

In addition to being backwards compatible with an MSO's infrastructure, we propose an additional middleware layer on top of the firmware that delivers a modern developers focused life cycle management layer for containerized components, albeit still under the control of the operator, but that gives full control of the development and deployment of these components in an end-to-end fashion to the IoT application solution provider.

By exposing an LCM API from the firmware this middleware layer would be a combination of a containerized agent of the CPE device and a backend infrastructure focused on deployment and management of containerized services. Typically, this backend solution will leverage a cloud based IoT platform that contains the service provisioning methods and monitoring tools that are independent of the managed applications that run in parallel on the CPE edge platform.

As such, this means that this middle layer needs to be in line with the IoT strategy of a particular operator, since an IoT cloud platform typically supports all the foundational building blocks to be able to communicate with the devices, the service/asset management services, device on boarding, authentication and integration with data monitoring and analytics infrastructure., and is needed to orchestrate the

deployment of containers, and to be able to do the operational management of the running containerized services.

By using a cloud-based control plane that interfaces with the existing CPE middleware functionality, and integrated with their existing back office infrastructure, operators can offer a multi service CPE platform that can deliver new, compelling use cases in line with their goals on enterprise-grade service assurance.

As previous stated, there is a heterogeneous set of IoT cloud solutions today that is probably not going away soon, and operators have or are in the process of making different strategic choices around technology and partners. Some of them are well versed and mature to build these IoT-based services in house, even in a hybrid fashion, while others will rely on commercial solutions and partners or have acquired an IoT platform.

The key here is to support easy integration on the CPE which will allow different operators to integrate their IoT backend with a set of standard local LCM API's and a virtualized IoT agent to deliver this type of orchestration around the LCM of containerized components in an end to end fashion.

This would allow operators to create an ecosystem around service components, much like the app stores of today, by exposing a set of high level services API from this middle layer for higher layer integration of additional services, similar in concept to the Northbound API of their current remote management infrastructure, but more open to third party solution providers.

Security and privacy controls, will be absolutely key, and included by design in this edge layer. As such there is a need for a secure execution environment, assisted by hardware mechanisms like Trusted Platform Modules (TPM) and a root of trust for safely storing secrets as supported by most modern CPU's like ARM Trust zone architecture. This will allow only trusted authenticated container code to run on these CPE devices, preventing hackers from compromising the systems, or running unintended malicious code. Set top box (STB) content providers today already mandate this to Digital Rights Management (DRM) violations.

This same proper isolation is needed for privacy related issues, hence the necessary mechanisms on access controls need to be in place to protect unauthorized access, and limit exposure of users' personal information.

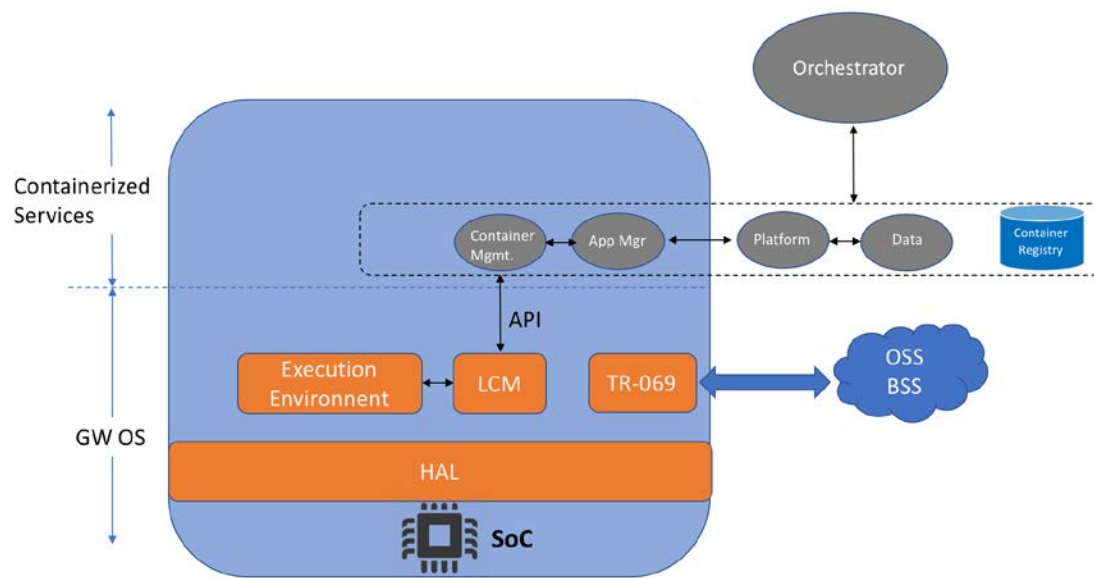


Figure 23 - Modern LCM layer for CPE

Conclusion

The decentralization of computing towards the edge, and the expansion of this edge compute paradigm towards the physical world of operator's customer premises equipment (CPE) is removing the last barriers of IoT application development and deployment to gain end to end control of their solutions in a more economical and viable way.

For the IoT project life cycle, implementing continuous delivery and being able to constantly add intelligence to these edge devices is absolutely mandatory to ensure that these solutions deliver on the promise of rapid value creation.

The intersection of DevOps practices and IoT Edge compute is still embryonic but growing day by day, to become an essential part of modern service delivery.

Furthermore, this Edge compute paradigm will become essential to unlocking the value of many IoT use cases, certainly for these use cases that are at the intersection of different application domains and industries.

The multitude and maturity of the current IoT platforms, and the advances in Big Data technology, will serve as the foundation, to build these new DevOps and life cycle management tooling and practices on, at an unprecedented scale. In this respect, it is important for the operator to include a strategy around modern life cycle management and containerization in their IoT initiatives.

When using lightweight virtualization techniques like containers on such CPE device start to become viable, they will provide a uniform, portable and secure environment for edge computing software components, solving the heterogeneous nature of today's embedded devices, as well as deliver a means to package all necessary dependencies in a uniform way in a more dynamic and agile way during the deployment phase.

Thanks to the combined industry effort of the last years in this area, the container technology has matured, both in terms of standardization effort, as from the technology perspective, by making open source container frameworks available for inclusion on these type of constrained devices, while still leaving room for innovation and improvements in various areas like security, operational management and privilege rights.

Albeit numerous researches have already been conducted on performance impact and footprint, further integration, evaluation and detailed measurement is required in this area on the applicability of operators' CPE equipment and their ecosystems.

While there are still numerous challenges around the immense scale of CPE deployment, security, privacy, standardization and interoperability issues, solving the infrastructural problems around automatization and the life cycle management of these new containerized edge computing components will be an absolute necessary first step.

Operators should expand their focus on edge and fog computing efforts to include CPE edge computing as a first-class citizen in their ecosystem.

Abbreviations

ACI	Application Container Image
AR	Augmented Reality
APPC	Application Container specification
Cgroups	Linux Control Groups
CRM	Customer Relationship Management
CPE	Customer Premises Equipment
CPU	Central Processing Unit
DBUS	DataBUS
DOCSIS	Data Over Cable Service Interface Specification
DRM	Digital Rights Management
eMTA	Embedded Multimedia Terminal Adaptor
ERP	Enterprise Resource Planning
ETSI	European Telecom Standards Institute
E2E	End to End
GDPR	General Data Protection Regulation
GW	Gateway
GPU	Graphics Processing Unit
IaaS	Infrastructure as a Service
IoT	Internet of Things
ISP	Internet Service Provider
IT	Information Technology
LCM	Life Cycle Management
LXC	Linux Containers
MEC	Mobile Edge Computing
ML	Machine Learning
MNO	Mobile Network Operator
NFV	Network Function Virtualization

ONAP	Open Network Automation Platform
OCI	Open Container Initiative
OT	Operational Technology
OS	Operating System
OSS/BSS	Operations and Business Support System
RKT	Rocket
SaaS	Software as a Service
SCTE	Society of Cable Telecommunications Engineers
STB	Setop-Box
SSH	Secure SHell
TCO	Total Cost of Ownership
TPM	Trusted Platform Module
UX	User Experience
VM	Virtual Machine

Bibliography & References

Bäckman, M., & Hagfjäll, F. (n.d.). *Application security for embedded systems*. (Department of Electrical and Information Technology Lund University) Retrieved from <https://www.eit.lth.se/sprapport.php?uid=1032>

CISCO. (n.d.). *Attaining IoT Value : How to move from Connecting things to Capturing Insights*. Retrieved from https://www.cisco.com/c/dam/en_us/solutions/trends/iot/docs/iot-data-analytics-white-paper.PDF

Deloitte University Press. (n.d.). *Tech Trends 2017 : The kinetic enterprise*. Retrieved from <https://www2.deloitte.com/content/dam/Deloitte/global/Documents/Technology/gx-tech-trends-the-kinetic-enterprise.pdf>

Financial model Edge compute. (n.d.). Retrieved from <https://wikibon.com/the-vital-role-of-edge-computing-in-the-internet-of-things/>

Forrester Consulting. (n.d.). *Understanding the roles of LoB practitioners and SoCs in securing IoT environments*. Retrieved from <https://www.forescout.com/wp-content/uploads/2017/11/Forrester-Survey-Fail-To-Plan.pdf>

Managing IoT devices the DevOps Way. (n.d.). Retrieved from http://sites.tcs.com/blogs/research-and-innovation/managing-iot-services-the-devops-way?_sm_byp=ivvnhk12n6ssvddf

Masek, P., & Thulin, M. (2016). *Container Based Virtualisation for Software Deployment in Self-Driving Vehicles*. Retrieved from http://publications.lib.chalmers.se/records/fulltext/237650/237650.pdf?_sm_byp=iVV0Q1KZQDTQk6SH&_sm_byp=iVVSJFH8FRk5R4p6&_sm_byp=iVVSJFH8FRk5R4p6&_sm_byp=iVVSJFH8FRk5R4p6&_sm_byp=iVVSJFH8FRk5R4p6

McKinsey&Company. (n.d.). *Beyond Agile : Reorganising IT for faster Software delivery*. Retrieved from <https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/beyond-agile-reorganizing-it-for-faster-software-delivery>

- MORABITO, R. (n.d.). *Virtualization on Internet of Things Edge Devices With Container Technologies: A Performance Evaluation*. (Ericsson Research, Ericsson AB, Jorvas 02420, Finland) Retrieved from <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7930383>
- Zhou, W., Zhang, Y., & Liu, P. (n.d.). *The effect of IoT New Features on Security and Privacy : New Threats, Existing solutions and Challenges yet to be resolved*. Retrieved from <https://arxiv.org/pdf/1802.03110.pdf>

Embracing Service Delivery Changes with Machine Learning

Change-Driven Segment Identification and Scoring Can Increase Operator Confidence in Network and Device Modifications and Upgrades

A Technical Paper prepared for SCTE•ISBE by

Andrew Sundelin

Director, Product Management & Cable Innovation

Guavus, Inc.

San Jose, CA

Andrew.Sundelin@guavus.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
1. The Change Imperative.....	3
1.1. Software, Devices and Combinations	3
1.2. Change Initiatives and Concerns	4
2. Change-Driven Segmentation.....	5
2.1. Timing Before and After States	5
2.2. Segmentation and Attributes.....	5
3. Change Scoring.....	6
3.1. Direct Measures	6
3.2. Indirect Measures	7
4. Operationalizing Change Metrics	8
4.1. STB Software Upgrades	9
4.2. Node Split Use Case	9
4.3. All IP Video	10
Conclusion.....	10
Abbreviations	10
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1- Attribute-driven Subscriber Identification.....	6

List of Tables

Title	Page Number
Table 1 – Before and After Attribute Sharing.....	5
Table 2 – Sample Data Value Analysis for Customer Experience Metric.....	8
Table 3 – Node-Split Change Events.....	9
Table 4 – Node-Split Deployment Schedule	9

Introduction

Cable operators face significant challenges in launching and improving services with agility and velocity within ever more complex service delivery architectures. The dilemma is clear: improvement requires change; but change drives performance incidents. So how can operators best test improvements and new services while minimizing unintended consequences?

Evidence-based methodologies, such as A/B testing, that are transforming other industries provide a good answer to the *what* needs to be done. The more challenging question is *how* can operators meaningfully apply these techniques to their operations? Creating a formal evaluation program around every change occurring in MSO access networks would be both labor-intensive and cost-prohibitive, if even possible.

Two techniques that operators can use to improve upon existing ways of tracking and maintaining customer satisfaction before and after changes in their network are the following:

Change-Driven Segment Identification. This is an automated technique for identifying populations and micro-populations of subscribers within the operator network whose service, customer premise equipment (CPE), or service delivery network has experienced the same change.

Change Scoring. Once groups with like changes are identified, the quality of a given upgrade needs to be evaluated. There are many ways to score changes, from straight customer experience measures to automatic feature selection models. Big data and machine learning provide operators with options for systemically evaluating potentially service-impacting change.

The goal of this paper is to describe ways that operators can use machine learning in combination with well-designed operational practices to quickly differentiate between upgrades that improve service quality and those that make it worse – thereby allowing operators to embrace change, not fear it.

1. The Change Imperative

1.1. Software, Devices and Combinations

Not long ago, broadband industry leaders talked about handling one or maybe two major initiatives in a year. While there are still major capital expenditures and large-scale projects, overall change has become more incremental and rapid, less exclusively waterfall and more agile, to use the terms that originated with software developers. That language is appropriate, because of all the modifications occurring in operator networks, software changes are the most common.

In the home, software upgrades and planned improvements can occur to the cable modem (CM), the set-top box (STB), the multimedia terminal adapter (MTA), the home security keypad – even the video camera or other internet of things (IoT) element. In the network, software changes also occur with growing frequency. The cable modem termination system (CMTS), the converged cable access platform (CCAP) (either CCAP core or even the return path demodulator - RPD), the DOCSIS® Provisioning of EPON (DPoE) system, the video controller can all have software changes that potentially impact service.

Combinations of CPE and network changes can also occur with a number of potential permutations that make it impossible to fully validate prior to production deployment. Furthermore, capacity-driven node splits, fiber-deep initiatives, redistribution of CPE across access equipment (or subscribers across shared video controllers) are also common changes to the actual network/service topology, all of which can negatively impact service performance and customer satisfaction.

1.2. Change Initiatives and Concerns

Change creates anxiety in many industries. In a 2015 study of some 300 clients in the area of information technology (IT) Operations Analytics, Gartner reported that 85 percent of performance incidents are traceable to changes. [Cappelli] The quicker cadence and more expansive scope of change in the cable industry is reflected in initiatives and operations involving several areas of technology, including:

Set-top boxes. Legacy devices remain in the field, but the market is no longer characterized by stodgy single-purpose devices with quarterly software release cycles. The reference design kit (RDK), or “platform behind the platform,” now in more than 40 million set-tops and gateways around the world, is a good illustration of how fast-paced, agile techniques are transforming the industry. Last year an industry executive stated that the RDK open-source consortium was shipping more than new 70 features per month. [Ismail]

DOCSIS infrastructure. As data over cable service interface specification - DOCSIS® and WiFi networking standards evolve, bringing greater speed and new architectures, both CPE and network hardware will continue to change rapidly. On the CPE side, new firmware for traditional devices such as CMs and MTAs is more frequent, and the types of IoT devices in the customer premises (most with remotely upgradeable software) are exploding. On the other side, access network devices are also a constant work in progress. Traditional CMTSs, CCAPs and DPoE systems all have software upgrades; and with distributed access architectures (DAAs), decomposed CCAPs are likely to increase this rate of change as the industry undergoes a significant transformation.

Fiber optic nodes. Changes in fiber nodes can have many potential motivations. An operator could be switching to a DAA architecture; “splitting a node” for capacity reasons; or transitioning to an n+0 architecture and going “fiber deep.” Comcast launched a major fiber-deep initiative two years ago, and last year one executive said that after this multi-year ramp-up is complete, the multiple system operator (MSO) could end up with eight to ten times as many nodes. [Breznick]

Video delivery. No more service may be undergoing more change and more difficult to manage than video. Adaptive bit-rate (ABR) technologies help preserve internet protocol (IP) video quality, but an operator has little visibility into those streams. Being one-way broadcast, quadrature amplitude modulation (QAM)-based video has no feedback, apart from subscriber complaints or telemetry from expensive probes, making it a necessary to find predictive metrics elsewhere. Changing architectures, such as cloud-based guides, can add new components to the service delivery architecture. Changes to many video delivery components can impact service in ways that can be difficult to detect or to test in the lab (due to both scalability and permutation challenges).

Apart from those four areas, there are other network elements that can also impact service delivery. The common thread connecting many of the techniques described in this paper is the ability to tie these devices (and, thus, the attributes of these devices, such as manufacturer, model and software version) to individual subscribers. That is the key to creating an anxiety-reducing change management system that is not only dynamic and extensible, but also service-centric and customer-centric.

2. Change-Driven Segmentation

2.1. Timing Before and After States

The idea behind segmentation of operational data is to identify groups of subscribers with identical service changes so as to best assess and compare differences across groups or across time (e.g. before/after a change).

Change-driven segmentation has a number of fundamental concepts. A change-driven segment has a common “before” state (i.e. the value of a subscriber attribute before a change), but possibly different “after” states. Change occurs in different portions of the network at different times and at different scales. Each change-driven segment has an associated time period for that change, typically, a day.

For example, 1 million subscribers might have a Cisco DPC3010 running software version 1.2.3. The operator may be migrating those subscribers from version 1.2.3 to version 1.2.4. The operator may upgrade 100,000 subscribers on January 1; then 250,000 subs on January 8; and finally, 650,000 subs on January 15. However, these 1 million subs, since they share a “before” and “after” state would be in the same “before” segment and the same “after” segment, regardless of the timing of this change for any individual subscriber in that group.

Alternatively, an operator might have the scenario seen in Table 1. Here some subscribers share the before state (e.g. pre-node split) and some share the same after state (e.g. post-node split or merge). But the change-driven segment is uniquely identified by a single attribute sharing the same before and after attribute value regardless of the timing.

Table 1 – Before and After Attribute Sharing

Subscribers	Before Attribute	After Attribute	Change Date	Change-Driven Segment #
100	Node = ABC	Node = ABC-1	August 1st	1
100	Node = ABC	Node = ABC-2	August 1st	2
100	Node = ABC	Node = ABC-3	August 1st	3
100	Node = ABC	Node = ABC-4	August 1st	4
23	Node = An34q	Node = ABC-4	August 12 th	5

2.2. Segmentation and Attributes

This type of segmentation can be driven by any number of service attributes:

CPE Attributes: One segment of subscribers might be those with the same make and model of CM or STB, being upgraded from software version X to software version X+1. In this case, the software version is the service attribute which is changing and driving their membership in that group.

Topological Attributes: Another segment of subscribers might be created when a service group is being divided by the splitting of a particular node. Or when a new CMTS/CCAP is deployed and subscribers are migrated to that CMTS. Here it is a common change in a subscriber’s network topology that is changing and driving their membership in that group.

Billing/Service Attributes: During service changes (e.g. going from QAM-based video delivery to IP-based video delivery) an operator might change a subscriber’s rate code. Or the rate limits on a particular HSD speed tier might get changes. Again, these common changes in subscriber attributes – whether they

happened simultaneously or not – allow for change-driven segmentation of subscribers with “like” changes from a “before” state to an “after” state.

Today these selections are today happening on an ad-hoc, manual basis, and so an exercise such as examining a fiber-deep initiative with speed and precision is difficult. Software changes are also happening to more devices with greater frequency. Thus, programmatic change evaluation programs are labor-intensive and expensive. The amount of change in the network makes them all the more difficult.

A lack of automation in identifying groups who have experienced change is one reason to fear rather than embrace improvement initiatives. Change-driven segment identification that leverages big-data analytics provides a way to automatically identify subscribers who have experienced change. (See Figure 1.)

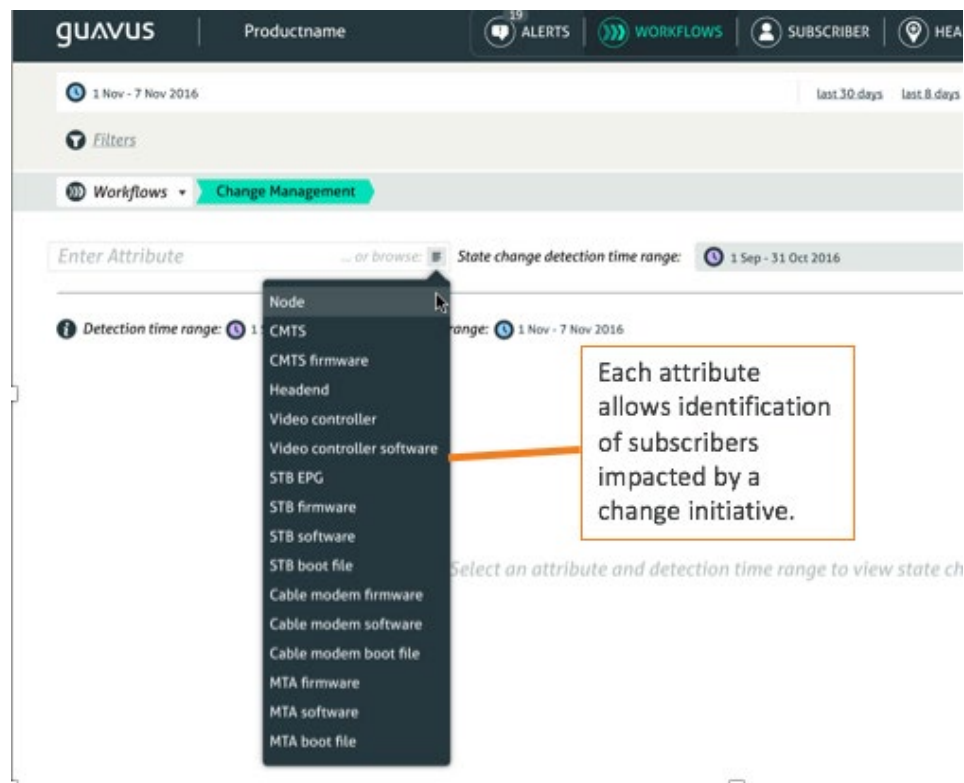


Figure 1- Attribute-driven Subscriber Identification

3. Change Scoring

Once operators can automatically identify groups of subscribers with identical changes, the next step is to evaluate those changes and determine which changes improved customer experience, and which changes made it worse.

3.1. Direct Measures

One way to score an upgrade is simply to look at direct measures of customer experience. To assess a change initiative or other operation, an operator would assemble data before and after to determine whether it yielded a positive or negative result. Every subscriber on the network has a set of attributes. As shown above, those can be device, service, topology, billing or other per-subscriber attributes.

It is common for operators to track all of the care events (e.g. technical support calls, customer-reported trouble tickets and scheduled truck rolls) from individual subscribers. This becomes even more powerful when operators leverage a big-data system to not just track this per subscriber, but also per subscriber attribute. Then operators can know how many subscribers with a given CPE are calling, how many associated with a given CMTS, how many with a specific rate code, and so on.

Thus, the simplest way to determine whether a change leads to a better or worse customer experience is to take a segment of subscribers identified by automatic, change-driven segmentation and look at the rate (by population) of direct measures of customer experience.

3.2. Indirect Measures

But there's another way to score an upgrade. Care events as a direct measure of customer experience are not ideal, because they indicate that a group of customers has already experienced enough pain to take the trouble to contact their operator. There are many other indirect measures of customer experience. These can be service-layer data (e.g. STB errors/reconnects, speed tester data, video probe data, CM reboots, ABR metrics, buffer size, etc.) or network-layer data (e.g. correctable codewords, uncorrectable codewords, RX power, TX power, dropped packets, etc.) None of these measures involve a human picking up a phone and calling the operator, but many of them can be predictive of that behavior.

There is a key challenge to leveraging these indirect measures of customer experience. At some level they do not negatively impact customer experience, but then at some point, the customer experience becomes so bad that customers begin to call, get tickets and instigate truck rolls. How can we determine when these various indirect measures (or features) start to impact customer experience?

The traditional way to determine at what level various network performance measures are contributing to negative customer experience is to convene a group of internal and external multi-discipline subject matter experts to perform a study. They then monitor any number of indirect measures of customer experience and develop a set of thresholds, which, when exceeded, lead to bad customer experience. However, these experts are expensive, and historically, the data has been difficult to accumulate for a holistic view of the entire network.

Perhaps, more importantly, these studies generally have limitations. Because they are expensive, they often happen over a limited portion of the network and over a limited period of time, which means that the resulting thresholds may not be universally accurate across the operator's network and will likely become inaccurate as they age and customer behavior, expectations and applications all change.

A better way to determine when these various indirect measures/features start to impact customer experience is through machine learning. This is a two-stage process.

The training stage begins when the machine-learning algorithm discovers the relationship between features and target variable(s), such as care events. Note that not all attributes are created equally; CPE/UE, subscriber health and network elements tend to be more highly predictive. (See Table 2). At the end of training comes validation, or confirmation through another set of feature data, not included in the training data, that the model is truly predictive of the target variable.

Table 2 – Sample Data Value Analysis for Customer Experience Metric

Attribute	Value	Notes
CPE/UE	High	Actionable, predictive
Plant/telemetry health	High	Directional, actionable
Usage patterns	Med	Noisy, requires proper handling
Subscriber health	High	Directional, actionable, predictive
Device characteristics	High	Actionable
Service tiers	Low	Noisy
Tenure of CPE/UE	Med	Temporally directional
Installation type	Med	Temporally directional
Network elements	High	Actionable, predictive
Plant capacity	Low	Varies based on service type

In the second stage of machine learning-driven analytics, live data is put into the model. The model then outputs a score. That value indicates how predictive of the target (e.g., care events) the input data set is.

For example, in the case of a node split, a machine-learning model could be developed where the features are indirect measures of customer experience (such as, correctable codewords, uncorrectable codewords, receive (RX) power, transmit (TX) power and signal to noise ratio (SNR)) while the target variable is technical support calls. By monitoring a time series of feature values before and after the change, the probability of technical support calls can be predicted (before and after the change). If the probability of support calls goes up, then this change made things worse for customers.

The machine-learning approach is better than a one-time study because the applications that subscribers use change fairly quickly over time, and different applications have different sensitivities to different types of impairments. Machine learning models can be periodically retrained to adjust to these application and behavior changes. This approach is called continuous learning.

To recap this discussion of change scoring, the obvious way to evaluate the user experience before and after a change is to look at direct measures of customer experience and to see if the changed attribute makes the experience of subscribers with that shared change better or worse. The problem with this approach is that it depends upon customer impact and input.

The second way, involving indirect measures of customer experience, can be driven by machine learning. Through a model that predicts the likelihood of poor customer experience, machine learning can output scoring on a time series of indirect measures before and after the change, creating scores that can be used to evaluate whether the change is likely to make the customer experience better or worse.

4. Operationalizing Change Metrics

The vision of this paper is that systems be developed to (a) automatically identify changes to a wide variety of per subscriber attributes; and (b) score changes so as to readily identify those that worsen customer experience or are predicted to do so.

The scoring aspect of the system could be a daily report that lists all of the change-driven segments identified in the network sorted by the difference in the score for that segment before/after the change. (See Figure 2.) Thus, a team could easily identify bad changes to the service or network in order to more quickly roll them back.

Table 3 – Node-Split Change Events

<input type="checkbox"/>	Oldest Change	Newest Change	Changed Subs	Old State	New State	Impact (Customer Experience Score)
<input checked="" type="checkbox"/>	9-Oct	9-Oct	10	AAA	BAAA1	+10.5
<input checked="" type="checkbox"/>	9-Oct	9-Oct	10	AAA	BAAA2	0
<input checked="" type="checkbox"/>	9-Oct	9-Oct	9	AAA	BAAA3	-2.5
<input type="checkbox"/>	7-Sep	7-Sep	46	PDQ	PDQ1	-3.7
<input type="checkbox"/>	12-Oct	12-Oct	23	XYZ	XYZ1	0
<input type="checkbox"/>	12-Oct	12-Oct	16	XYZ	XYZ2	-1

4.1. STB Software Upgrades

At some operators, this basic model is already being followed for modern STB software upgrades. However, this usually requires a complex program of phased upgrades where larger and larger groups of subscribers receive the same upgrade over time. It also requires an equally complex method of evaluation, where a priori knowledge of the groups being upgraded at specific times is used to create a group to be monitored. This entails tracking which devices are being upgraded when and where, as well as specific monitoring to see if the upgrade is good or has unintended consequences.

Operationalizing this approach also permits broadband operators to apply the web technique of A/B testing of alternative software upgrades, allowing them to readily determine which of two implementation options is better. Thus, change analytics can increase agility, enabling a business to “lean in” to quick learnings.

To summarize, phased software upgrades remain a best practice, but a continuous and comprehensive approach to evaluating changes across many different attributes removes the need to specifically track when and on what specific devices these phased changes are being made. Plus, it brings these benefits to other types of changes, beyond software upgrades.

4.2. Node Split Use Case

The industry’s hunger for fiber in the access plant is growing, yet with no common mechanism today to automatically identify “bad” node splits, there is no quick and easy way to know how well these initiatives are going. A pace of 25,000 or 50,000 per year leads to a nightly deployment rate that is too numerous to track manually. (See Table 2.)

Table 4 – Node-Split Deployment Schedule

Per year	Per week	Per night
25,000	481	120
50,000	962	240

Suppose there are bad splits. A scenario requiring an additional truck roll for each split could add \$3 million in the case of 50,000 deployments per year. (\$3m/year = 50,000 x 1 x \$62.) If only one in six is bad, that’s another \$516,000. Adding in service calls that might follow a defective node upgrade could easily and significantly increase that amount.

Given the high stakes, operators need to identify bad nodes as quickly as possible. Without scoring, however, there is no baseline for subscribers following a split, and with no baseline, you’re looking at up to three weeks to begin detecting anomalies. No scoring mechanism also means no before/after metrics,

or some way of understanding how customer experience on the original nodes differs, if at all, from customer experience on the new or fiber-deep nodes.

4.3. All IP Video

Many changes to the network are designed to escape subscriber notice. They may not even be visible when looking at subscriber-related attributes such as STB make/model/software version.

If a subscriber has a STB which is capable of both IP- and QAM-based video, then device-oriented attributes will not allow an operator to differentiate between subscribers' being delivered all-IP video and those being delivered primarily QAM-based video. However, they may be visible when looking at the billing system, or other subscriber attributes. Some larger operators are using a different rate code for their subscribers on all-IP video, and this is one way they can be identified on a per-subscriber basis.

Thus, operators can leverage Change-Driven Segmentation to identify subscriber populations with like changes to a vast array of subscriber-related attributes. This can either be part of a formal phased-roll out process or more ad hoc, but the technique works equally well in either situation – even when the like changes happen at different times. Once these subscriber populations are automatically identified their before/after experience can be scored (Change Scoring) such that changes which make customer experience worse can be quickly backed out or otherwise mitigated.

Conclusion

Network and device upgrades and other change initiatives cause concern among technical operations and customer care personnel. This is the case among service providers and tech companies of all stripes, but it is especially so in cable as the pace of change and complexity of the industry's service delivery platforms increases. But evaluating every change occurring in MSO access networks using traditional means is a non-starter. The costs are too high, resources too limited and workload overwhelming.

An attractive, efficient and cost-effective alternative is to leverage change-driven segmentation techniques and machine learning-based change scoring. The benefits include a reduction of uncertainty and anxiety surrounding change; greater clarity about which changes promote and which detract from customer experience; and advancement of the long-sought goal of transforming a network-centric view of operations into one focused more on services and the subscriber.

Abbreviations

ABR	adaptive bit rate
CCAP	converged cable access platform
CM	cable modem
CMTS	cable modem termination system
CPE	customer premises equipment
DOCSIS	DOCSIS over Cable Service Interface Specification
DAA	distributed access architecture
DPoE	DOCSIS Provisioning over Ethernet
HSD	high-speed data
IP	Internet protocol
IoT	internet of things

MSO	multiple systems operator
MTA	multimedia terminal adapter
QAM	quadrature amplitude modulation
RDK	Reference Design Kit
RPD	Remote PHY Device
RX	receive level
SNR	signal-to-noise ratio
STB	set-top box
TX	transmit level
UE	user experience

Bibliography & References

Breznick, A. Why Cable's Feasting on Fiber. Light Reading, Nov. 9. 2017.

Cappelli, W. Causal Analysis Makes Availability and Performance Data Actionable. Gartner, Oct. 7, 2015.

Cheikhrouhou, A., and J. Davenport, A. Kelkar. The Imperative of Customer-Centric Operations. SCTE/ISBE, 2017.

Cunha, G. Approaches for Proactively Managing Customer Experience and Reducing OPEX in a Cable Operations Environment. SCTE-ISBE 2017.

Dorairaj, S, and C. Basian, B. Burg, N. Pinkernell. Simplifying Field Operations Using Machine Learning., SCTE/ISBE, 2017.

Ismail, L. RDK Product Roadmap Update, Inside RDK. Summary of session offered at European RDK Summit, Sept 17, 2017.

Sundaresan, K. and N. Metts, G. White, A. Cabellos-Aparicio. Applications of ML in Cable Access Networks, NCTA, 2016.

Sundaresan, K., and J. Zhu. Access Network Data Analytics. SCTE/ISBE, 2017.

Sundelin, A. Leveraging Machine Intelligence and Operational Analytics to Assure Virtualized Networks and Services, SCTE/ISBE, 2017.

Enabling Smart Cities By Leveraging IoT Sensors, Multi-Building Modeling And Analysis, And Smart Energy Business Case Analytics

Digital Transformation For Energy Savings Programs

A Technical Paper prepared for SCTE•ISBE by

Mark Stratton

VP Global Digital Solutions Enablement

Hitachi Consulting Corporation

Denver, CO

303-253-2521

mark.stratton@hitachiconsulting.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. Portfolio Energy Savings	3
2. Building Energy Modeling at the Portfolio Level	4
2.1. Energy Conservation Measures and Cost Simulation	5
2.2. Data Science and Cloud Computing	6
3. Enabling Investment Grade Savings.....	8
4. Hardening the Business Case with Pilots/POC	8
4.1. Leveraging IoT for M&V	9
Conclusion.....	10
Abbreviations	11
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1 - Example: Large Cable Operator Portfolio	4
Figure 2 - Building Energy Modeling Mapping To Portfolio	5
Figure 3 - Building Energy Model Process.....	6
Figure 4 - Example Calibrated Modelings.....	7
Figure 5 - Cloud Computing Environment.....	7
Figure 6 - Building Case Output & Recommendations	8
Figure 7 - IoT System Deployment Onsite.....	9
Figure 8 - Actual and Relative Baseline For Savings Performance.....	10

Introduction

Modern IoT smart energy platforms have the potential both to power smart city intelligence via sensors, analytics and services, and to drive cities to become smarter. Platform providers and ultimately cable operators themselves can do this by leveraging their knowledge and experience in traditional large-scale smart building practices that have been undertaken in recent years in both telecom and non-telecom infrastructure. They can also leverage their experience in selecting the right energy efficiency improvements for small- and medium-sized cable edge facilities that can be applied to small and medium businesses with similar size and payback constraints. The platforms and tools already used for large scale analysis of cable facility portfolios can now be applied to any building type at the portfolio, campus, city and national scales. Recent use of these tools has resulted in business cases being made for early retirement of aging HVAC equipment, deployment of cost-effective energy conservation measures that payback on the order of 3 years (even for edge facilities in states with low utility rates), and analysis and prioritization of LED lighting improvements. And these are just some of the potential energy conservation measures that can be analyzed, prioritized, and planned with a comprehensive platform for helping buildings, portfolios and cities get smart about energy use. In this paper, the experience and lessons learned in cost-effectively selecting and deploying energy conservation measures in cable infrastructure and facilities will be shown to be directly applicable to new services and revenue streams that platform providers and cable operators can develop or partner on for smart city initiatives using IoT technologies and smart energy tools. Coupling these tools with diverse funding models will further enable companies and entire cities to deploy smart energy management and energy conservation measures cost-effectively and even as a service.

Content

1. Portfolio Energy Savings

The facility portfolio of a cable provider often has a wide breadth of building types, a split on leased and owned buildings and typically covers a diverse geography and can scale from as small as a few hundred facilities to as large as tens of thousand facilities when including retail. Characteristics of the portfolio include a wide distribution in facility age, equipment age, mixed use facilities and various facility types like headend, hubs, service centers and administrative office space. Not to mention what also gets included from mergers and acquisitions. More than ever before there is pressure to reduce and control energy spend as a function of these facilities, with energy and maintenance typically being two of the larger expenses to the P&L from a building lifecycle perspective. Two important business directives from the top down are cost savings (aka energy savings) coupled with sustainability reporting or what is sometimes called carbon dashboards. Relevant energy savings achieved can be leveraged in quarterly and annual company financial reporting to help demonstrate the cable operator cares about the environment and to keep sustainability goals in front of their subscribers and key financial stakeholders.

When portfolios are large like this cable operators seek to identify and run energy savings programs against the portfolio that are designed to produce a relatively fast return on investment. Technologies in the form of energy conservation measures, e.g. lighting upgrade, often produce a simple payback in 2-3 years. The sought out win here is to apply company investment in a way that delivers a quick near-term result generated from the savings the program can achieve. This approach often works well for a handful of large facilities or select targeting of facilities that have a proportionately larger than average energy bill. After capitalizing on what's often referred to as the low hanging fruit, the problem gets harder when

trying to make the determination of which facilities to treat next and how far to go into the long tail of the building portfolio. This is where finding the energy opportunity for savings can become a more difficult challenge. Two examples for cable operators are cable hubs, which is often one of the highest count in the total number of building in a portfolio and headend and retail locations. The table below provides a representative example facility distribution of a typical cable operator portfolio in the USA.

**Large Facility Portfolio :
Cable Operator**

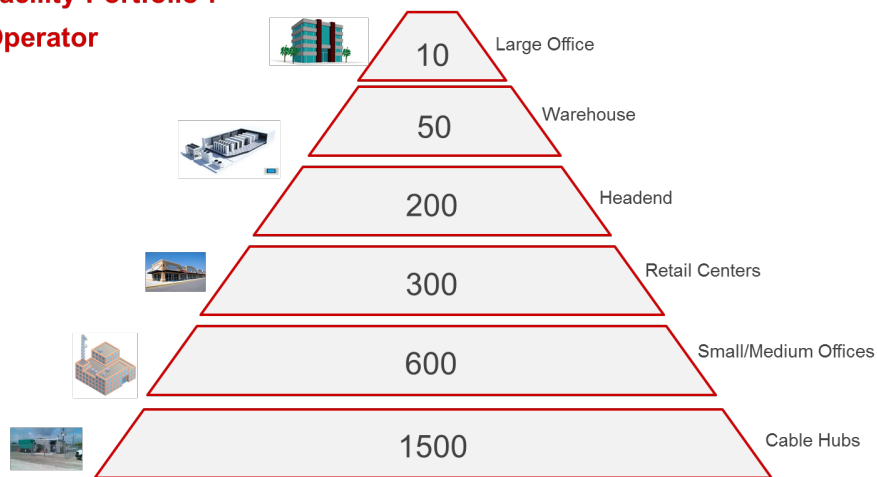


Figure 1 - Example: Large Cable Operator Portfolio

A large building portfolio like this has vast differences between mission critical facilities and general office buildings. The need remains to find energy savings potential deeper within the facility portfolio at scale for a larger number of typically smaller, medium size and specialized facilities like outside plant or other edge facilities. The next section demonstrates a new approach to solve this problem by leveraging building energy modeling and simulation designed to help yield the best potential energy savings for cable operators with large building portfolios.

2. Building Energy Modeling at the Portfolio Level

Building energy modeling is not new, however, the application of it at the portfolio level and focused on the energy savings economics and efficiency programs is new. The age-old mindset in the industry is that every building is a snowflake, for which there is some valid truth, but we have found through specialized building modeling approaches that information can be analyzed to help make better investment decisions when looking at this problem from the perspective of the entire building portfolio. Utilizing a basis of Energy Plus modeling from the Department of Energy, and leveraging data science and cloud based computational scale, as a team we modified the models and embodied them in a software application that allows building energy engineers to load data representing the building portfolio and run various scenarios to determine what are the lowest risk, highest paying investment options that yield the highest potential for energy savings within the portfolio. Going well beyond the Energy Plus models, new modeling of multiple energy conservations measures and modeling implementation cost based on materials and labor is included to help drive data centric business decisions. In the industry we call this investment grade decision making. The more information we have the better the decision we can make.

Examples of inputs at this level, but insufficient to give away our proprietary approach, are number of buildings, building addresses, building size and actual energy history. Once we have the client data on the building portfolio, we initiated a mapping exercise. If you are familiar with building energy models, rather than develop a unique model for each building, we group the facility portfolio into facility types and then customize a building energy model (BEM) per building type. This proto-typical model becomes the base model that is utilized to inherit the individual information of all facilities of that type, for example cable hubs or medium size office buildings. The below figure is an example of just such a mapping of the facilities within the portfolio grouped together by type and energy spend.




Figure 2 - Building Energy Modeling Mapping To Portfolio

To ensure our modeling was correct an engineer performed a review on the buildings to determine typical approaches to cooling, heating, weather-based location and many other factors to ensure the attributes of each building in the portfolio are accurately represented at this level. It is important to note at this level a deep audit or assessment at the facility is not required, as the purpose here is to model at the portfolio level to achieve a business outcome.

2.1. Energy Conservation Measures and Cost Simulation

Energy conservation measures are improvements that can be made to existing facilities to improve their overall energy performance and optimization. Often technical in nature, examples are HVAC controls, new refrigerant technologies, air flow optimization, efficient lighting and many others. Building energy models do a great job at modeling the existing building and its core infrastructure. However, the Energy Plus software provides limited ECM modeling capabilities today. We enhanced the building energy models to include select ECM technologies and more importantly modeled their cost characteristics.

While building modeling leverages statistical techniques to converge on both overall building calibration and sensitivity analysis. The modeling of the ECM and cost allows us to go further and simulate the

impact of potential energy improvements to the building. Cost modeling and simulation like this is new and has the extremely valuable result of gauging which ECM or combination of ECM enhancements make the most sense across a diverse and large building portfolio. We call this the optimal yield, that is a data driven approach to making better informed decisions about where to make portfolio improvements versus the more standard industry approach of vendor quoted savings and payback terms based on a flat savings rate.

2.2. Data Science and Cloud Computing

Data scientist were leveraged to develop software algorithms to create a statistical approach to the building energy modeling and the associated convergence process. Historically, buildings are modeled one by one on what is often called a desktop analysis by individual engineers to help determine where in the facility to save energy. In our approach cloud computing was leveraged to run 1000s of iterations enabling statistical convergence of the solution at scale for a large number of facilities in a reasonable time and within very specific confidence levels to ensure results are accurate at this level. We leverage AWS as the platform and elastically we run as many EC2 instances as required to generate results quickly, often within days for a large portfolio.

The data science becomes extremely interesting at this level. For those familiar with this subject, there are guidelines in the industry around approaches and methodology for ASHRAE guidelines, regional building codes and many other factors. The data science gets applied in the statistical approach, for those fascinated by these things, like non-sorting genetic algorithms (NSGA2), Strength Pareto Evolutionary algorithm (SPEA), Linear Discriminant Analysis algorithm, hierarchical clustering algorithm. It's amazing to think that genetic algorithms are being utilized to improve building efficiency and embarking us on the path to introduce more machine learning and artificial intelligence.

As part of this work, we also had to develop specialized tools for automation, enabling the launch of multiple parallel jobs, multiple docker containers, a common data repository and we are still working on things like automating geometry generation.





Figure 3 - Building Energy Model Process



Calibration Statistics

CV(RMSE) = 9.4%, NMBE = 2.7%

Figure 4 - Example Calibrated Modelings

Energy engineers guide the entire process to check and confirm results make sense. The immense value of this approach is again the ability to approach an entire building portfolio independent of size and drive meaningful investment grade savings decisions.





Figure 5 - Cloud Computing Environment

3. Enabling Investment Grade Savings

Those investing in energy efficiency and savings programs are business owners. They are making business decisions and want the information and data that helps them become more informed in the way they make decisions. The information below is an example of some output and recommendations that comes from a portfolio analysis at a very high level when using building energy modeling at the portfolio level. The essential take away is how to determine where in the portfolio you can yield the most savings and what it will cost to run a program, understanding the simple payback and return on investment.

What is net Energy Efficiency Impact of Recommended ECMs?



What is the estimated program cost and payback by ECM?

(\$ in millions)	Cost Estimate		Savings Estimate		Payback (Year)	
	Low	High	Low	High	Low	High
LED Retrofit	2.30	2.80	0.75	1.25	1.84	3.73
Intelligent Lighting	1.20	1.50	0.60	0.70	1.71	2.50
Packaged HVAC	0.75	1.20	0.60	0.80	0.94	2.00
Compressor Control	1.10	1.30	0.40	0.50	2.20	3.25
Variable Frequency Drives	3.30	4.00	0.35	0.55	6.00	11.43
HVAC Replacement	1.40	1.60	0.40	0.50	2.80	4.00
Refrigerant Replacement						

What is the program benefit, cost and simple payback by geography, division or facility type?

Division	Total Cost			Total Savings			Simple Paybacks		
	Est. Total Cost (Low)	Est. Total Cost (Medium)	Est. Total Cost (High)	Est. Total Savings (Low)	Est. Total Savings (Medium)	Est. Total Savings (High)	SPB (Low)	SPB (Medium)	SPB (High)
Central	\$3,235,339	\$5,392,232	\$8,627,571	\$999,214	\$1,665,357	\$2,664,571	1.2	3.2	8.6
HQ	\$216,266	\$360,443	\$576,709	\$69,431	\$115,718	\$185,149	1.2	3.1	8.3
Northeast	\$2,236,480	\$3,727,466	\$5,963,946	\$695,205	\$1,158,675	\$1,853,880	1.2	3.2	8.6
Other HQ Ops	\$1,795,968	\$2,993,280	\$4,789,248	\$604,830	\$1,008,050	\$1,612,880	1.1	3.0	7.9
West	\$2,152,211	\$3,587,019	\$5,739,230	\$683,620	\$1,139,367	\$1,822,987	1.2	3.1	8.4
Grand Total	\$9,636,264	\$16,060,440	\$25,696,704	\$3,052,300	\$5,087,167	\$8,139,467	1.2	3.1	8.4

Figure 6 - Building Case Output & Recommendations

The results are designed to help guide facilities management and corporate energy and sustainability teams to select which approach makes the most sense based on priorities, funding and business climate. However, in all cases a next step in our scenarios is to recommend a pilot, proof of value or proof of concept.

4. Hardening the Business Case with Pilots/POC

Before moving to a larger scale investment across a sub-section of the portfolio, it is highly recommended that a limited pilot of a specific building type associated with a prototypical be run as the next step in this journey. The reason for this is that the recommendation can be implemented on a limited scale and tested to ensure that in the real-world the program at scale will produce empirical results that validate the savings potential. This is used as a methodology to harden the business case with more specific site values and conditions, which as a second phase helps to further verify and provide information that further reduce risk and ensures savings are still aligned with the investment strategy.

4.1. Leveraging IoT for M&V

At the pilot level limited deployment of the proposed energy conservation measures is typically insufficient. Often a M&V process is utilized to measure before and after results of the technology improvements. Doing this on a limited scale helps prove the savings really exist. More importantly, doing this on a limited scale help prevent the issue of pre-maturely investing and scaling without knowing the return will manifest itself over time. Reducing investment risk is the key concept we are pursuing here.

This is where IoT comes into play and helps with the digital transformation of the M&V process but also sets the stage for measuring and monitoring energy for the long term to ensure continuously programs are generating the savings they set out to accomplish.

This is where our approach and methodology becomes more interesting. Leveraging the building model as a digital twin and through the pilot process, a deeper audit and assessment of the facility is performed. This improves and deepens the data that goes into the building energy model. This is important because it lays the groundwork for what's call IPMVP Option D or calibrated simulation, which leverages the ASHRAE guideline 14 for M&V procedures. Only for our purpose, we were pursuing this for real-time calibrated simulation using real-time telemetry by using IoT submetering of the facility and ECM technology.

By introducing real-time submetering associated with the assets which had ECM improvements we could measure a more accurate before and after. Why this is important is due to the dynamic load changes, workload changes, upgrades, and equipment changes typical of what cable facilities go through over time. Straight line regression is no way to pursue an investment for 3 or more years. It's become downright impractical. We wanted to truly measure the before and after and apply IPMVP methodologies to make it more standardized.

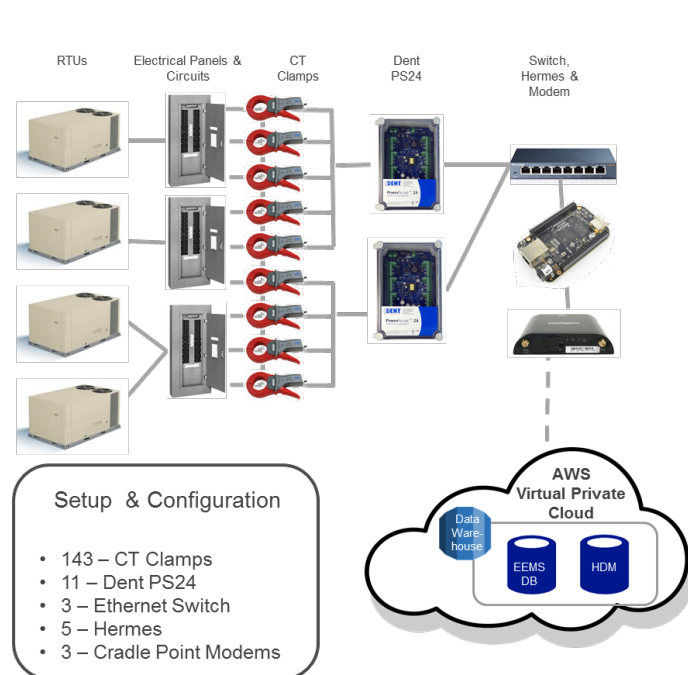


Figure 7 - IoT System Deployment Onsite

The data captured onsite is transported via a raspberry PI device to the Amazon cloud where specialized dashboards were created to track the energy use trend from the facility and sub-metered at several points within the facility. The facility building energy model in this circumstance is utilized to forward model the previous state of the building creating a relative baseline that changes with the building, to model the facility in a predictive way to model the behavior in the future. We call this a relative baseline. Why this is so important is simple, far too often we have seen improvement made to a building and the savings get washed out because of other changes and modifications to the facility over time. This makes it impossible for the energy or facility manager to prove the investment was worthwhile when their boss asks the question as to why the energy bills are now higher a year into the program.

Data through IoT can solve this problem. It demonstrates where all the changes are coming from and helps to differentiate between weather normalization, new IT load, facility improvements/changes and many other factors that could impact energy consumption over time.




Figure 8 - Actual and Relative Baseline For Savings Performance

IoT is wonderful in this way, as the system telemetry tells the truth about what happening and provides a fresh perspective on how operations are performing over time. This has the added benefit of enabling better insight into long term optimization. For example, things like performance drift can be spotted related to maintenance conditions allowing for appropriate corrective actions.

Ultimately, if the pilots go well and prove successful, then this can lead into a scale-out program to introduce changes across the majority of the buildings under consideration. The IoT and telemetry are optional at scale, however, it clearly can provide significantly more insight than previously.

Conclusion

In conclusion the use of building energy modeling with enhancements and paired with IoT becomes a very powerful approach to understanding and allowing facility and energy managers to better understand and have firsthand data and information for making energy related decisions. The portfolio approach helps in selecting programs with the best potential savings yield for the business. The building digital twin combined with IoT through pilot ECM deployment helps further reduce the risk and greatly improves the understanding of potential deployments. Then keeping the IoT in place helps with the process of continuously optimizing and understanding relative changes over time. It doesn't make the

building any smarter, but it certainly makes those responsible for energy and facility management a lot smarter and much more insightful regarding that status of building operations over time.

It is our goal to help introduce this energy insights capability as a more productized offering. If we can scale this solution across the entire building portfolio to measure and monitor energy consumption and savings in real-time, the same holds true for the cable operators to potentially offer the same services to their business clients. The commercial potential is virtually limitless when considering leveraging the cable operator network to enable new services like this in their own solutions portfolio.

Abbreviations

ECM	Energy conservation measure
IPMVP	International Performance Measurements & Verification Protocol
ASHRAE	American Society Heating, Refrigeration and Air-Conditioning Engineers, Inc.
BEM	Building Energy Model
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

ASHRAE Guideline 14-2014: Measurement of Energy Demand, and Water Savings; 2014 ASHRAE, ISSN 1049-894X

EnergyPlus™ Version 8.9.0 Documentation, Using EnergyPlus for Compliance, U.S. Department of Energy, March 23, 2018

International Performance Measurement & Verification Protocol; Concepts and Options for Determining Energy and Water Savings; Volume 1, March 2002, International Performance Measurement & Verification Protocol Committee, DOE/GO-102002-1554

M&V Guidelines: Measurement and Verification for Performance-Based Contracts Version 4.0, U.S. Department of Energy, Federal Energy Management Program, November 2015

F. Campolongo, J. Cariboni, A. Saltelli, An effective screening design for sensitivity analysis of large models, Environ. Model. Softw. 22 (10) (2007)1509–1518.

Menberg. K, Heo. Y, Chourhary. R (2016), Sensitivity analysis methods for building energy models: Comparing computational costs and extractable information. Energy and Buildings 133 433-445

Enhancing Service Agility for the Enterprise Customers using an Integrated Orchestration and Test Automation Solution

A Technical Paper prepared for SCTE•ISBE by

Shiby Parayil

VNF Certification and Cloud Deployment Business Leader,
Ericsson North America

Earl Villanueva

Head of Solutions and PMO for Cloud, Orchestration, and NFVI
Ericsson North America

Ericsson

6300, Legacy Drive, Plano, Texas
Shiby.Parayil@Ericsson.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Industry Context	3
Increasing the Service Velocity for Enterprise Services	4
1.1. Business Considerations for the Automation Solution	4
1.2. Ericsson Solution Approach	4
1.3. Integrated Orchestration and Test Automation Solution	6
1.4. New Operational Context: NetOps – DevOps for the network.....	9
1.5. Key Benefits to the MSO from the approach	10
Conclusion.....	11
Abbreviations	11
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1: Integrated Service Orchestration and Solution Phases.....	5
Figure 2 :Interface points for the SLA validation	5
Figure 3: Integrated Service Orchestration and Test Automation Solution High-Level Functional Architecture Diagram.....	7
Figure 4 :Continuous Deployment and Delivery Facilitated by the Solution	9

List of Tables

Title	Page Number
Table 1: Interaction work flow for the Continuous Deployment and Delivery	10

Introduction

The ability to rapidly launch and operationalize new services is a critical success factor for Multi-Service Operators (MSOs) competing in the enterprise and business space. Software Defined Networking (SDN) and Network Functions Virtualization (NFV) are key technology enablers for this. Deploying and managing the multi-vendor network cloud services is a very complex task due to the degree of change across virtualized service chains. Standards are still maturing, Virtual Network Function (VNF) vendors are releasing software and patches more frequently, and the components of the underlying NFV infrastructure are also subject to change. Managing these changes require a highly automated and integrated approach to network service orchestration and test automation. Ericsson has addressed these key challenges for our customers globally to accelerate the velocity of new service introduction.

Increasing the service velocity requires a holistic transformation for MSOs across many aspects of the MSO's operating context. MSOs need to transform from a network centric organization to a customer centric organization. In line with this, there are two key indicators:

- Service Agility – resulting in quicker time to revenue
- Operational Efficiency- Reducing the CapEx costs by utilizing the network equipment better and leveraging automation for OpEx improvements.

In this paper, to achieve these key indicators, Ericsson shall focus on the processes, automation approach, and best practices and risk mitigation an operator must undertake. Ericsson shall also examine the phases of the transformation journey for the MSOs, and key considerations for the MSOs, to make a successful transition. Specifically, we examine how an operator can leverage an integrated orchestration and test automation solution framework that can significantly accelerate the service velocity of launching new enterprise services in an operationally efficient manner.

Industry Context

Based on the global industry analysis, key drivers of Cloud (NFV/SDN) adoption that are improving Service Agility and Operational Efficiency [1] include the following:

- Introduce new services and gain revenue faster
- Improved customer experience using on-demand and self-service
- Scale services up or down quickly.

Some of the barriers for the Cloud adoption include

- Products available in the market or in open source communities are not carrier grade
- Multi -vendor VNF integration can be expensive, complex, and risky
- The changes to Operations Support System (OSS) and Business Support System (BSS) required for NFV.

In the NFV Adoption strategies, we are seeing three approaches for the deployment and delivering Cloud solutions. These are:

1. **Fully decoupled approach:** MSO takes the responsibility of building the cloud stack. This includes building the entire cloud stack in house or selecting specific vendors for Network Function Virtualization Infrastructure (NFVI), SDN provider, VIM layer, orchestration (resource

and service). This approach provides the maximum flexibility and has the highest inherent risks. 47% of the top 20 operators are embracing this approach.

2. **NFVI Integrated Stack Provider:** MSO selects a single vendor to provide the cloud stack. The vendor provides the NFVI/VIM and is responsible for providing the cloud and infrastructure for the application providers (VNF providers). Typically, orchestration is not included as part of the NFVI Integrated Stack. This approach provides reasonable flexibility and eliminates substantial risks. 51% of the top 20 operators are going down this approach.
3. **NFVI Full Stack:** MSO selects a single vendor to provide the cloud stack including the orchestration and the VNFs. This approach provides the least risk for the customer. Based on the marker survey, majority of the Tier 2/3 customers are considering this approach to meet their business goals.

Ericsson sees that the pendulum for the cloud adoption swinging towards more risk averse options (Option 2 and Option 3). One of the motivations for the shift is the increase in service velocity associated with Option 2 and Option 3.

Increasing the Service Velocity for Enterprise Services

MSOs are looking to increasing the service agility to serve their customers. However, we are seeing that many VNF vendors are moving from few VNF software releases to frequent releases, few products to many products. Unlike the physical world, the cloud infrastructure is also undergoing significant change. This results in the increase in certification cycles triggered by application upgrade, patch releases or OpenStack upgrades.

MSO needs to develop an automation solution that addresses the service onboarding, service validation and Life Cycle Management of the multi-vendor network services and NFVI.

1.1. Business Considerations for the Automation Solution

The automation solution needs to address three key aspects

- Reduce Operation Expenses (OpEx) associated with testing
- Increase Service Agility for Customers
- Enable Innovation for MSOs by enlarging the ecosystem of VNF vendor partners

1.2. Ericsson Solution Approach

Ericsson proposes an integrated service orchestration and test automation solution to address the MSO business needs. Integrated approach provides a solution-based approach to address service design, service onboarding, service validation and service LCM.

The automation solution will be realized on multiple phases as given below. One of the key solution components is the end to end service orchestration solution that is integrated with the test automation libraries for the validation of the NFVI and VNFs

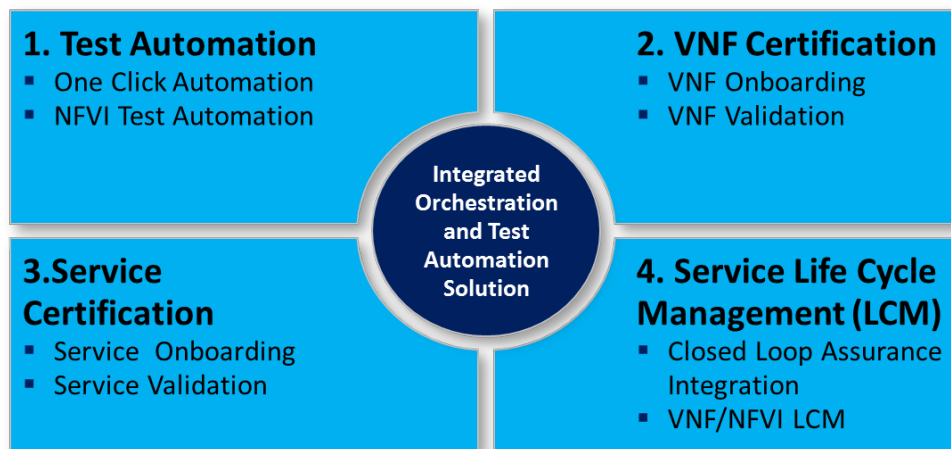


Figure 1: Integrated Service Orchestration and Solution Phases

Step 1: Test Automation

VNF Test Automation

MSO implements a test automation solution for testing the network functions. The goal is to create a **One Click Automation solution** wherein a test automation layer executes a series of automated test suites for testing the VNF, network services and NFVI. This step is applicable for the Physical Network Functions and/or Virtual Network Functions. This phase can happen today if the operator has PNFs or yet to deploy VNFs. Typically, this is done using licensed test platforms like IXIA.

NFVI Test Automation

Clear technical Service Level Agreements (SLAs) at Integration Points (red dots above) are required for VNFs to be able to commit to their performance SLAs. The red dots in the figure represent the interface points where the NFVI SLAs need to be verified.




Figure 2 :Interface points for the SLA validation

Therefore, the NFVI's SLA compliance becomes the operator's commitment to the VNF vendors. Consequently, NFVI SLAs must be verified and guaranteed for Operators to be able to require VNFs' SLA compliance from the VNF vendors. Automated NFVI SLA compliance tests are needed in pre-operation validation as well as during the full Life Cycle.

Step 2: VNF Certification

In Step 2, MSOs develop a process automation methodology to onboard the VNFs to the MSO cloud environment. Typically, MSOs develop self-service portals which are integrated with the NFVO MANO functionality to instantiate the VNFs on the MSO cloud. Orchestrator will use the VNF artifacts to instantiate the VNF on the MSO cloud. Once they onboard, the next step is to automate the validation of the VNFs by triggering the test automation suites developed in Phase 1. VNF Vendors have varying levels of maturity. Hence, standardizing the entry criteria helps the MSOs to streamline the operations.

Step 3: Service Certification

In Step 3, MSOs extend the process to include the instantiation and the validation of Virtual network services. Virtual network services include multi-vendor VNFs. (e.g. VERSA SD-WAN with Palo Alto Firewall as a service to customers). Typically, this stage involves the integration of the end to end service orchestrator, the creation of service contexts for the orchestrator like Network Service Descriptors (NSDs)/TOSCA templates for the service instantiation and service configuration. Service Orchestration component of the solution creates the end to end service. Once the service is up and running, the service validation is automatically triggered by running end to end service tests

Step 4: Service Life Cycle Management (LCM)

MSOs implement a closed loop automation mechanism to monitor the functional and performance aspects of the multi-vendor service chains, VNFs and network infrastructure. Closed loop Assurance can be implemented at multiple levels”

- Infrastructure level: Monitoring is done for the data center and resources
- VNF& Network Services: Closed loop is implemented at the end-to-end service level with integration to Service Orchestrator. It can also be done at VNF level where the VNFs.
- Analytic recommendations for closed loop based on insights.

This is a highly evolved state where the service assurance engine is integrated with the end to end service orchestrator and performs highly complex functions like the auto-healing of network service (scaling up and down a network service etc.). The solution helps in managing the software versions of NFVi, VNF and network services

The business value and the technical complexity increases with each phase. Each of these phases is a logical extension of the previous phase and needs to be carefully designed and architected. Choosing the right strategic partner to assist the end to end transformation journey, is critical to the business success.

1.3. Integrated Orchestration and Test Automation Solution

Integrated Service Orchestration and Test Automation Solution should provide a vendor agnostic platform to facilitate the process automation to onboard, validate VNFs and virtual network services and do the LCM for VNF, NFVI and network services.

Figure 2 provides a functional architecture diagram of Integrated Service Orchestration and Test Automation Solution.

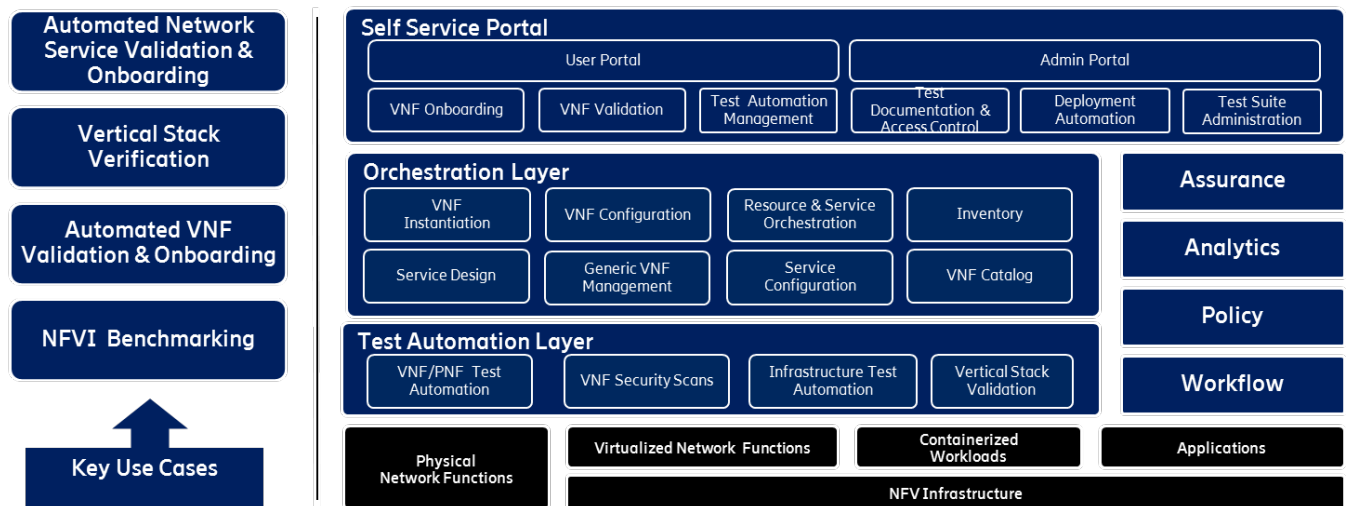


Figure 3: Integrated Service Orchestration and Test Automation Solution High-Level Functional Architecture Diagram

Some of the key components include:

Self Service Portals

The Integrated Service Orchestration and Test Automation Solution platform provides the self-service portals for vendors to onboard the VNFs into the MSO cloud environment. It also provides the administrator portal for the MSO to authorize vendor access and to define the validation criteria for the vendors and VNF types. The solution needs to have a dashboard that displays the results of the onboarding step and the validation tests.

Some of the key functions that are enabled by the self-service portal include

- VNF Onboarding – Uploading the VNF images and artifacts and then using the orchestrator to instantiate the VNF on the MSO cloud.
- VNF Validation – Functional validation of the VNFs
- Test Automation Management- Validating the VNF against multiple levels of tests like cloud compliancy tests, functional tests, smoke tests and others.

Some of the key functions that are enabled by the admin portal include

- Test Documentation and Access Control – Uploading the test documentation for the VNF types and managing the access to the self-service portal to the VNF vendors.
- Deployment Automation – Ability to push the golden image for the production deployment.
- Test Suite Administration- Ability to manage the test suites for the various test levels for the VNF type.

Orchestration Layer

Orchestration layer provides the following functions.

- VNF Instantiation – Using the VNF image and the templates, instantiate the VNF in the operator cloud.
- VNF Configuration – Automating the configuration of the VNF in the operator environment.
- NFV Orchestration- This provides the resource orchestration functionality as specified by the ETSI MANO specification.
- Inventory – Provides the inventory management of the VNFs. This layer provides the repository for the VNFs
- Service Design – Service Design Center provides the ability to design the network services. This is realized by the integration with the SDN Controller
- Generic VNF Management – Provides the G-VNFM functionality and talks to the specific VNFM to instantiate the VNFs.
- Service Configuration – Use the TOSCA templates to define the network service chain
- VNF Catalog- Provides the list of the VNFs which are part of the VNF Catalog.

Validation Layer

The Validation layer provides the ability to validate the multiple layers of the cloud stack, network service, network function, orchestration, and infrastructure layers.

Some of the key functions that are enabled by the validation layer integration include

- VNF/PNF Test Automation – Automating the PNF/VNFs for functionality verification
- VNF Security Scans- Doing vulnerability scans on the VNF to make sure that they are compliant to the security requirements
- Infrastructure test automation – Validating the infrastructure layer to make sure that the infrastructure layer meets the SLAs for the VNFs and network services.
- Vertical stack validation – Validating the entire network service with multiple layers of verification which includes infrastructure, orchestration, and network service level tests

Assurance

Integrated Service Orchestration and Test Automation Solution can be integrated with the service assurance engine to monitor the network service and closed loop automation.

Analytics

Integrated Service Orchestration and Test Automation Solution should have an analytic engine that provides the dashboard views of the network service. This present a global view of the service validation in conjunction with the NFVI performance.

Policy

The native policy engine of Integrated Service Orchestration and Test Automation Solution can be integrated with the policy engine for specific operator policies

Workflow

Integrated Service Orchestration and Test Automation Solution can be integrated with the existing work flows associated with the onboarding and validation of specific workloads

Some of the key use cases enabled by the Integrated Service Orchestration and Test Automation Solution platform include.

- Automated VNF Onboarding and Validation: Ability to onboard the VNF using orchestrator on to the MSO cloud. Integrated Service Orchestration and Test Automation Solution also has internal traffic generators to validate the VNFs in a functional and performance context.
- Automated Network Service Validation and Onboarding: Ability to work with the end to end Service orchestrators to instantiate a network service on the operator cloud
- Vertical Stack Validation: The wholistic validation of network function or network service, in conjunction with the infrastructure and orchestration layers to verify the integrity of the cloud stack.
- NFVI Benchmarking: Ability to monitor the performance of the NFVI to support the VNF SLAs.

1.4. New Operational Context: NetOps – DevOps for the network

The integrated orchestration and test automation will be configured for continuous deployment and delivery to accelerate the service agility. The solution can be integrated with software release repository of any VNF vendor or it could be integrated to a central repository of the MSO where the software images are stored.

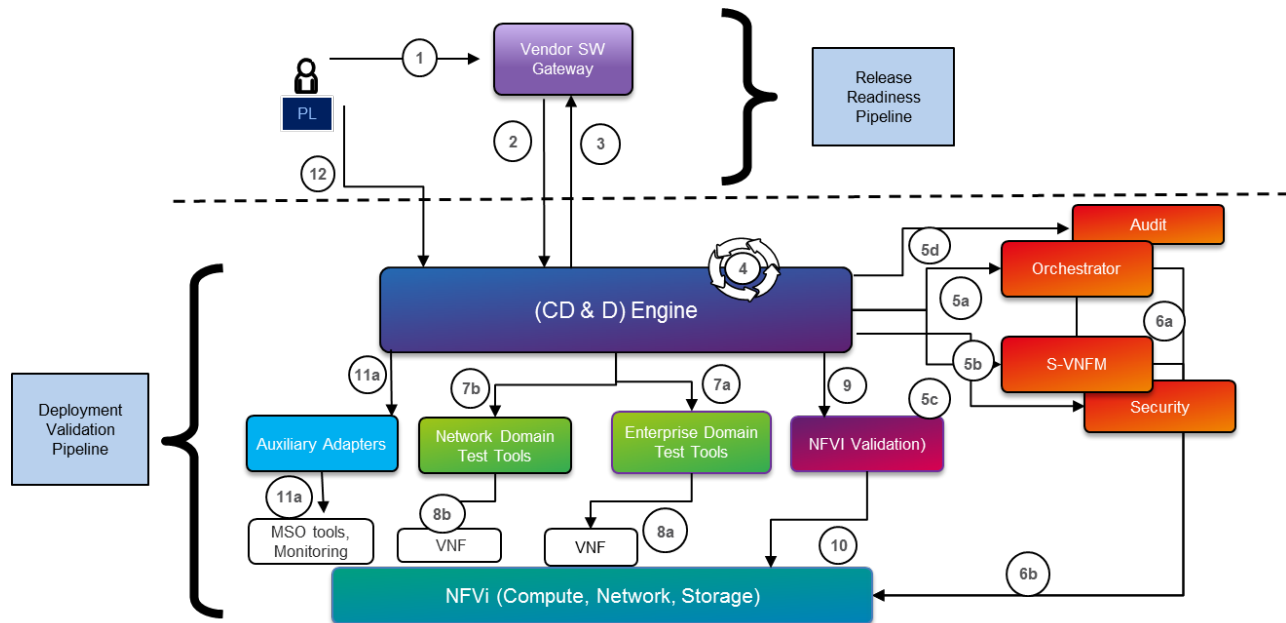


Figure 4 :Continuous Deployment and Delivery Facilitated by the Solution

The interactions in the figure above are as following:

Table 1: Interaction work flow for the Continuous Deployment and Delivery

Interface Number	Description of the interface
1	VNF Vendor development unit releases the software for verification to the Software Gateway
2, 3	SW Gateway notifies the Integrated Service Orchestration and Test Automation Solution platform and the Solution downloads the VNF package
4	Integrated Service Orchestration and Test Automation Solution uploads (manually) the VNF package using the web portal.
5a, 5b	Integrated Service Orchestration and Test Automation Solution selects the orchestrator for onboarding the VNFs.
5c, 5d	The solution can be integrated with the MSO audit engine or security policies if required. (Optional)
6a, 6b	Orchestrator acts as G-VNFM and works with S-VNFM to instantiate the VNFs on the NFVI. (This is an Optional Step)
7a, 7b	Integrated Service Orchestration and Test Automation Solution triggers the validation cycle by triggering the test engines to validate the VNFs
8a, 8b	Test Engines generate the traffic to test the VNFs under test. VNF validation reports are provided to Integrated Service Orchestration and Test Automation Solution from the test engines
9	Integrated Service Orchestration and Test Automation Solution platform can trigger the OpenStack level and infrastructure level tests using specific ONFV tests and some specific tests that are designed for the NFVI performance
10	Test Automation Solution triggers the NFVI tests for the OpenStack and infrastructure tests. Test Validation reports are provided to Integrated Service Orchestration and Test Automation Solution
11a, 11b	Develop adapters/work flow engine for MSO processes like ticket tracking systems, monitoring systems. Service assurance engines etc.
12	Vendor uploads a VNF patch to the Integrated Service Orchestration and Test Automation Solution. Integrated Service Orchestration and Test Automation Solution runs through the onboarding and validation cycle as mentioned above.

1.5. Key Benefits to the MSO from the approach

Ericsson has significant experience deploying orchestration solutions including Integrated Service Orchestration and Test Automation Solution in multi-vendor environment in the operators. Some of the key business benefits that the Integrated Service Orchestration and Test Automation Solution provide are the following

- The Solution has been able to accelerate the service velocity by 30 to 40% for enterprise service deployment for a Tier 1 customer in North America.
- Integrated Service Orchestration and Test Automation Solution provides end to end process automation which can reduce the total cost of ownership for the MSO.
- The solution also provides the ability for the MSO to enforce the MSO standards on the VNF vendors and ensure consistency across vendors.
- It provides scalable validation infrastructure with integration to test platforms for functional and performance validation of VNF and NFVI.

- It provides the ability to enlarge the ecosystem of the VNF partners for the MSOs. Integrated Service Orchestration and Test Automation Solution provides an alternative to the traditional RFQ process by enabling VNF vendors to onboard and validate the VNFs to see if they meet the desired Service Level Agreement (SLAs) prior to any commercial agreement.

Conclusion

Service velocity has emerged as the most critical benchmark of success for an MSO. To achieve this goal, it is necessary for much more advanced collaboration to occur between the MSOs and the vendor community. This paper explored the challenges in improving service velocity and developing advanced new services. Extending the automation domain to include service instantiation, validation and life cycle management is critical for improving the service velocity.

Choosing the right partner for this complex journey is critical for the success. The wholistic approach requires a cross functional solution which includes subject matter expertise and delivery experience in OSS (Orchestration and Assurance), Cloud and Infrastructure and VNF validation domains. Ericsson has successfully partnered with operators globally to enhance service agility and achieve operational efficiency using an Integrated Orchestration and Test Automation Solution

Abbreviations

VNF	Virtual Network Function
ETSI	European Telecommunication Standards Institute
E//	Ericsson
OPNFV	Open Platform for NFV
LCM	Life Cycle Management
NFV	Network Functions Virtualization
NFVI	NFV Infrastructure
SDN	Software Defined Network
MSO	Multiple System Operator
OSS	Operations Supports Systems
BSS	Business Support Systems

Bibliography & References

- [1] Ericsson Elaboration on IHS Technology, “NFV Service Provider Strategies” 2017, 2018
- [2] ETSI GS NFV-MAN 001, Network Function Virtualization (NFV) Management and Orchestration

Enhancing Wi-Fi QoE With Targeted Approach

A Technical Paper prepared for SCTE•ISBE by

Bart Vercammen

CTO Customer Premises Equipment BU
Technicolor
Prins Boudewijnlaan 47, B-2650 Edegem, Belgium
+32 3 443 6 519
Bart.Vercammen@technicolor.com

Jos Delbar

Director Product Management Managed Services
Technicolor
Prins Boudewijnlaan 47, B-2650 Edegem, Belgium
+32 3 443 64 16
Jos.Delbar@technicolor.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Wi-Fi QoE: a stepped approach.....	4
1. A Burning Platform	4
2. Case for Customer Wi-Fi Experience Management	4
3. Steps to a whole home Managed Wi-Fi solution.....	8
3.1. Comprehensive Wi-Fi diagnostics.....	9
3.1.1. Measuring Wi-Fi QoE.....	9
3.1.2. Segmentation of the issues.....	11
3.1.3. Define the right diagnosis and its importance.....	11
3.1.4. Define the right cure.....	13
3.2. RRM (Wi-Fi network optimization)	13
3.3. CAPEX Strategies	16
3.4. Self-Installation.....	17
3.5. Intelligent Client Steering	17
Conclusion.....	19
Abbreviations	21
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 - Relative contribution of interference, radio path (coverage) and saturation issues to Wi-Fi quality of experience problems	6
Figure 2 - Band usage by Stations (Technicolor, 2018)	7
Figure 3: % of STAs connecting on 2.4GHz only	7
Figure 4 - The Wi-Fi QoE wheel : the 5 Steps for a succesfull Wi-Fi QoE	8
Figure 5 - Wi-Fi performance of Mr. Robinson's home.....	10
Figure 6 - Wi-Fi performance of the Johnson family.....	10
Figure 7 - Distribution of amount of root causes in large deployments	12
Figure 8 - An example on the effects of RRM on Wi-Fi QoE: Pamela's home	14
Figure 9 - Quantitative improvement of Wi-Fi QoE by means of RRM.....	15
Figure 10 - The impact of RRM on optimal Wi-Fi across several deployments.....	15
Figure 11 - % of installed base benefitting from an extender due to coverage issues	16
Figure 12: Instable roaming behaviour	19
Figure 13 - How the different steps help in reducing the Wi-Fi QoE	20

Introduction

A subscriber's Wi-Fi quality of experience (QoE) has been an industry focus for many years. On the CPE side, the approach has long been to increase the MIMO configurations with each generation as well as to maximize the power budget by using better components, to provide maximum coverage at adequate speeds from a single access point location.

Now, as the industry pivots towards access point disaggregation, by adding Wi-Fi extenders in multiple locations in the home, we must find a way to deploy these additional devices in CAPEX and OPEX efficient ways. Typically, multiple access points immediately trigger a quest for the best roaming solution.

However, Technicolor believes this is a misplaced area of focus, firstly because an extender is often not a solution to the problem, on the contrary. It is imperative to first proactively determine which households will actually benefit from an extender, allowing a Service Provider (SP) to better plan its CAPEX investments. Secondly, these roaming solutions will commoditize over time, as standardization increases and client devices (STAs)¹ implement better roaming algorithms and standardization increases.

Technicolor recommends a five-step plan to allow a Service Provider to further enhance a subscriber's Wi-Fi QoE, in CAPEX and OPEX strategic and efficient ways:

- 1) Diagnostics: Create Wi-Fi QoE visibility into your installed base.
- 2) RRM (Radio Resource Management): Optimize the installed base. Technicolor has field data from millions of devices that shows how effective RRM/SON (Self Optimizing Networks) activity can substantially reduce OPEX cost thanks to increased QoE.
- 3) CAPEX strategies: The ability to target specific older gateways models whose replacement would significantly improve subscribers' QoE, as well as proactively identifying subscribers who can benefit from extenders.
- 4) Self-installation: Deploying solutions that are zero touch, combined with the right security, are essential to drive mass market adoption without the need for professional installation.
- 5) Client steering: Finally, for the percentage of subscribers requiring multiple access points, effective steering solutions provide additional capacity and provide the final boost to a yet higher QoE.

¹ A Wi-Fi client is also called an STA (station). We will use the word STA consistently when referring to Wi-Fi clients. However we will keep on using the word client steering as this is a commonly used term in the industry.

Wi-Fi QoE: a stepped approach

1. A Burning Platform

In today's consumer's world, with ever increasing content sources and streaming services, with a rising number of consumer devices that need to be supported in the home, Service Providers around the world are facing more and more challenges to deliver the optimal performance of the in-home network that satisfies subscribers' user experience in the home.

With Wi-Fi being the most predominant – head and shoulders above alternatives – distribution medium in the home, there's an increased focus on how to achieve the most optimal Wi-Fi performance in the home.

Since several years, for Service Providers, Wi-Fi quality of experience has been a problem. What makes it worse is that there is no uniform way for Service Providers to uncover and tackle the issues stemming from poor Wi-Fi QoE. More than often these problems get funneled through the helpdesk as “I have no Internet” or “My Internet is slow”. In reality the root causes can vary from Wi-Fi interference coming from Wi-Fi and non-Wi-Fi sources, through poor Wi-Fi coverage, to saturation of the Wi-Fi medium or any combination of the above. In addition, Service Providers are facing difficulties to diagnose those Wi-Fi issues at their subscriber's residence.

But, recently, retail Wi-Fi is now driving a wedge between the Service Provider and their subscribers, threatening to disintermediate the subscriber's broadband access from the Wi-Fi network. This imposes the Service Provider's to lose...

- Customer intimacy
- Brand recognition
- Revenue opportunities

... while subscribers still expect the Service Provider to provide increasing bandwidth (CAPEX investment) and technical support (OPEX cost). Together this might result in the Service Provider losing control of the Connected Home.

Technicolor, however, believes that Service Providers still have a window of opportunity to gain control of the Wi-Fi performance in the home by rolling out a Managed Wi-Fi solution consisting of:

- Wi-Fi insights and diagnostics
- Wi-Fi network optimization (channel planning)
- Better/more Wi-Fi access points
- Easy Wi-Fi set up and configuration
- Wi-Fi client mobility (active steering)

As we will show in the next sections these elements help our Service Provider customers to achieve dramatic improvements in customer experience.

2. Case for Customer Wi-Fi Experience Management

Consumers have, over the last few years, progressively and frustratingly, suffered from a lack of full home and high-quality Wi-Fi coverage from their Service Provider. Service Providers have been tackling this challenge head-on, rolling out a steady progression of routers whose Wi-Fi performance exceeds that

of the generation prior. It is fair to say that the industry is asymptotically approaching the maximum coverage and quality that can be provided from a single device location, and now attention starts to turn to the next wave, which is adding one or more additional Access Point devices – Wi-Fi Range Extenders – to the home.

The beautiful coincidence is that each additional Wi-Fi Extender is more likely to be in a “public” area of the home. This allows a Service Provider to blend a need – to further enhance the delivery of its broadband service via Wi-Fi – with an opportunity – to insert a device in the living areas of the home which creates an emotional link with the end user. Next, to the increasing competitive threat from retail “Mesh”-systems, this generates the trigger for Service Providers to invest in extenders and “Mesh”.

Is the Wi-Fi Extender the next Holy Grail? Yes and no. Yes, because we all know that Wi-Fi coverage is an issue that is present in the market and extenders provide a solution to that. Even more, when the average access and download speeds are increasing gradually, the need to have that same speed delivered anywhere in the home increases as well. And that is going to increase the need for extenders. So, having extenders as part of the solution is good.

However, from our real-world and large-scale collection of Wi-Fi quality of experience data (see Figure 1), we have learned that 41% of Wi-Fi QoE problems are caused by Wi-Fi interference issues, and 49% are caused by coverage issues; leaving 10% are caused by saturation issues (See paragraph 3.1.2 for an explanation on these categories). This means that 51% (= 41% + 10%) of Wi-Fi problems will not be solved by bringing in Wi-Fi Range Extenders. In fact, for these cases, Wi-Fi Range Extenders will only make matters worse. See paragraph 3.1.3 for an explanation on why this is the case.

Hence deploying Wi-Fi Range Extenders to the wrong part of the subscriber base will not only make matters worse but will also be a wasteful allocation of CAPEX. Deploying Wi-Fi Range Extenders to the right part of the subscriber base, and deploying the minimum necessary number of them, is the golden ticket.




Figure 1 - Relative contribution of interference, radio path (coverage) and saturation issues to Wi-Fi quality of experience problems²

It is crucially important for the Service Provider to know whether it is necessary or appropriate to install a new Wi-Fi Range Extender in the home, or whether such action will simply aggravate the issues. The key is knowing which subscriber suffers from which problems and how to solve them.

Is the subscriber suffering from interference from neighboring networks (so called "hidden nodes")? A hidden node is a problem that is very common in Wi-Fi and the root cause of many problems typical of denser areas. The problem arises when two neighboring access points are using the same Wi-Fi channel, but are not visible to each other. A STA (station) that is somewhere in between can however see both networks. Since both networks are on the same channel, the station starts experiencing issues because both access points are happily sending traffic at the same time. Adding Wi-Fi Range Extenders will only make it worse, by creating additional interference. Instead, deploying macro Wi-Fi network optimization to coerce the existing Wi-Fi routers to select clearer channels is a better solution. Additionally, deploying intelligent client band steering to coerce Wi-Fi STAs to use the clearer 5GHz frequency as a solution. Let's take an example to illustrate this. From the deployment mentioned (Technicolor, 2018) we see that nearly 60% of all STAs are only connecting to 2.4GHz. This of course does not tell the full story as only 38%³ of the STAs are dual band capable⁴. But it indicates that 5GHz is not used ubiquitously.

² (Technicolor, 2018)

³ (Technicolor, 2018)

⁴ There are two comments to make here. First of all not all STAs support 5GHz, some mainly older ones might only support 2.4GHz. Secondly this deployment is still using 2 different SSIDs so this shows that people often only fill in their 2.4Gh credentials and not the 5GHz ones.




Figure 2 - Band usage by Stations (Technicolor, 2018)

This becomes even more visible when analyzing the stations that are dual band capable. When we analyse this population (Technicolor, 2018)⁵ in Figure 3 it is clearly visible that 42% of those clients is not connecting to 5GHz even though they are capable to do so. The reason for this is two-fold. The main root cause is that in this deployment (Technicolor, 2018) the 2.4GHz and 5GHz bands have different SSIDs which indicates consumers don't bother filling in the 5GHz SSID and hence connect only to 2.4GHz. Another reason is that certain phones prefer being on 2.4GHz whether due to stickiness, environment conditions... This clearly demonstrates the need to unify the Wi-Fi experience in the home by combining single SSIDs with the correct band steering techniques.




Figure 3: % of STAs connecting on 2.4GHz only

⁵ We only analyzed dual band capable STAs that were connected to dual band APs

Is the main issue indeed stemming from a lack of coverage? Of course, deploying Wi-Fi Range Extenders will alleviate this problem. But which homes lack coverage? How many extra Wi-Fi Range Extenders are needed to provide the optimal level of coverage? This is not only important for efficient Wi-Fi network performance, but also for corresponding Wi-Fi extender CAPEX efficiency.

By taking a step back, and first gaining more insight into your Wi-Fi installed base, performance metrics, and root causes of Wi-Fi issues, you will be able to make better decisions on how and where to invest CAPEX. By deploying a good RRM-SON solution, you will be able to dramatically decrease the amount of Wi-Fi issues, without deploying additional hardware. This stepped approach to a whole home Managed Wi-Fi ecosystem to manage the Wi-Fi experience is the essence of the success of Service Providers in the battle for “who owns the SSID?”, and the SSID must be owned to be able to capture many of the applications and services that are on the horizon.

3. Steps to a whole home Managed Wi-Fi solution

For the Service Provider, delivering Wi-Fi as part of a broadband service is a double-edged sword. On one hand, the promise of superior Wi-Fi performance is a key selling argument for operators and a top reason to buy for subscribers. On the other hand, day-to-day Wi-Fi issues are a major source of frustration and Wi-Fi consistently appears in the top of the ranking for helpdesk calls. If operators do not succeed in managing their Wi-Fi service in a better way, subscribers will eventually turn to alternative solutions offered by the direct competition or even by the retail players. Along with alienating the subscriber, opportunities for marketing additional value-added services beyond broadband are lost.

Continuously improving the performance and the user experience of Wi-Fi products and services should be a firm objective. Service Providers need to understand that being successful requires taking a holistic approach that goes far beyond the deployment of additional Wi-Fi access points. Equally, if not more important, are the software, the user interfaces and the deeper integration with network systems. All of these elements come together in the five steps to deliver a successful Managed Wi-Fi solution. This is what Technicolor refers to as the Wi-Fi QoE wheel.



Figure 4 - The Wi-Fi QoE wheel : the 5 Steps for a successful Wi-Fi QoE

Obviously, each individual step executed in the right order is key to ensure the maximum Wi-Fi QoE for consumers. Most RRM-SON solutions go directly to step 2 of the wheel – network optimization – and most mesh solutions typically focus on step 5 – client steering, but we wish to stress the importance of getting the full spectrum of Wi-Fi use cases right. If the basics are not covered (i.e. primary access point management, first time extender installation, getting devices properly connected to Wi-Fi, discriminating between subscribers who need mesh and who do not ...) then most subscribers will still perceive a poor Wi-Fi user experience no matter how advanced other aspects of the solution are. A comprehensive Wi-Fi QoE solution addresses the full set of use cases.

3.1. Comprehensive Wi-Fi diagnostics

The foundation of a successful managed Wi-Fi solution that delivers a great QoE is insight and this insight stems from data. The appropriate type and amount of data is needed to identify what is right and wrong with the Wi-Fi QoE in the home and to determine the root cause of the issues. Only then can the right course of action be determined to help improve the Wi-Fi.

In the following paragraphs we are going to use some examples from real life use cases. In order to better understand the figures and how they are measured, let's assume you have a monitoring system that is able to indicate and measure what the full theoretical Wi-Fi performance is of the STA taking into account the capabilities (Wi-Fi standard, MIMO) between AP and STA, for each STA. In each figure, you will typically see a normalized view of the Wi-Fi performance over time for a device. In the X-axis you have the time scale, on the Y-scale you have the %-scale (0-100%). 100% means the full theoretical performance that the device can have. Every vertical line indicates a time where that device was connected, and the colors indicate where and how performance was used/available/lost. Blue indicates lost performance due to physics (too far from AP), Red means the same but due to interference, Yellow means performance lost due to other Wi-Fi traffic in your network. Dark green means used capacity (real bandwidth used by that device) and light green means available capacity. This should help in reading and understanding the examples better that follow in the text.

3.1.1. Measuring Wi-Fi QoE

The first thing is of course to measure the Wi-Fi QoE correctly. This might seem obvious, but the reality is often different. Let's take as an example the home of Mr. Robinson⁶, who is using his tablet frequently. As one can immediately see from Figure 5, Mr. Robinson is using his tablet quite far from his AP and has a lot of issues due to physics⁷ if you look at the amount of blue present in the picture. Physics typically results in a bad RSSI that results in a low modulation speed on the Wi-Fi medium and as a consequence a low throughput. Hence looking at this from a pure speed perspective, which is common practice in diagnosing Wi-Fi, the immediate conclusion is that this person needs immediate attention and care, and even that this customer should be offered an extra access point to fix his issues.

However, from a Wi-Fi QoE perspective the story is much different. In reality, Mr. Robinson is perfectly happy with his Wi-Fi at home, as he is only using it for casual surfing and mailing. Mr. Robinson is not a heavy service intense user and was quite surprised when he was contacted to “solve” his Wi-Fi issues. Although an extender would greatly improve his Wi-Fi performance, from a QoE perspective, there is currently no need for it. Once his user profile changes through using Netflix e.g. the story may change. But today, Mr. Robinson is a content customer, he is not likely to churn due to Wi-Fi issues, he is

⁶ Fictitious name used for this example

⁷ This is e.g. typically very visible through a bad RSSI (Received Signal Strength Indication)

unlikely to call the helpdesk, and likely has no appetite for buying an extender or a managed Wi-Fi service.




Figure 5 - Wi-Fi performance of Mr. Robinson's home

A similar example can be taken by looking at the home of the Johnson family in Figure 6. They like watching Netflix movies from their PlayStation and consume quite a lot of Wi-Fi bandwidth. This is visible through the large dark green lines in the left of the figure. These lines and the corresponding bandwidth consumption is much higher than Mr. Robinson from above. The bandwidth capacity they have however is also larger than Mr. Robinson (visible through the light green). When analyzing both Figure 5 and Figure 6 one can quickly see that the total bandwidth as the sum of both the used (dark green color) and the available bandwidth (light green color) in both homes shows a big difference. A technician that looks at this data would conclude that the Wi-Fi performance of the Johnson family is more than adequate (around 70-80% of the theoretical bandwidth during the majority of time) and, if you compare them with the Robinson's home, they would not get attention as their Wi-Fi performance looks good compared to the Robinson's. The Johnson family is however not doing as great as we would expect from their sheer Wi-Fi performance. At 10:45PM their AP switched to channel 11⁸ and suddenly a lot of interference⁹ starts to appear. Although the Wi-Fi performance they achieve is still good, the capacity they need is more than what is available in the system with the interference being present. The result is that the movies watched on Netflix start to get impacted and the Johnson's suffer from pixelization or serial buffering. Contrary to what may be expected, this family is not happy with their Wi-Fi, and is more likely to place an unhappy call to the helpdesk or even churn.




Figure 6 - Wi-Fi performance of the Johnson family

⁸ A typical action done by AP by means of their ACS (Automatic Channel Selection) function.

⁹ The red bars on the graph represent the interference

These two examples show that it is essential for every managed Wi-Fi solution to correctly measure the Wi-Fi QoE before anything else can be done. Looking just at effective speed is not enough, looking at range is not enough and taking a “in the moment” snapshot of performance is also not enough.

3.1.2. Segmentation of the issues

When measuring the Wi-Fi QoE correctly and accurately over time, the consequence is that it allows the Service Provider to distinguish where to focus their attention. First and foremost, it allows the Service Provider to filter the false positives from the operational flows. From various deployments, Technicolor measures that roughly around 50%-70% of a population¹⁰ has no Wi-Fi performance issues at all at any given moment in time. Typically, with a lot of deployments various “No Internet” issues are impossible to diagnose and, in some cases, assumed to be Wi-Fi related with a consequence of long investigations, box-swaps or even truck rolls. This can all be avoided with a good, time-based Wi-Fi QoE monitoring solution. When doing diagnostics through continuous monitoring of Wi-Fi QoE, a helpdesk operator can look at summary information of that subscriber when a call comes in. An automatic verification could then alert the operator that the Wi-Fi for that subscriber shows no deficiencies. As such with one glimpse of an eye a helpdesk operator can determine that Wi-Fi performance issues have nothing to do with the call and can focus their attention on other problems.

Throughout this paper we always talk about Wi-Fi QoE as heavily linked to the performance that is achieved with the devices that are being used, at the times where they are being used. However not all Wi-Fi issues seen by customers are always performance related. Typing wrong Wi-Fi passwords when connecting for the first time to a network is an often-recurring issue, despite various attempts in the industry to make this as zero-touch as possible¹¹. From recent deployments we can see as much as 7% of all households (Technicolor, 2018) suffer from failed connection attempts due to wrong passwords¹². These items have a corresponding negative effect on Wi-Fi QoE, but they require a more dedicated treatment which we will not tackle here in detail. The exception is the 4th step of the Wi-Fi QoE, which is heavily linked to this, and we discuss this in paragraph 3.4.

Measuring correctly the Wi-Fi QoE also provides a good indication of the severity within a certain population. This can be very helpful for Service Providers who wish to proactively approach and fix the most affected end-users thereby making a good judgment on who to tackle first.

3.1.3. Define the right diagnosis and its importance

As soon as a Wi-Fi QoE problem has been determined and segmented for further investigation, the most crucial step becomes determining the right diagnosis of the Wi-Fi QoE issue as well as its importance.

The root cause for Wi-Fi QoE issues can vary and fall into 3 kinds of domains:

1. **Physics issues:** A typical physics issue is a lack of coverage (too far away from the AP) or an environment that is degrading the Wi-Fi signal heavily (concrete walls, metal reinforcements...). This is the issue that most people are thinking of when it comes to Wi-Fi QoE issues
2. **Saturation issues:** Another issue that is commonplace, is a Wi-Fi QoE that is heavily degraded due to a complete overutilization of the medium. This can be due to one STA completely filling the Wi-Fi network with P2P downloads or the sum of the bandwidth of all STAs in the network that is completely filling the medium. Although this is a severe issue in certain specific

¹⁰ Data taken from (Technicolor, 2018)

¹¹ Items like QR codes, WPS (Wi-Fi Protected Setup) ...

¹² Data taken from (Technicolor, 2018)

environments, it remains a relatively minor group of issues compared to the others as can be seen in Figure 1

3. **Interference issues:** This is the most underestimated category of issues, not only because of its contribution to the overall amount of issues (see Figure 1), but even more so due to its complexity to correctly diagnose and identify the root cause. Interference can have many forms, ranging from a simple RF interferer in a certain Wi-Fi band, too far away Wi-Fi networks generating collisions as they don't see each other¹³. The location of the interferer being closer to the AP or closer to the STA (Station) also matters when determining the root cause.

A Wi-Fi QoE problem can have multiple reasons that are widespread. Having the diagnosis wrong can lead to massive discrepancies in how to tackle the issue. It is obvious that a coverage problem requires a different approach than an interference issue.

Let's use the example of a hidden node (as we explained in paragraph 2) to illustrate this. A hidden node is created because the AP and the STA see a different list of nearby APs and due to the fact the STA sees another AP on the same channel than the one he is connected to this gives collisions. This is seen as an interference issue. By deploying extenders you risk to create hidden nodes or make things only worse. Because when you install an extender this extender has to follow the channel as indicated by the main AP¹⁴ and due to the function of an extender you make that Wi-Fi network larger. Hence the probability that a STA and an AP see things differently and risk collisions is only getting bigger. That is why using extenders should only be done in clear cases where physics are the root cause of the issues.

When we all think of Wi-Fi, we typically reason around one category of issues as mentioned above. The reality is often much more complex. With the lessons learned from a number of large deployments (Technicolor, 2018) we know that in 92% of the cases (see Figure 7 the problems that are being diagnosed for a given home are not linked to just one of these three specific categories of issues. In only 8% of the cases the Wi-Fi QoE issues can be reduced to a single root cause.




Figure 7 - Distribution of amount of root causes in large deployments

¹³ So called hidden nodes

¹⁴ Assuming the hidden node appears on the band to be repeated. When using different bands on the backbone link and the link to the STA the situation can be different

This is one of the most important findings we as Technicolor have seen in our field experience. There is no such thing as one solution fits all. There is no such thing as a single solution to a problem. This only happens in lab environments and to give easy and clear examples. In real life it happens but rarely. A structured and methodological approach needs to be followed.

When the problem involves more than one kind of issue, it becomes imperative to determine the relative contribution of each category to the overall Wi-Fi QoE. This allows the Service Provider to focus his attention on these issues that contribute most to the degradation of the user experience. If that information is not available, one risk is a focus on fixing issues, with only marginal gain and at an enormous cost. Seeing the amount of multi-category issues, the structuring of priorities becomes critical

3.1.4. Define the right cure

Diagnosing the Wi-Fi QoE is the hardest and most critical part in order to improve the Wi-Fi QoE. What remains is determining and applying the right cure to the problem. If we go back to the different categories of issues, the spectrum of possible cures is wider than one might think.

Physics issues are not only solved by placing extenders in the home but, as we will see in paragraph 3.3, moving AP's to optimal locations and/or defining the suitable AP model for each subscriber. This can assist in solving a lot of issues.

Interference issues will get solved by selecting the right bands and channels to operate in as an AP. But the situation is more complex than that and a possible solution could also be to let an end-user activate its 5GHz profile on its devices to work around certain interference issues

Solving Saturation issues requires a bit more attention. In most cases, adapting the bands and channels is also the right solution. However, when there is already a plethora of devices in the air, the right solution can be to spread devices around to the available bands (load balancing). Another solution for that problem could be to prioritize certain devices over others when the overall available bandwidth remains the bottleneck.

Before diving into the fix, it is equally important to identify what those cures are, and which ones can be tackled proactively without end user action or consent. Therefore, the next step in enhancing the Wi-Fi QoE is attacking those issues that can be dealt with remotely and even proactively.

3.2. RRM (Wi-Fi network optimization)

Now that all Wi-Fi QoE issues have been diagnosed and the cure(s) have been identified, it is also time to start tackling those issues that can improve the Wi-Fi QoE automatically without end-user notification or involvement. Such automatic improvements help massively in working proactively with customers. Instead of waiting for the call or the identification of issues at the end user, working proactively in the background to improve his Wi-Fi QoE is a great mechanism to avoid churn and increase NPS.

Such proactive techniques typically fall in the domain of RRM. Through AI techniques and advanced analytics, the best channel/band is determined for the AP's in your home. This optimization of the radio network in the home aids in reducing many of the QoE issues in the interference category. Let's look at another example: the home of Pamela on Figure 8. Until shortly before 7PM the AP in Pamela's home was suffering quite badly from interference. This interference has eaten up almost all of her capacity. This combined with the fact that Pamela's laptop is not really close to the AP, leaves almost no data left to

consume. Pamela has not yet noticed this issue as she is not using any bandwidth¹⁵. The Managed Wi-Fi solution has identified this potential problem, has diagnosed it as an interference issue, and now the RRM function is stepping in to proactively fix it. Just before 7PM the RRM function switches the channel of the AP to channel 13¹⁶ and as you see from the graph the interference disappears¹⁷. Suddenly Pamela enjoys more than enough capacity for her needs. This action has been performed just in time as the fix occurs shortly before she starts consuming bandwidth. Without the RRM intervention this would have caused Wi-Fi QoE issues and would have generated a helpdesk call and frustration.




Figure 8 - An example on the effects of RRM on Wi-Fi QoE: Pamela's home¹⁸

The fact that these actions can be performed anytime without any need for end user interaction allows for continuous and proactive optimization of the Wi-Fi QoE when it comes to interference related issues. This optimization generates a better and more consistent end user experience without any hassle for the consumer. This optimization can be done both on an individual home basis and can also work with perfect effectiveness for complete MDUs (Multi-dwelling units) and neighborhoods that are under control of the RRM function. More than their retail competitors, a Service Provider often has control over a large part of the visible AP's. Where a Service Provider controls a large number of APs in one neighborhood, MDU, or campus, they have a clear advantage in improving the Wi-Fi QoE over retail solutions which are more scattered and unlikely to be able to play the wide area RRM game.

The reduction on Wi-Fi QoE issues can be significant and have a huge impact on both helpdesk call rates as well as customer satisfaction. If we do an analysis on how much such an RRM function brings in real life, let's go back to the data from (Technicolor, 2018). In that installed base, there were two groups of people monitored¹⁹. For one group there was merely monitoring and no intervention (we called that the control group) and in the other group we started intervening after 1 week (the proactive group). The first week we only monitored both groups to demonstrate that both populations were equal in behavior. In Figure 9 one can see for every day monitored the amount of Wi-Fi QoE issues in the Y-axis²⁰. A higher bar in the graph means there have been more QoE issues in that time period than during a period with a lower bar. After one week, the RRM function started to optimize the channels for the proactive group. As one can see from the graph this drastically improved the Wi-Fi QoE, in the order of 20%-30% improvement²¹. If we translate this into OPEX savings this resulted in a reduction of 30% in first line helpdesk calls (Technicolor, 2018).

¹⁵ There is no dark green visible in the graph that is representing the consumption of traffic

¹⁶ For avoidance of doubt: channel 13 is taken from an example in EU where channel 13 is allowed. In NAM this is not allowed

¹⁷ The physics issue remains and this is logical as nothing is changing to the position of the laptop

¹⁸ For explanation of the colors and bars, see Figure 5

¹⁹ Each of them around 50% of the population and statistically relevant (more than 0.5M households)

²⁰ We normalized the Y-axis to 100% to clearly indicate the delta's

²¹ In (Technicolor, 2018) we witnessed an improvement of 24%




Figure 9 - Quantitative improvement of Wi-Fi QoE by means of RRM

Another way of demonstrating the value can be seen in Figure 10, where we have aggregated the amount of homes where the Wi-Fi QoE is optimal and suboptimal. This figure shows a comparison of the optimal QoE across several deployed ISPs and with clear indication whether RRM is activated or not. The results show a clear effect of enabling RRM across the several ISPs. The amount of homes with an optimal Wi-Fi QoE is in the range of 20% higher with RRM enabled.




Figure 10 - The impact of RRM on optimal Wi-Fi across several deployments

With the arrival of extenders and mesh systems, the importance of an RRM function will only increase. The explosion of managed and unmanaged AP's in homes increase the usage of the Wi-Fi spectrum²² and the presence of interference issues impacting Wi-Fi QoE. So, the numbers as shown in 8 will only increase the effect of proper RRM.

We did not cover yet the aspect of variations over time which adds another dimension in the complexity that needs to be considered. Wi-Fi issues in the interference category are in many cases intermittent in time. This means that certain interference sources appear and disappear in regular or even irregular intervals. An MDU, for example, will be largely empty during day time when people are at work, but will be very busy during the evenings and weekends. Hence the interference and disturbance that is generated by the inhabitants will show different patterns during the days vs the night. The same applies to office buildings but in the opposite time intervals. There, usage will be higher during the day but very limited in the evenings and weekends. Therefore, it is important that these RRM functions run at regular intervals as

²² Even an AP that is not consuming BW from STAs has an impact on the Wi-Fi network

the bands and channels that need to be chosen will continuously change over time. This also requires that these RRM functions monitor the Wi-Fi environment on a continuous basis, so they can intervene whenever needed.

3.3. CAPEX Strategies

While many Wi-Fi issues can be attributed to interference and to saturation, both of which can be dealt with automatically and without the deployment of additional devices, some issues come down to a lack of coverage and require manual intervention: the deployment of better and/or more Wi-Fi access points.

The majority of the readers will immediately think of rolling out extenders to address this, however the reality is often different. First, an extender should not be sold/given to any/every subscriber. A massive advertising campaign about the new model of extender is not going to contribute to a substantially improved Wi-Fi QoE. As we demonstrated in paragraph 3.1.3, end users suffering from mainly or only interference and saturation will not be helped by installing an extender. On the contrary, their Wi-Fi QoE will get worse. If we look at (Technicolor, 2018) we can distinguish that only 33% of the installed base is benefiting from an extender. This is visible in Figure 11 where we looked at the full installed base and looked at the % of households where coverage issues could be solved by adding an extender.




Figure 11 - % of installed base benefitting from an extender due to coverage issues

Second, it is better to first look at the existing installed base of AP's. Every year, Service Providers are struggling to identify how much CAPEX should be spent in replacing their legacy base of gateways/AP's. This is very costly, and the benefits are typically unquantifiable/unknown. With a clear root cause analysis of the Wi-Fi QoE issues, it can be identified which impacted subscribers will actually benefit from replacing their AP by a more recent and better performing one. This will lead to CAPEX savings due to a better replacement strategy. It also avoids unnecessary CAPEX that is spent on extenders that might not be needed anymore, once the main gateway is swapped out. In (Technicolor, 2018) we managed to prove that only 17% of the older installed based needed to be proactively swapped out to avoid Wi-Fi QoE issues. This leads to a saving of more than 80% on a budget to swap gateways in homes.

And finally, this stepped approach allows the Service Provider to have a more targeted and successful approach in deploying extenders. It not only saves CAPEX by supplying them only to consumers who have a need for them, but it also allows it to work in a more proactive way to eliminate churn. In current deployments, Service Providers have no other option than to push extenders to all subscribers through advertising and retail shops or wait until the customer calls the helpdesk with problems, or churns. With advanced Wi-Fi QoE analytics, Service Providers can proactively utilize a targeted approach that will decrease CAPEX and increase customer satisfaction.

3.4. Self-Installation

A step in the process that is often seen as obvious is the self-installation of Wi-Fi extenders. Reality however is different. Once shipping managed Wi-Fi systems to customers, either by sending extra AP's or even by sending a complete kit, the manner of install is, in many cases, not trivial.

In the majority of active deployments today Service Providers still use multiple SSID names and passwords for the different Wi-Fi bands on the main AP. When moving to a full managed Wi-Fi system there also needs to be a migration to a single SSID deployment. Albeit technically quite easy to do, in practice it is more difficult to execute. Customers might have installed devices on both bands or only one and you cannot just remove one of the SSIDs without informing the end user or risking that certain STAs don't connect anymore²³

Next to that, a good self-install starts by giving subscribers a 'plug and play' experience that leverages smart industrial design and frictionless user interfaces. Onboarding procedures need to be designed to be simple and clear, so that even subscribers who are not technologically savvy will succeed. For the more "connected" subscribers, one could offer an intuitive smartphone application, and this could be complemented with voice control AI. Common scenarios for daily use (e.g. setting up a guest network) need to be made easy. Finally, all devices must be able to be managed remotely by the Service Provider.

3.5. Intelligent Client Steering

The last step in optimizing the Wi-Fi QoE is done by roaming STAs from one AP or band to another when certain criteria are met. In the industry, this is often the first topic that is interrogated, so why do we put this as the last topic in our quest for optimizing the Wi-Fi QoE?

The reason is quite straightforward and has already been mentioned a couple of times. Steering only brings benefits to certain type of issues, namely interference and saturation.

1. **Interference:** As we discussed already in paragraph 3.1.3, interference can have many root causes. Client steering can help in reducing immediate interference issues, but it can't completely solve them. When moving from one access point to another one (or to another band) the client steering mechanism relies on the fact that every AP has chosen its optimal channel. If not, you might jump or get steered to another AP that has similar, or even worse, issues. That is why it is imperative that this steering happens only after channel optimizations have been done through the RRM function. Client steering can indeed help to overcome sudden jumps of interference but then only for a little while till the RRM optimizes the overall situation again and this depends on his frequency and scalability

²³ In the case that a customer only had installed a 2,4GHz SSID on his phone and the 5GHz SSID is the one that was maintained after the migration, the station will not connect anymore to Wi-Fi as the STA does not have the right SSID and password configured

2. **Saturation:** Client Steering will help in addressing issues that are completely saturated by acting in two ways. First of all, it will spread STAs around the various access points and, even more importantly, around the various bands. This is called load balancing. Second, it can also take action in extreme cases by prioritizing certain traffic (eg. Video) over normal traffic in order to guarantee a certain QoE

In summary, client steering can bring further marginal improvement to the Wi-Fi QoE after all other, more significant optimizations have been done, to deal with more instantaneous and intermittent interference problems. It can also deal with saturation issues through load balancing and through prioritization of certain traffic types and use cases. This what we also call Intelligent Steering and is the cherry on the top of the cake. The important thing to note here is that this steering needs to be done to improve the Wi-Fi QoE proactively (and even reactively) and that this does not require any end user intervention. Especially that last part is essential. Due to various implementations on client devices such as smartphones, an intelligent steering solution needs to take care in how/when to do this with certain types of devices. Due to the STAs' own intelligent implementations, they often have their own preferences of where they want to connect²⁴. When not being careful, this can lead to a user experience that is heavily impacted by trying to steer too frequently or quickly.

Let's illustrate this with a real-life example as depicted in Figure 12. Assume a home environment with two access points installed together with a Wi-Fi roaming solution that determines which AP is best for the STA to be connected to and that will ensure the STA is roamed to the correct AP when needed. The first AP (called AP1) is on channel 6 and the second AP is on channel 1²⁵. As an STA we have taken an iPhone with IOS11 for this test. We have deliberately chosen this device as this is publicly known²⁶ to have an intelligent behavior when it comes to selecting the Wi-Fi networks it wants to be connected to.

In the starting position for the test the phone is connected to AP1 without any issues and with a more than decent RSSI. The user experiences no issues whatsoever. At a certain point in time (Step 2) there is an RF interferer arriving on the Wi-Fi channel 6 of AP1. This can typically be a hidden node as explained before which will be invisible to AP1 and will not trigger an ACS rescan²⁷. A good roaming solution will detect this disturber and take immediate action to roam the STA away from AP1. Seen the fact that AP2 is in range of AP2 it will decide to roam the phone to AP2 (Step 3). The phone follows and connects to AP2 however the link is a bit worse in terms of RSSI because the phone is further away from AP2. But this is still better than being connected to the disturbed AP1. However, the phone does not know anything about this disturber on AP1 and only sees that AP1 has a better RSSI than AP2. Because of its implementation²⁸ the phone will disconnect from AP2 as the RSSI is below the threshold and will reconnect to AP1. The roaming solution will react and move the phone back again to AP2 (Step 4).

²⁴ See (Chowdhry, 2017)

²⁵ For simplicity in explaining we assume to be using the 2.4GHz frequency band but the conclusions are equally valid on 5GHz or a mixture of both

²⁶ See (Apple, 2017)

²⁷ ACS rescan: the automatic channel selection of the AP when triggered will scan the environment and select an other and better Wi-Fi channel

²⁸ See (Apple, 2017)




Figure 12 - Instable roaming behaviour

The implementation of the roaming solution will determine the reaction speed to roam the phone back to AP2 as well as the mechanism of roaming. The end-user experience is heavily dependent on it. The faster and more frequently the roaming solution will intervene the more interrupts the end-user will notice. If the mechanism is restricted to 802.11v this behavior will go on forever. The result is even worse when the roaming solution uses MAC ACL mechanisms²⁹ to deauthenticate the phone from the AP. In that case the phone will even decide that it is so bad and will never connect autonomously again to any of the two APs and instead revert back to the cellular network.

This test clearly demonstrates the unwanted effects of a roaming solution that does not take into account the specific implementations of STAs which are becoming more and more intelligent. Not doing so can result in a very bad user experience.

Conclusion

In this paper, we demonstrated that a 5-step approach³⁰ to proactively and reactively tackle Wi-Fi QoE dramatically improves the situation. If we look back at the distribution of categories we introduced in Figure 1, each of our 5 steps contribute, in a coordinated way, to reducing Wi-Fi QoE issues as depicted in Figure 13. This figure illustrates when the necessary root causes have been identified (Step 1), proactive channel planning both in the home and in the neighborhood, work on the interference category. CAPEX strategies mainly have an impact on radio path issues and client steering can diminish the remainder of Wi-Fi QoE issues in all categories.

²⁹ MAC ACL is well known in Wi-Fi and is a access control list containing the MAC addresses of devices that are allowed to connect to the AP and which ones are blacklisted and can't connect to the AP. When using a MAC ACL mechanism to roam this means that the roaming solution will throw off the STA from the AP and then put the MAC address of the STA on the blacklist of the MAC ACL hence prohibiting the STA to connect back to the AP.

³⁰ See Figure 4

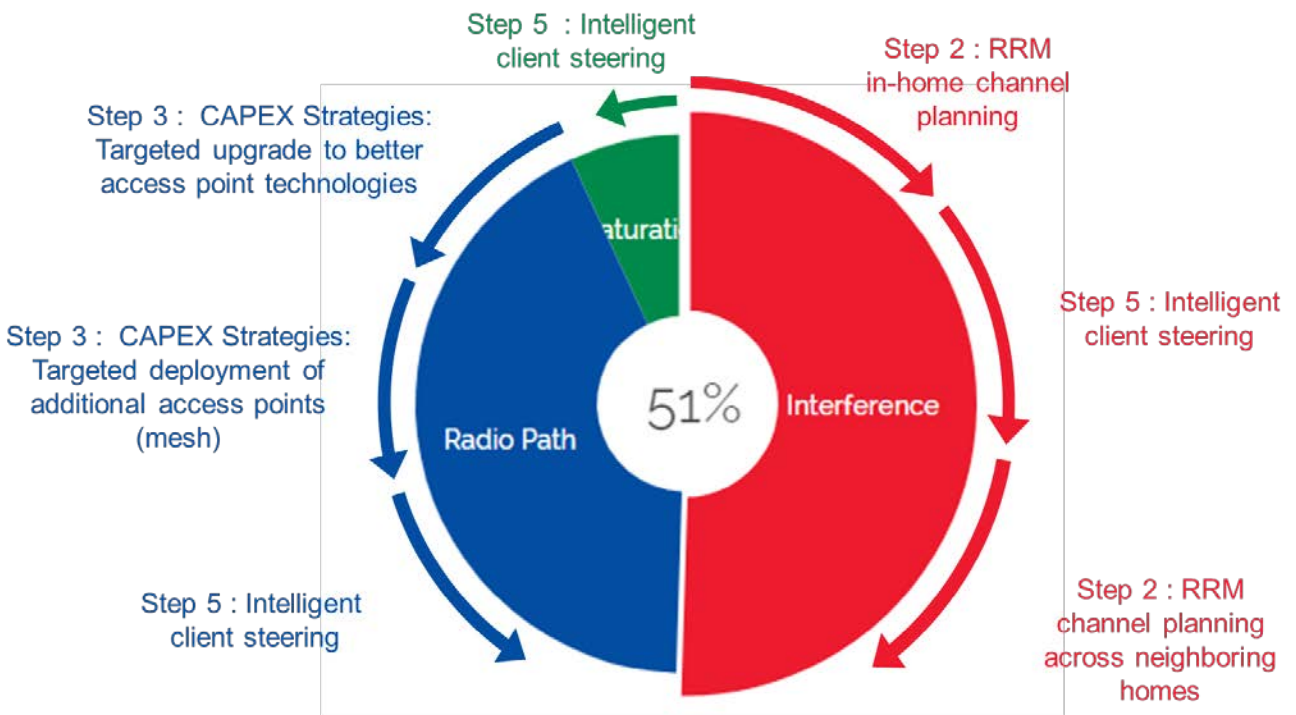


Figure 13 - How the different steps help in reducing the Wi-Fi QoE

The impact on CAPEX and OPEX expenditures are massive. Experience in deploying this Wi-Fi QoE wheel demonstrates that savings can be made in a variety of areas (Technicolor, 2018)

- Reduction in % of first calls: up to 30%
- Reduction in % of repeat calls: up to 40%
- Reduction in the amount of unidentified Wi-Fi issues: up to 90%
- Increase in the amount of first call resolution
- Reduction in the call duration
- Reduction in the amount of truck rolls
- Reduction in the amount of CAPEX spent in box swaps
- Reduction in the amount of CAPEX spent on extenders

The most important thing that is increased is the amount of customer satisfaction and the NPS (Net Promoter Score).

Abbreviations

AP	access point
SP	Service Provider
CAPEX	Capital Expense
OPEX	Operational Expense
SSID	Service Set Identifier – another name for the name of your Wi-Fi network
RRM-SON	Radio Resource Management – Self Optimizing Networks
QoE	Quality of Experience
RSSI	Received Signal Strength Indication
ACS	Automatic Channel Selection
BW	Bandwidth
MDU	Multi-Dwelling Unit
NPS	Net Promoter Score
STA	Station
MAC ACL	Medium Access Control Access Control List
ISP	Internet Service Provider

Bibliography & References

Apple. (2017, November 8). *macOS wireless roaming for enterprise customers*. Retrieved from Apple Support: <https://support.apple.com/en-us/HT206207>

Chowdhry, A. (2017, July 26). *Apple iOS 11 Is Going To Fix A Pesky Wi-Fi Issue*. Retrieved from Forbes: <https://www.forbes.com/sites/amitchowdhry/2017/07/26/wi-fi-auto-join/#7aae030e73cf>

Technicolor. (2018). Data taken from a 2.5M live deployment of households.

Evolving The “Box”: The Smart Set-Top Box

A Technical Paper prepared for SCTE•ISBE by

David Goodwin

Director of Product Management
ARRIS

101 Tournament Drive, Horsham, PA 19044
+1 215 323 2215
david.goodwin@arris.com

Charles Cheevers, CTO, ARRIS

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
The Primary Job & Evolving the “Box”	4
Anatomy of the Smart STB	7
Audio / Video Processing	8
CPU / GPU	9
Wi-Fi 10	
IoT - Bluetooth / Zigbee / RF4CE.....	14
Microphone / Speaker	15
1. Broadside Approach.....	16
2. Endfire Approach.....	17
3. Independent Component Analysis (ICA) Approach	18
4. DUET Approach	19
Sensors	21
LTE / 5G	22
Form Factor and Placement.....	22
Services.....	24
Benefits to the Service Provider.....	25
Conclusion.....	25
Abbreviations	26
Bibliography & References.....	27

List of Figures

Title	Page Number
Figure 1 – Video Playback (the Primary Job)	4
Figure 2 – Where STBs are Located	5
Figure 3 – Time Spent in Locations	6
Figure 4 – Evolve the Box	7
Figure 5 – Audio/Video Processing.....	8
Figure 6 – CPU/GPU.....	9
Figure 7 – Wi-Fi.....	10
Figure 8 – Bandwidth and RSSI.....	11
Figure 9 – STB Location and RSSI.....	11
Figure 10 – STB / Extender Location and RSSI	12
Figure 11 – STB Location / RSSI	12
Figure 12 – IoT	14

Figure 13 – Broadside Approach	17
Figure 14 – Endfire Approach	18
Figure 15 – ICA Approach.....	19
Figure 16 – DUET approach	20
Figure 17 – Sensors in a Smart STB	21
Figure 18 – LTE / 5G in a Smart STB	22
Figure 19 – Form Factor and Placement Examples	23
Figure 20 – Smart STB solutions	23
Figure 21 – Smart STB services	24

List of Tables

Title	Page Number
Table 1 – Wi-Fi Configurations.....	13

Introduction

The set-top box (STB) has long been the cornerstone of video decoding TV services for the service provider. However, now it is being recognized and marketed as a device that does a great deal more. This paper explores the value of the device that is in multiple rooms, is the gateway to the pixels on the screens, and offers the scope to be leveraged as an Internet of Things (IoT) Hub and a smart assistant – 4 for 1 (or more) functionality in a single box.

This paper offers readers options and challenges for leveraging the STB for video, IoT, natural language-based voice services and more. The paper discusses the technologies of voice detection in a noisy environment, out of phase surround sound, and the challenge of making the smart STB support IoT, microphone, and speaker technology from lower end devices all the way up to high end offerings with high quality speaker solutions. How will service providers maintain control of the pixels and continue to provide value to their subscribers?

The Primary Job & Evolving the “Box”

The STB has had the primary job of helping the subscriber navigate and play back video content that is offered by the service provider.



Figure 1 – Video Playback (the Primary Job)

The STB fills the following traditional roles:

- Acts as a consistent / known end point for the provider’s video delivery network
- Provides a UX for accessing video services
- Connects to the Television (TV) / home theater

The STB performs that job securely and reliably to assure a level of service expected by the service provider and to ensure all content provider / content protection requirements are met.

Over the years certainly features & connectivity have changed:

- Analog to Digital tuners and CableCARD to Internet Protocol (IP) delivery
- Standard-definition to High Definition (HD) and now Ultra High Definition (UHD) / 4K
- Single tuner HD boxes to multi-tuner Digital Video Recorders (DVRs) to IP clients & cloud DVR
- Linear broadcast to On-demand to IP video apps such as Netflix and YouTube
- Analog audio/video outputs to High-Definition Multimedia Interface (HDMI)

However, the primary job of the STB has remained mostly the same.

Meanwhile many new consumer devices & new experiences have emerged, and the home is filling up with many single-purpose devices with varying capabilities & non-integrated user experiences:

- Smart assistant devices
- Wi-Fi access points & extenders
- Soundbars and small Bluetooth speakers
- 4K and HDR capable TV displays
- IoT hubs and devices, cameras and connected ‘smart’ sensors

How does the STB continue to provide value to the subscriber and to the service provider with this multitude of gadgets and technology?

The STB is a unique device – it is in multiple rooms and where consumers spend a lot of time.

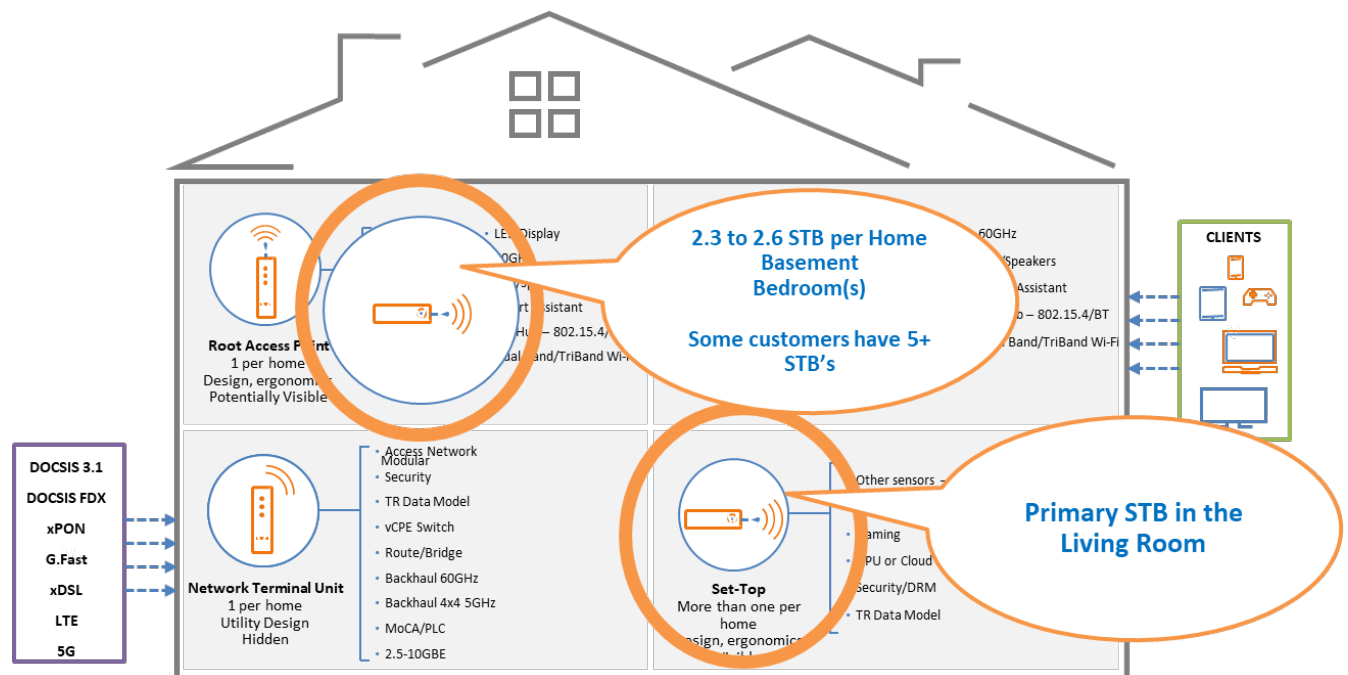


Figure 2 – Where STBs are Located

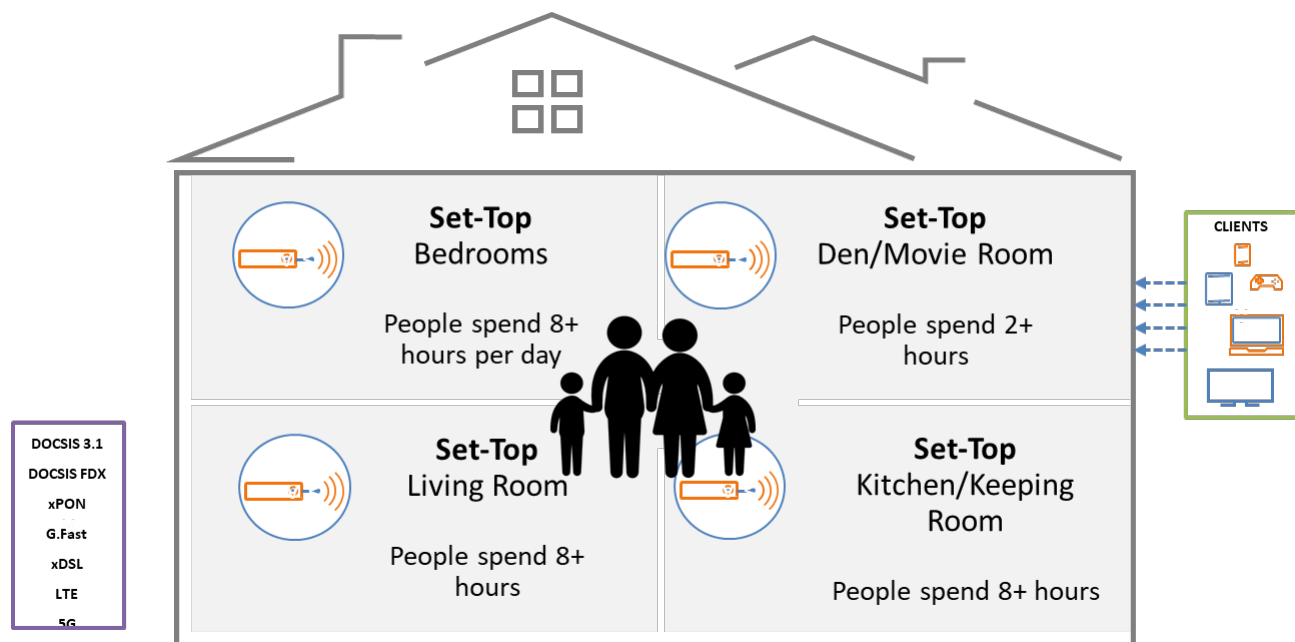


Figure 3 – Time Spent in Locations

It is hard to get a consumer device located in multiple rooms – the consumer must value its purpose in each room separately.

Given that advantage of the STB, what more can we do with the STB device and its functions?

- People prefer less devices ----- Avoid device clutter
- People prefer simple than complex ---- Ease of use
- People like using the large screen for communal and family things. Small screens are used for private activities
- The “Large Screen” in the home is evolving to more of a home control panel with visual and audio outputs
- Consumers care about Household monitoring functions like Carbon Dioxide (CO) or Radon – but maybe haven’t considered this safety feature integrated into another device
- Consumers show a growing interest in audio throughout the home in multiple rooms for multi-room listening




Figure 4 – Evolve the Box

We must ask: **How should the STB evolve?**

Think of the “box” as a connected node...a **smart STB**:

- Network connected
- Located in multiple rooms & where consumers spend time
- Connected to the “Large Screen” and audio output
- Capable of additional interfaces & sensors
- Interactive

By leveraging this ability, you can evolve the traditional STB to a smart STB!

Anatomy of the Smart STB

The anatomy of the smart STB includes all the core elements that are evolving and can be integrated in new products. We will explore each of these in further sections.

- Audio/Video Processing
- Central Processing Unit (CPU) / Graphics Processing Unit (GPU)
- Wi-Fi
- Bluetooth / Zigbee
- Microphone / Speaker
- Sensors
- Long-term Evolution (LTE) / 5G
- Form factor & placement

Audio / Video Processing

Compelling A/V experiences & more codec options for the provider to leverage:

- 4Kp60 – p120 High Frame Rate
- High Dynamic Range (HDR)
- High-efficiency Video Coding (HEVC), AV1
- More immersive audio
- 8K

The cornerstone of the device is providing the highest quality video rendering to the display and audio output. The value and likely usage of the device is significantly reduced if the consumer gets ‘trained’ to want the perceived better video quality directly from their Smart TV apps or other connected device.




Figure 5 – Audio/Video Processing

Newer and emerging video technologies can benefit both the subscriber via improving video quality but can also benefit the service provider via more efficient use of bandwidth and adapting to needs of various content providers. New video and audio codecs may be required for new services or can help reduce cost of licensing and implementation.

Challenges:

- Most of these advancements require newer / emerging silicon
- TV / Home theater interoperability issues will arise
- Consumer confusion over nomenclature and support
 - Example: Many consumers know they have a “4K” TV they purchased over the past several years. Newer TVs will claim support for HDR and/or wide color gamut, etc. but TVs can range from 400 to 2,000+ nits of brightness and varying video processing capability.
- Content availability / delivery

The end-to-end delivery including content production, delivery, smart STB, HDMI cable and TV, and even any HDMI repeater or an Audio/Video Receiver (A/VR) home theater system, must all work together to ensure the best experience. A smart STB must provide premier A/V capabilities, adapt to the consumer’s connected TV and audio system, and provide meaningful information to the service provider and consumer to identify any interoperability issues.

CPU / GPU



Figure 6 – CPU/GPU

The CPU and GPU of the smart STB help keep the consumer engaged with:

- User Experience (UX) - fundamental
- Content Navigation with high-quality artwork - fundamental
- 360 / Immersive video – new and untapped
- Gaming - untapped
- Virtual Reality (VR) / Augmented Reality (AR) – new and potential parallel experiences

Silicon in the STB space is evolving to lower power (watts) but dual and quad core CPUs providing 15-20K+ Dhrystone MIPS (DMIPS) of processing power for applications. 2D/3D GPUs and overall system memory bandwidth and improved video display pipelines can support higher-end graphics to keep up with the 4K HDR video experience.

Some of the challenges:

- Potentially higher cost as the very latest silicon is made available
- Need to identify value-add experiences and focus on those
- Accessory/headset compatibility
- Leveraging the remote control (or even voice) to navigate
- Ecosystem (content, apps, partnerships)

It is natural to consider dedicated gaming consoles and high-end gaming PCs as competitors in this area. However, the smart STB has the advantage of being in multiple rooms, in a small and quiet form factor

and well-integrated with the service provider experience for video content. The smart STB could be a client to streaming gaming services or in-home streaming from consoles and high-end PCs.

Wi-Fi




Figure 7 – Wi-Fi

Wi-Fi in the smart STB can be used for video delivery from a gateway, but can also be designed as Wi-Fi Extender (STB can be wired, or wireless backhaul)

- Flexible installation – not tied to coax outlet locations
- Located in a room likely to have other Wi-Fi clients
- Sense other Wi-Fi devices in the room
- Provide health of the Wi-Fi in that part of the home
- Emerging 802.11ax for latest Wi-Fi
- 802.11ad 60 GHz for highest same room performance

Challenges:

- Extender adds cost & power and impacts size of the STB
- Backhaul connectivity & performance




Figure 8 – Bandwidth and RSSI

It is not recommended to go beyond -65 dB RSSI (Received Signal Strength Indicator) when deploying a Wi-Fi STB to ensure sufficient performance for video stream delivery.

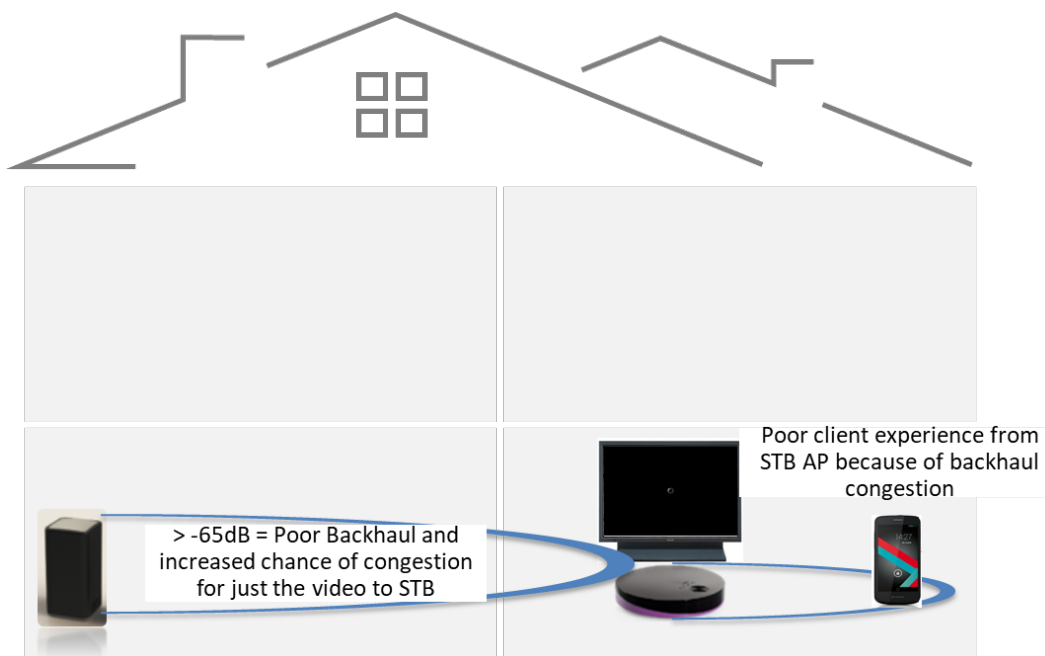


Figure 9 – STB Location and RSSI

STBs go where the RSSI may be > -65 dB typically which can still allow for video delivery but may result in backhaul congestion for extender use case.

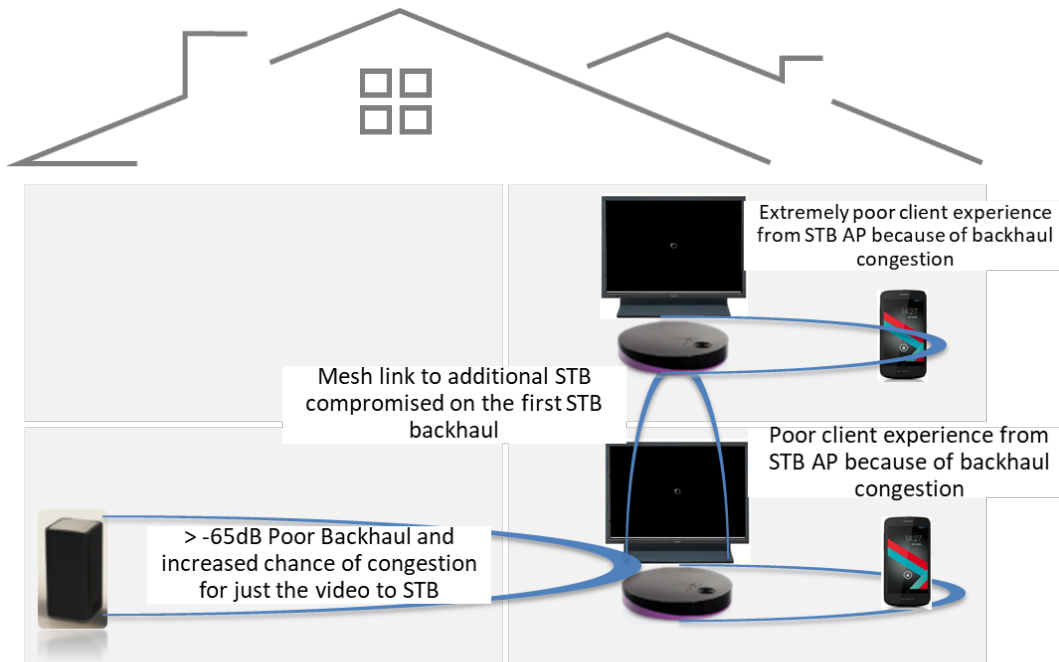


Figure 10 – STB / Extender Location and RSSI

Poor Wi-Fi client performance / experience is likely if a mesh link between 2 STBs acting as extenders is used due to backhaul congestion. Even the prioritization for the video streaming may suffer and be inadequate.




Figure 11 – STB Location / RSSI

Due to TV placement, the STBs end up in remote locations in the home driving RSSI concerns.

Figure 11 shows a large dataset of ~280K 4K 4x4 Wi-Fi STBs.

- Majority > 50 dB path loss from primary AP
- At -65 dB it is recommended to add a Wi-Fi extender even to support primary 4K video feed @ 25 Mbps

The STB can be a secondary extension device once primary extender strategy has been deployed.

Another consideration/challenge is that the STB can also be turned off. STBs need to sleep to drive energy saving metrics (~1W).

The form factors of the STB can also influence using the device as a Wi-Fi Access Point (AP). For example, a larger soundbar enclosure can allow for additional space for antenna diversity, but placement within the enclosure can also face challenges in certain directions due to the speakers and other components that may be in the way.

A small form factor encourages poor placement for Wi-Fi such as behind a TV, or on the wall.

A very small dongle device may only allow 2x2 solutions with very minimal antenna diversity and also encourages poor placement for Wi-Fi.

Table 1 – Wi-Fi Configurations

Single-band 5G	Lowest \$	Smallest size	Focus on video
DBS 2.4 / 5G	Mid \$	Mid size	Avoid 2.4 for video
DBC 2.4 / 5G	Highest \$	Largest size	Avoid 2.4 for video

There are several considerations for selecting a Wi-Fi radio and antenna configuration for a video + extender scenario.

A wired back-haul may be considered, but interoperability on both ends is required and may not be feasible (e.g., MoCA in the gateway + set-top). A Tri-radio (dedicated 5 GHz) backhaul that is found in various solutions in the marketplace may be considered but adds cost and size to the STB.

- A single band 5 GHz radio can suffice if a separate 2.4 GHz Extender architecture exists; this can still provide value to push 5 GHz reliably throughout the home
- Dual-band switched (DBS) is not ideal as video should not be on 2.4 GHz
- Dual-band concurrent (DBC)
 - Use 5 GHz + QoS for video delivery
 - For wireless repeating a 2x2 or 4x4 halves the bandwidth (4x4 preferred)
 - Using 2.4 GHz as the backhaul protects the 5 GHz video delivery but may not meet performance expectations

IoT - Bluetooth / Zigbee / RF4CE



Figure 12 – IoT

Numerous experiences can be supported with the (re)use of existing STB radios being used for remote control support. In particular a Bluetooth LE remote can be used for IoT device connectivity as well. Addition of low power ZigBee / Z-Wave radios are also feasible.

- Wireless remote (no line of sight)
- IoT hub, Security & Presence detect
- Audio streaming (in/out)
- Push to talk voice control of Smart Assistant - near field mic in remote

Wireless remotes are already deployed in millions of STBs – far superior to IR traditional remotes and support additional data/connectivity (such as voice in remote) capabilities already deploying with service providers.

For IoT Hub / Security the smart STB can connect to devices to extend coverage as part of the mesh network and can display status on TV and provide other interactions and controls (such as controlling lighting or turning off cameras).

Presence is a capability long mentioned in the STB space but not well-integrated today. The standard is the profile selection via remote control used with OTT services today. Why not take advantage of the smart STB Bluetooth or Wi-Fi connection that can detect personal device, namely smartphones that are rarely far away from the individual consumer?

Audio streaming, when described to consumers, is an under-appreciated feature. Imagine using your wireless Bluetooth headphones or Bluetooth connected audio speaker system with your smart STB.

- Privacy – listen to your favorite program at night without disturbing your family, or be able to listen to the news while your family chaos spins around you
- Accessibility – have a dialog enhancement version or descriptive video service track sent privately to your headset while rest of the family can listen on the TV
- Second language – send a different language track to the headset vs. the TV
- Voice chat – gaming headset style or for standard Voice over IP (VOIP) phone conversation via your smart STB

Challenges:

- Interoperability / Pairing & Connection, (Ease of use)
 - Any connection (Bluetooth or Zigbee, etc.) to a device needs to be automated where possible or UX assistance may be needed to guide the consumer
- 2.4 GHz coexistence between Bluetooth, Wi-Fi and other 2.4 GHz radios must be considered
- RF4CE only STBs may not be dual-stack and therefore may only support remote control use cases and may not be capable of also communicating to Zigbee devices

Microphone / Speaker

Far-field voice and optional speaker or soundbar integration provide new interactivity

- No remote needed
- Smart assistant
- Voice calling
- TV audio / soundbar

Adding a microphone to the STB has challenges due to determining the best location and number of microphones. Placement behind a TV, in an enclosed shelf or oriented vertically on-wall may lead to insufficient performance for voice input. However, they are small and may not materially affect size of STB.

Adding a speaker to the STB can substantially affect the size and the power consumption of the device. The speaker(s) can be sized and designed just for basic tone and voice feedback or provide full-range output for music in varying performance levels. Specific design experience for acoustic tuning is crucial.

Speakers in the STB can be omitted in favor of using the built-in TV speakers or home theater, but voice assistant use cases that require feedback would need the TV / audio device to be on. Using HDMI Consumer Electronics Control (CEC) or other capability to turn the TV on to provide feedback is likely not a desirable experience / not expected by the consumer.

2-4 watts of speaker amplification for mass market smaller devices may suffice. A premium audio experience can be delivered in a soundbar form factor that is integrated with the STB. A soundbar with 20 watts+ of amplification provides both an improved audio experience and a more consistent install performance for far-field voice & speakers (vs. small STB placed behind TV). The consumer is much more likely to perceive higher value from the soundbar.

Another benefit to including far-field voice is that consumers no longer need to worry about a lost remote or changing the remote's batteries.

Adding microphones, voice Digital Signal Processing (DSP) chips (if not integrated into the STB chip), optionally speakers along with all the mechanical considerations for mounting and isolating the components certainly adds cost to the STB.

Of course, audio privacy is critical in this growing area of consumer adoption. The need for a physical mute button with very clear indication of when the microphones are disabled is well understood.

One of the challenges in far-field voice that we will dive deeper into within this paper is how to detect the barge-in / key word to trigger the voice assistant functions.

When the TV is playing, and this sound is generated by the STB decoder output – then the DSP in the STB voice processing unit – can invert the sound in – to be able to effectively cancel it from the sounds in the room detected by the microphones.

When the TV uses a surround sound system – it's slightly out of phase to the original source and amplified – it is so much harder to cancel.

Similarly, for the Smart Assistant if the sound on the TV comes from an additional HDMI input source it is harder to cancel.

To support voice control & barge/keyword detection in a home environment, several techniques can be considered and balanced against cost such as DSP, mics, and memory along with needed performance for the solution.

- Beamforming
 - Broadside
 - Endfire
- Blind Source Separation (BSS)
 - Independent Component Analysis (ICA)
 - Degenerate Unmixing Estimation Technique (DUET)

1. Broadside Approach

An array of microphones is oriented perpendicular to the general direction of the command sound sources. It isolates sound sources in a room by removing sounds coming from angles relative to the orientation of the microphone array. This is done by time shifting the microphone signals and adding them together.

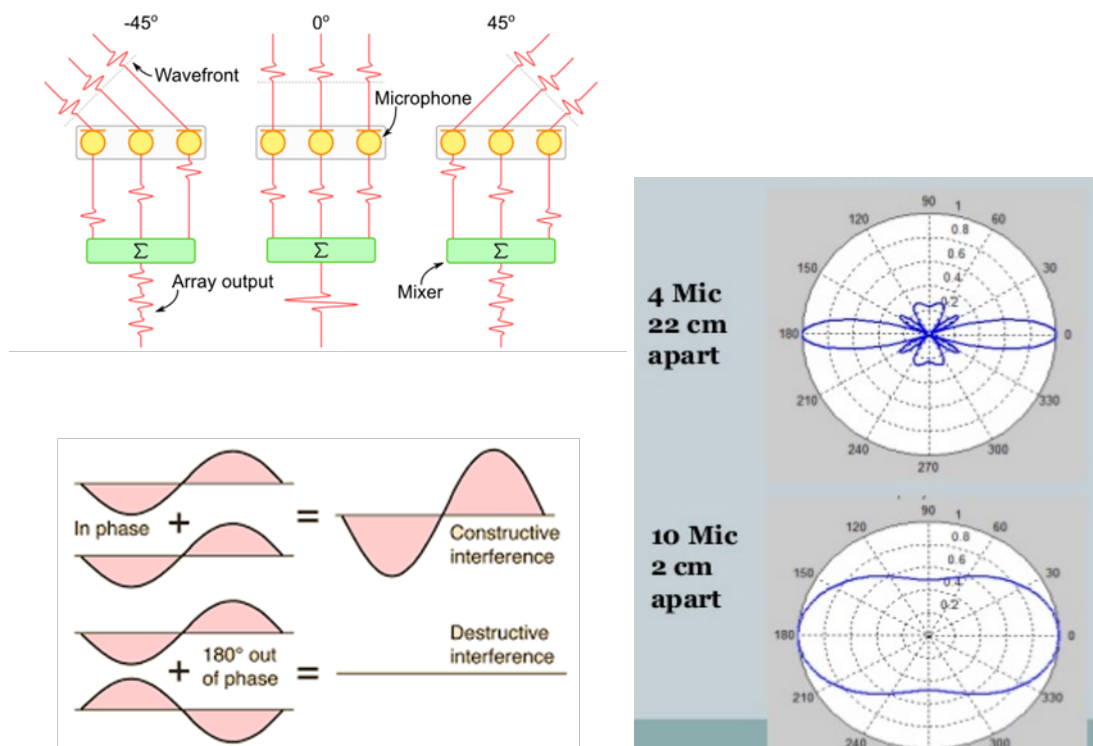


Figure 13 – Broadside Approach

Advantages

- The direction from which sounds are picked up is ‘steerable’ to maximize receptivity from that location, while minimizing sounds from other directions
- The DSP is straight forward and requires little memory

Disadvantages

- Picks up sounds equally well at 180° from the ‘steered’ direction at the same time (Example: If a TV is behind the microphone array, it will pick up the TV audio along with the command voice in front)
- 2 mics needed to ease the spatial selectivity, (narrow window angle over which sounds can be picked up)

2. Endfire Approach

An array of microphones is oriented in the general direction of the command sound sources. Sounds from behind are removed while sounds from in front are detected. This is done by time shifting the microphone signals and subtracting them. The further you place the microphones the narrower the sound receiving angle becomes, at the expense of less sound removal from behind.





Figure 14 – Endfire Approach

Advantages

- Can completely remove sound noise from almost the half of a room (e.g., remove interfering TV sounds if product is in front of the TV speakers)
- DSP is straight forward and requires little memory

Disadvantages

- Picks up sound at an angle in the vicinity of +/-90o, depending on the spacing of the microphones
- Adding microphones will not improve the spatial selectivity much

3. Independent Component Analysis (ICA) Approach

The vocal source signal amplitude varies with a Super-Gaussian distribution.

A source signal at a moment in time can be thought of as a sample in an N-dimensional grid space, where each orthogonal dimension axis is a microphone reading. A source signal sample is described by an N-element coordinate value in this space.

If the basis vectors are rotated so that one lines up along a direction of Super-Gaussian distribution, and we assume that all other sound sources are independent (orthogonal) to this direction, then only data on this axis with Super-Gaussian distribution is assumed to belong to a single voice.

The voice signal can be isolated by only taking the sample values along this axis.

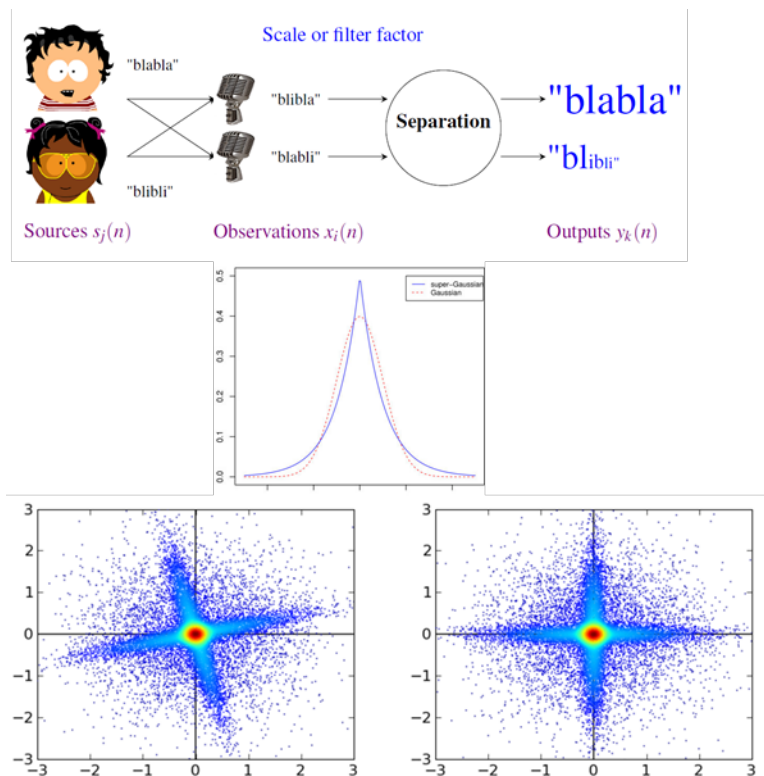


Figure 15 – ICA Approach

Advantages

- As few as 2 microphones can be used to separate two sound sources (Example: One command voice and a TV)
- The source sounds that are separated can be placed anywhere in the room as long as they do not come from the same direction of arrival
- The mic arrangement and distance is flexible, as long as their relative positions are known. This allows one to place mics away from known noise source locations, to improve the signal to noise ratio (SNR) of extracted command signals.

Disadvantages

- The number of sound sources in the room, must not exceed the number of mics used, in order to cleanly separate the sounds (Example: many mics required to isolate a command voice during a party)
- DSP is very involved and requires much memory

4. DUET Approach

A sound sources direction of arrival (DOA) is uniquely related to the delay time (phase difference) from when it reaches one microphone to the next.

Excited frequencies belonging to a single sound source will fall along a constant phase angle. Therefore, we can reconstruct a single sound source audio from just those spectral components associated with a particular source phase angle.

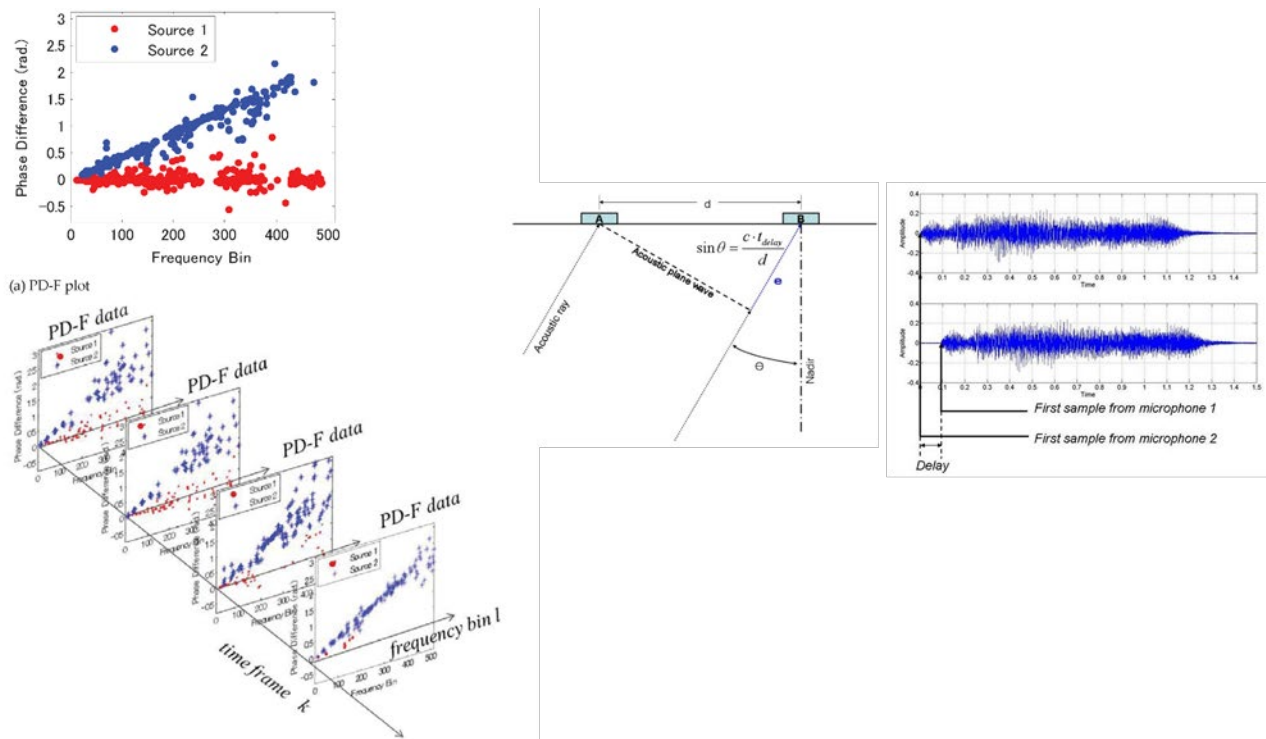


Figure 16 – DUET approach

Advantages

- As few as 2 mics are necessary to separate any number of sound sources - although the quality of sound source separation diminishes with an increase in # of sources
- The source sounds that are separated can be placed anywhere in the room as long as they do not come from the same direction of arrival

Disadvantages

- Traditional use of DUET requires that the mics be placed no more than ~2cm apart, because the source location is determined in the frequency domain (as opposed to the time domain, such as ICA)
- Undesirable artifacts can be created and may be difficult to remove because of the conversions to/from the time/frequency domains
- DSP is very involved and requires much RAM

Sensors




Figure 17 – Sensors in a Smart STB

A connected node smart STB is part of the home network and can help with the “health of the home”. A few sensors to consider include:

- Smoke / heat detector
- Radon monitor
- CO / carbon monoxide detector
- Security / glass break & presence

It is worth exploring new revenue opportunities from these sensors. Some states in the US now mandate a 90 day radon report before selling a house. Consumers may also value continuous monitoring and notifications/reports. Could this be a \$1 per month service from STB location?

A smart STB with far-field voice could possibly leverage the mic and DSP for security monitoring such as glass break, door opening, gun shot, or presence detection. False alerts could be managed by checking for known / trusted devices (Bluetooth or Wi-Fi to smartphones).

When integrating these sensors, the design must consider placement and size. A small plug-mount smart STB could contain a CO detector, but a ceiling mount smart STB may be more appropriate for smoke detection and could be combined with Wi-Fi Extender and 60 GHz wireless HDMI to the TV. There is opportunity to explore the right form factors and designs.

Challenges include any unique installation factors or setup and for certain sensors the needed safety & compliance evaluation & design must be considered. Sensor lifetimes should be evaluated to match with expected STB deployment lifecycle (e.g., if a CO sensor is rated at max 5 years) and consumers need a means to be notified of sensor status and if any replacement capability exists.

LTE / 5G




Figure 18 – LTE / 5G in a Smart STB

The smart STB can connect to the service provider network over 4G LTE and emerging 5G WAN network.

Fixed Wireless Access (FWA) allows for rapid deployment and no waiting for cable wiring to the home.

The download bandwidth should be able to sustain at least 5 Mbps for video and up to ~18 Mbps for 4K video when using FWA.

LTE modules are now available in small M.2 form factors to minimize the additional space required inside the STB. Antennas can be installed inside the plastic enclosure to avoid external antennas visible to the subscriber. CAT-4 and CAT-6 and even higher performance modules are available that exceed 150 Mbps capabilities with use of 2 or more carrier aggregation support in LTE-A.

Challenges:

- 5G networks still in early phases
- Placement in the home needs evaluation – will there be sufficient signal deep into a home
- Subscriber Identity Management (SIM):
 - A SIM slot can be included, although this requires either the subscriber installing one or the service provider pre-installing one and possibly having to deal with replacements if needed
 - An eSIM can be used, but the service provider and network need to support this technology & associated provisioning

Form Factor and Placement

The smart STB can take on many forms and be placed in different way. You may wish to make a bold statement ‘out front’ or blend / hide the device and focus on the TV UX. Features such as voice or Wi-Fi or sensor performance capabilities / cost may dictate the form and placement as well.

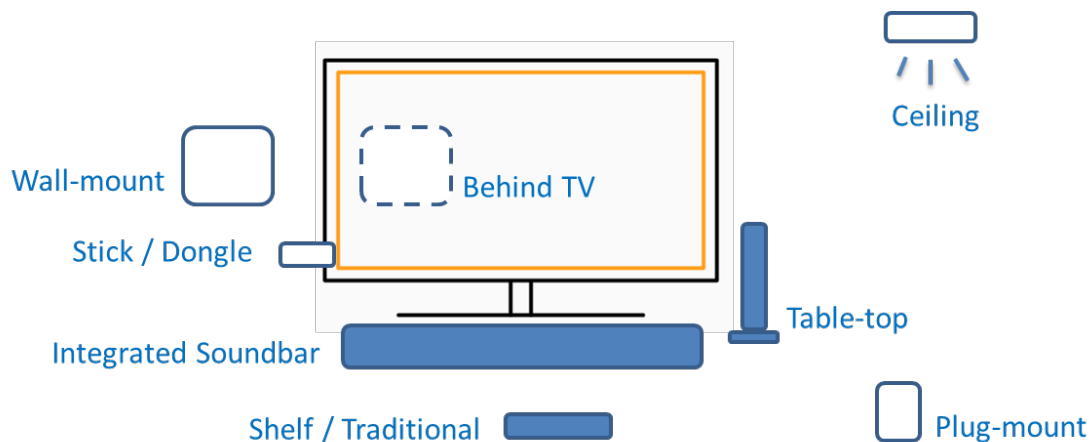


Figure 19 – Form Factor and Placement Examples

A portfolio of devices can be explored to adapt to different market and consumer needs. A lower / mid-range target for mass market may have select smart STB features and lower power audio if speakers are integrated.

Higher-end devices can push the performance barrier with latest CPU/GPU capabilities for VR and immersive experiences or can be a soundbar smart STB that complements the latest 4K HDR displays.



Figure 20 – Smart STB solutions

Consumers value the benefit in ergonomics – one device vs. 4 or 5 in a high traffic room. The smart STB that combines multiple functions can provide this value.

Services

The smart STB can provide new services on the Large Screen in the home.

- Room console – control your lighting and room environment and Wi-Fi
- Home security controls and camera viewing / notifications
- Education
- Visual UX smart assistant feedback
- Shopping
- Home / energy management
- Aging in place / health management
- Child management




Figure 21 – Smart STB services

Aging in Place:

- Health stats & notifications (time for medication) – pause TV for compliance
- Connect to sensors for blood pressure, weight, pulse

Energy Management:

- Show latest energy bill and usage on-screen
- Compare usage to last month / last year

Home Console for the room:

- Control lighting, security, sound, and connectivity (Wi-Fi health)
- Visual (UX) Smart Assistant feedback
- No need to rely just on voice feedback – the smart STB can present information and responses / actions on-screen

Education:

- Integrate online courses with an on-screen UX
- Family experience with children & reward with TV time after

Child management:

- Parental controls for TV and Wi-Fi / devices
- View reports and status

Benefits to the Service Provider

The smart STB provides differentiation and enables the ability to offer additional services & value. This could provide additional revenue opportunities and subscriber retention and higher perception & ratings of the service.

- Extend services deeper into the home (beyond the gateway/network demarcation)
- Multiple end-points throughout the home that can have new leverage and services
- Service provider ‘owns’ the largest screens in the home
- Controls the quality of connection of service
- Adds new combined Internet/Visual UX/IoT/Lifestyle Services to the TV
- Drives new opportunities like commerce on TV, Education on TV, Aging in Place on TV – all with connectivity control

Conclusion

The smart set-top box provides an integrated experience and contains multiple devices in single device:

- 4K HDR video playback
- Smart assistant / voice
- IoT hub & sensors
- Audio / soundbar

The Smart STB. The “Box” can Evolve.

Abbreviations

A/V	Audio/Video
A/VR	Audio/Video Receiver
AR	Augmented Reality
bps	bits per second
BSS	Blind source separation
CEC	Consumer Electronics Control
CO	Carbon Dioxide
CPU	Central Processing Unit
DBC	Dual-band concurrent
DBS	Dual-band switched (or selectable)
DMIPS	Dhrystone MIPS
DOCSIS	Data Over Cable Service Interface Specification
DSP	Digital Signal Processing
DUET	Degenerate Umixing Estimation Technique
DVR	Digital Video Recorder
FWA	Fixed Wireless Access
GPU	Graphics Processing Unit
HD	High-definition
HDMI	High-Definition Multimedia Interface
HDR	high dynamic range
HEVC	High-efficiency Video Coding
ICA	Independent Component Analysis
IoT	Internet of Things
IP	Internet Protocol
ISBE	International Society of Broadband Experts
LTE	Long-term Evolution
LTE-A	Aggregated Long-Term Evolution
MIPS	Machine Instructions Per Second
MoCA	Multimedia over Coax Alliance
RSSI	receive signal strength indicator
SCTE	Society of Cable Telecommunications Engineers
SIM	Subscriber Identity Management
SNR	Signal to noise ratio
STB	Set-top box
UHD	Ultra high-definition
UX	User experience
VOIP	Voice over IP
VR	Virtual Reality
xDSL	Variations of Digital Subscriber Line
xPON	Next generation Passive Optical Network

Bibliography & References

STMicro. “AN4426 Application Note: Tutorial for MEMS Microphones.” January 2014.

A. Greensted. “The Lab Book Pages: An Online Collection of Electronics Information.” (<http://www.labbookpages.co.uk/audio/beamforming.html>), November 29, 2010.

Y. Zheng. “A Broadband Adaptive Beamformer Using Nested Arrays and Multirate Techniques.” Dept. of Systems and Computer Engineering, Carleton University, Ottawa, K1S 5B6, Canada. 1998.

STMicro. “UM1967 User Manual: Getting started with osxAcousticBF real-time beam forming software expansion for STM32Cube.” March 2016.

J. Griffith, C. Jim. “An Alternative Approach to Linearly Constrained Adaptive Beamforming.” IEEE Transactions on Antennas and Propagation, Vol. AP-30. NO. 1, January 1982.

M. Puigt. “A Very Short Introduction to Blind Source Separation.” Foundation for Research and Technology – Hellas Institute of Computer Science. May 2011.

N. Bryan. “Source Separation Tutorial Mini-Series II: Introduction to Non-Negative Matrix Factorization.” DSP Seminar. Center for Computer Research in Music and Acoustics, Stanford University. April 2013.

J. Thiemann. “An Experimental Comparison of Source Separation and Beamforming Techniques for Microphone Array Signal Enhancement.” 2013 IEEE International Workshop on Machine Learning for Signal Processing, Sept. 22–25, 2013, Southampton, UK.

V. Thormundsson. “Voice as an Interface in the Smart Home: Can you hear me now?” EDN Network. February 2015.

A. Barinov. “Voice Samples Recording and Speech Quality Assessment for Forensic and Automatic Speaker Identification.” AES Convention Paper. Speech Technology Center Ltd., Saint Petersburg, Russia. November 2010.

FDX & D3.1 Capacity Scenarios

Understanding overheads and RBA configurations

A Technical Paper prepared for SCTE•ISBE by

Karthik Sundaresan

Principal Architect

CableLabs

858 Coal Creek Circle, Louisville, 80027

3036613895

k.sundaresan@cablelabs.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
FDX Overview	5
1. FDX Overview	5
1.1. FDX Frequency Plan.....	6
1.2. FDX Initialization Overview	7
2. FDX Sounding.....	9
2.1. Sounding	9
2.1. Sounding methods	9
2.2. IGs & TGs.....	10
2.3. Initial and Periodic Sounding.....	11
2.4. CWT Sounding Overhead	11
2.5. OUDP Sounding Overhead.....	12
3. FDX Echo Cancellation	13
3.1. FDX Echo Cancellation at Node	13
3.2. FDX Echo Cancellation at CM	14
3.2.1. Foreground and Background Training	14
3.2.2. Initial vs Periodic Training	14
3.3. EC Training & Overhead	14
4. FDX Resource Block Allocation	15
4.1. Fast and Slow RBA Switching.....	16
4.2. RBA Switching Use cases.....	17
4.3. RBA switching & Overhead.....	18
FDX Scenarios	21
5. FDX Capacity Basics	21
5.1. FDX Modulation orders	21
5.2. FDX TG formation	21
5.3. FDX Node, FDX CM & FDX-L CM Capacity	22
5.4. CM Channel Support.....	23
5.5. FDX MER and Capacity assumptions.....	23
6. FDX RBA configurations	24
6.1. 5 Sub-band configurations	24
6.1. Static US & DS.....	25
6.2. Duty Cycle RBAs.....	26
6.3. Dynamic RBAs and the Need for RBA Management.....	29
7. FDX + D3.1 Channel Capacity	29
7.1. FDX CM Common Use case (Use Case 1)	29
7.2. Additional OFDM channel for 1 or 2 FDX sub-band (Use case-2).....	30
7.3. Additional Optional OFDM channel Support (Use case-3)	31
7.4. FDX-L CM (Use case-4).....	31
Conclusion.....	32
Abbreviations	33
Bibliography & References.....	33

List of Figures

Title	Page Number
Figure 1 – FDX spectrum	7
Figure 2 – FDX CM Initialization	8
Figure 3 – Sounding Opportunities (a) CWT Sounding (b) OUDP Sounding	10
Figure 4 – RBA Example.....	16
Figure 5 – RBA Switch in different TGs	17
Figure 6 – RBA Switching Time Calculation	18
Figure 7 – CM RBA Switch/Channel Pause Time.....	20
Figure 8 – RBA Overhead.....	20
Figure 9 – FDX Channel Pause/Cycle Time (Overhead).....	20
Figure 10 - CM-to-CM Interference Levels and Interference Groups	22
Figure 11 – RBA SubBand Configurations & Capacity.....	24
Figure 12 – Different view of the RBA capacities.....	25
Figure 13 – Duty Cycle RBA example 96 MHz FDX Band	26
Figure 14 – Duty Cycle RBA example 192 MHz FDX Band	26
Figure 15 – Duty Cycle RBA example 288 MHz FDX Band	27
Figure 16 – Duty Cycle RBA example 384 MHz FDX Band	27
Figure 17 – Duty Cycle RBA example 576 MHz FDX Band	28
Figure 18 – FDX CM Total channel Capacity- Use Case 1	30
Figure 19 – FDX Total CM channel Capacity- Use Case 2	30
Figure 20 – FDX Total CM channel Capacity- Use Case 3	31
Figure 21 – FDX-L CM CM channel Capacity- Use Case 4	32

List of Tables

Title	Page Number
Table 1 – CWT Sounding Size.....	11
Table 2 – CWT Sounding Overhead.....	12
Table 3 – OUDP Sounding Size	13
Table 4 – OUDP Sounding Overhead.....	13
Table 5 – EC training Overhead	15
Table 6 – RBA OverHead	19
Table 7 – RxMER needed for Modulation orders.....	21
Table 8 – CM-to-CM Interference Levels and Interference Groups.....	21
Table 9 – FDX Channel Support.....	23
Table 10 – D3.1 Bit Loading and Channel Capacities	23
Table 11 – D3.1/FDX Bit Loading and Channel Capacities.....	24

Introduction

Full Duplex DOCSIS 3.1 (FDX) represents the next evolution of DOCSIS 3.1 (D3.1) technology, significantly increasing upstream capacity and enabling multi-Gbps symmetric service tiers over HFC networks.

Over the last many years, service offerings and IP capacity needs have continued to grow at a rapid pace. As a result, operators will require cost effective means of adding capacity to their HFC networks in order to provide the services that their customers will expect, eventually reaching gigabit and multi-gigabit speeds. D3.1 was designed to be a cost-effective way of meeting these performance targets, especially for typical consumption patterns where downstream consumption is much higher than upstream consumption. FDX is an evolution of D3.1 that increases the upstream capacity to similar levels as the downstream.

FDX fundamentally changes the nature of information delivery across the cable plant, and how it will be maintained and managed. FDX significantly increases the upstream capacity by enabling upstream and downstream channels to concurrently exist over the same spectrum without the need to time share the use of the spectrum. The upstream and downstream channels each fully access the same spectrum at the same time, practically doubling the use of the spectrum. Using D3.1 as a foundation, FDX accomplishes this by using a combination of interference cancellation and intelligent scheduling at the CMTS through enhancements to the existing D3.1 technology. This allows a migration path that allows operators to cost effectively migrate to FDX, while still maintaining and leveraging their existing installed base.

New FDX procedures such as sounding and echo cancellation happen at FDX initialization as well as periodically. The FDX channel may not be available for data transmission during this for some amount of time. These processes place some overhead on the raw bandwidth available. D3.1 allows for multiple OFDM/OFDMA profiles each tuned to account for plant conditions experienced by a set of CMs, using different modulation orders on the same channel. Aggregate channel capacity varies with the CMs and the profiles in use at a time. In FDX operation interference from other CMs, and Echo Cancellation training at the CM and at the FDX Node, the MER signature at the Node & CM will be different, the question is how the full duplex operation will affect the capacity of the channel and the network.

The FDX band is divided into sub-bands and the CMTS assigns sub-band(s) for upstream or downstream operation. Now this Resource Block Assignment (RBA), where a sub-band can change direction at a given time, directly impacts the available DS and US bandwidth seen by the CM in the FDX Band. Different CMs will have different bandwidth demand for both the upstream and downstream directions which can change over time, and FDX allows for the RBA to be changed dynamically to match. FDX CMs are grouped into Interference groups and Transmission groups each with a unique RBA.

This paper helps understanding how an operator can estimate US and DS capacity for FDX channels in different scenarios. The channel capacity affects how many subscribers can be assigned to use the same set of channels and affects traffic engineering and operational decisions (such as when an operator would need to split the node to increase available capacity). This paper presents a framework to understand the FDX/D3.1 downstream and upstream channel capacities across a range of MSO operational scenarios.

FDX Overview

1. FDX Overview

Full Duplex DOCSIS 3.1 (FDX) technologies significantly increases the upstream capacity by enabling upstream and downstream channels to concurrently exist over the same spectrum without the need to time share the use of the spectrum. The upstream and downstream channels each fully access the same spectrum at the same time, practically doubling the traffic-carrying capacity of the spectrum. Using D3.1 as a foundation, FDX accomplishes this by using a combination of interference avoidance, echo cancellation and intelligent scheduling at the CMTS. The evolution to an FDX network is an incremental evolution of D3.1 technology and will support both backward compatibility and coexistence with previous generations of DOCSIS technology deployments.

The FDX band will occupy a subset of the RF spectrum, 108-684 MHz. The FDX spectrum is divided into one, two, or three sub-bands. Each sub-band contains 1 OFDM downstream channel and 1 or 2 upstream OFDMA channels. From the CMTS perspective, traffic will be simultaneously flowing upstream and downstream in each sub-band. From the CM perspective, the spectrum will still be frequency division multiplexed, i.e. each CM will use a sub-band only for upstream or downstream operation for a given time. But one set of CMs can use the sub-band for upstream at the same time that a different set of CMs has been assigned to use that sub-band for downstream.

In order to transmit and receive at the same time in each sub-band, the CMTS/FDX Node will use echo cancellation techniques to separate the upstream and downstream transmissions. An FDX Node supports simultaneous upstream and downstream communications over each FDX channel enabled by echo cancellation techniques for self-interference and echo-cancellation. FDX cable modems will operate in frequency division duplexing (FDD) mode, where on any FDX sub-band, the CM is either transmitting in the upstream or receiving in the downstream. An FDX CMTS allocates FDX channels to cable modems by providing modems access to upstream and downstream channels through FDD; a CM's operation on an FDX band in either US or DS can be changed by the CMTS. FDX channels can be bonded with non-FDX channels and with other FDX channels.

Due to the lack of complete isolation between a pair of CMs, if one CM is transmitting in the upstream in one sub-band while another CM is trying to receive in that same sub-band, energy from the first CM upstream transmission can leak into the location of the second CM and prevent it from successfully receiving downstream transmissions. Simultaneous transmission and reception within the same frequencies in different interference groups introduces co-channel interference (CCI) lowering spectral efficiency.

Adjacent Leakage Interference (ALI) refers to the power that leaks from an upstream transmission of the CM into a downstream channel of the same CM in another part of the FDX spectrum. The CM has to transmit at a relatively high-power level to be received by the FDX Node, and as a result the power of the out-of-band components of this upstream transmission are comparable to the power of a downstream signal in an adjacent channel at CM input. Some of this upstream out-of-band power gets coupled into the receiver path through the coupler within the CM. Further out-of-band power gets added to the received signal through reflections in the drop cable and at the connection with the main cable. The sum of all these out-of-band components of the upstream transmission that gets added to the downstream signal is referred to as ALI.

Adjacent Channel Interference (ACI) refers to the power that remains in the same band as the transmitted signal but gets added into the receiver path through the coupler within the CM as well as through

reflections in the cable and its taps. This is significantly stronger than ALI, but it is not an in-band interference like ALI. Its main effect is in overloading the receiver circuitry. Hence precise cancellation is not needed as in the case of ALI, though some cancellation is beneficial to reduce the load on the receiver and analog-to-digital conversion circuitry. In the context of FDX, this ALI and ACI are interferences resulting from upstream transmissions of the specific receiving CM, and hence these can be categorized as self-ALI and self-ACI, respectively. The reception at a CM can also be impacted by ALI and ACI from upstream transmissions of other CMs in the cable plant, in particular, other CMs in the same IG.

To avoid the risk of co-channel interference (CCI) and adjacent channel interference (ACI) between CMs, the CMTS schedules transmissions and grants such that a CM does not transmit at the same time as other CMs that are susceptible to interference and are receiving. CM to CM interference susceptibility is measured through a sounding process. A sounding method is used to identify groups of CMs, called Interference Groups (IGs), that would interfere with each other if they were allowed to transmit and receive at the same time in a sub-band. After measuring CM to CM interference susceptibility, the CMTS creates IGs, and schedules transmissions and grants to CMs to avoid having a CM transmit when other CMs in its IG are receiving.

IGs will be grouped together into a small number of Transmission Groups (TGs). These TGs will be used to load balance the upstream and downstream traffic within each sub-band. Each TG will be given a Resource Block Assignment (RBA), which assigns the direction of traffic in each sub-band for that TG. A TG can use some sub-bands in the upstream direction while using other sub-bands in the downstream direction. While a TG can only use a sub-band in one direction at a given time, the RBA for a TG can be changed, allowing the direction of traffic for that TG in the sub-band to be changed. The CMTS coordinates the change of the RBA to ensure that the traffic in one direction is stopped before starting traffic in the opposite direction in order to prevent interference.

There is a significant difference in power levels between data transmission and reception at a given CM. Normally, diplex filters keep the upstream channel transmissions from interfering with neighboring downstream channel reception in the CM. However, FDX CMs will not have diplexers between FDX sub-bands, in order to allow the CM to efficiently change the direction of the spectrum in a sub-band. In order to prevent upstream channel transmissions from interfering with adjacent downstream channels in the CM, the CM will use echo cancellation techniques to reduce the upstream interference ACI & ALI. Because the CM does not have control over downstream transmissions, the CMTS will assist in coordinating upstream and downstream transmissions to allow the CM to train its echo canceller.

1.1. FDX Frequency Plan

The FDX Band refers to the spectrum where FDX operation can occur. Occupied FDX Band refers to the part of the FDX Band where FDX operation is occurring, this can be a subset of the FDX Band. FDX Downstream Channels are downstream channels in the Occupied FDX Band. FDX Upstream Channels are upstream channels in the Occupied FDX Band. FDX Sub-band refers to a single FDX Downstream Channel and the associated FDX Upstream Channel(s)(1 or 2) sharing the same spectrum.

The frequency range defined for FDX is 108 MHz to 684 MHz. The upper limit of 684 MHz is derived from starting with the lower band edge of mid-split (108 MHz) and allowing for three OFDM channels at 192 MHz

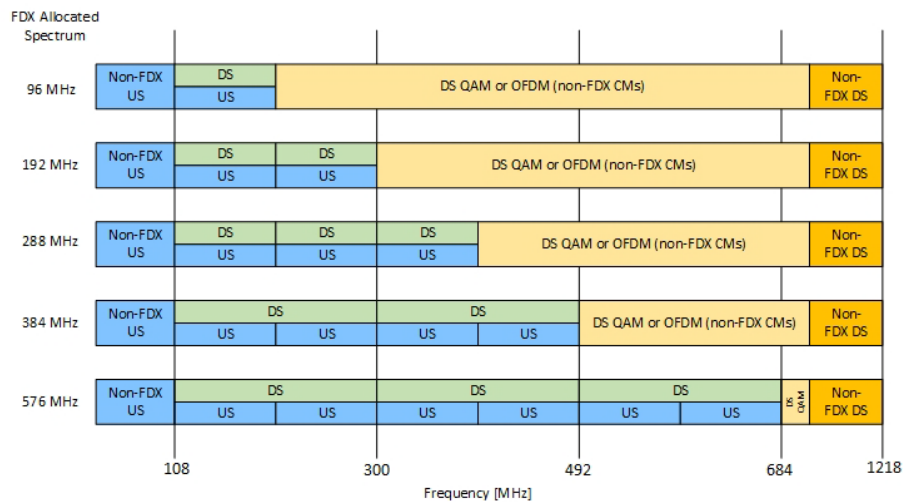


Figure 1 – FDX spectrum

The FDX CM can receive 4 total OFDM channels both FDX and non-FDX combined and optionally may support 5 OFDM channels. The FDX Node/CMTS support a minimum of 6 independently configurable OFDM channels.

The FDX CM supports a minimum of 7 independently configurable OFDMA upstream channels. The FDX Node /CMTS supports 8 configurable OFDMA upstream channels, each occupying a spectrum of up to 95 MHz.

FDX CMs are new purpose built CMs, and are designed with hardware and software capable of supporting FDX functionality. A D3.1 CM can be software upgraded to support a number of features in order to be capable of handling limited FDX functionality, these CMs are known as FDX-L CMs. FDX-L CMs are D3.1 CMs software upgraded with limited capabilities for operating within the FDX Band. Since there are two types of D3.1 CMs: mid-split and high-split, those are the flavors of the FDX-L CMs.

There are three FDX Operational Modes in the Occupied FDX Band.

- FDX-only: only FDX CMs in operation.
- FDX/High-split Coexistence : FDX CMs and high-split FDX-L CMs operate simultaneously.
- FDX/Mid-split Coexistence : FDX CMs and mid-split FDX-L CMs operate simultaneously

FDX CMs can transmit or receive in any FDX Sub-band and any FDX Operational Mode. In FDX/High-split Coexistence mode, high-split capable FDX-L CMs can transmit in FDX Sub-bands located up to 204 MHz and can receive in FDX Sub-bands that are positioned at or above 258 MHz. In FDX/Mid-split Coexistence mode, mid-split capable FDX-L CMs can receive in any FDX Sub-band.

1.2. FDX Initialization Overview

An FDX-capable CM first becomes operational as a D3.1 CM. The FDX-specific CM initialization is commenced under direction of the CMTS. The FDX initialization needs to be complete before the CM can transmit or receive data within the Occupied FDX Band.

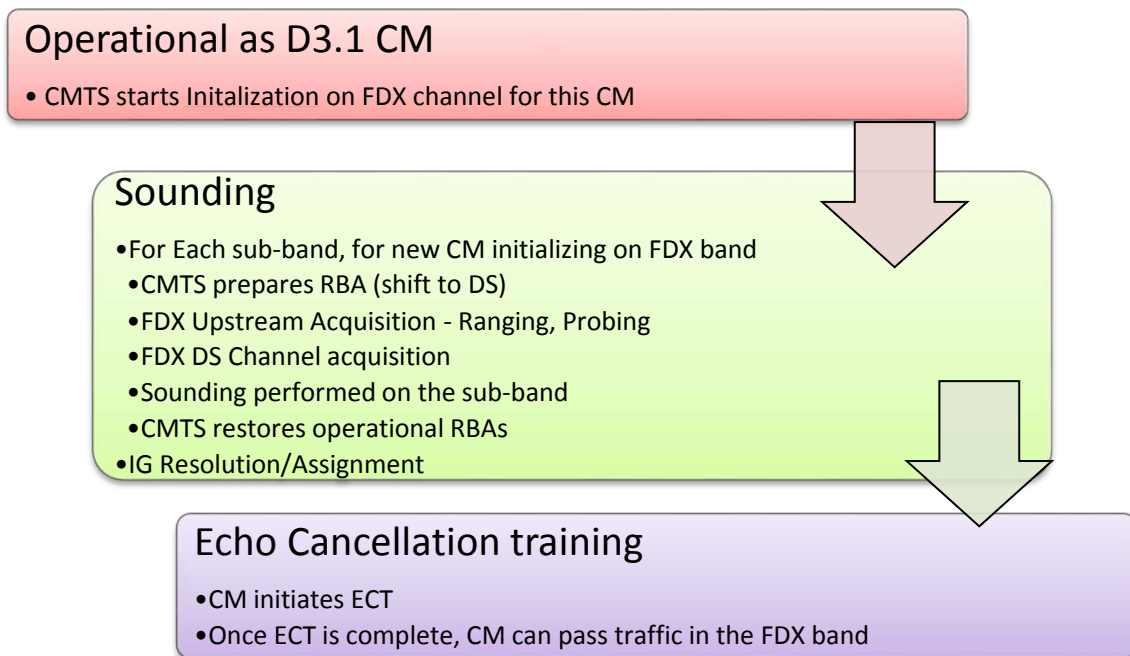


Figure 2 – FDX CM Initialization

The framework for FDX-specific CM initialization is shown in Figure above. The CMTS directs an initializing FDX-capable CM through specific procedures for each FDX Sub-band. The procedures depend upon the FDX Operational Mode, the current FDX Sub-band under consideration and the type of CM that is being initialized.

For each FDX Sub-band where an initializing FDX-capable CM will transmit and/or receive, the CMTS first coordinates RBA reconfigurations as needed so that Sounding may be performed on the sub-band. FDX Upstream Channel acquisition is then done for CMs that will transmit in the FDX Sub-band. FDX Downstream Channel acquisition is done for FDX-capable CMs that will receive in the FDX Sub-band. The Sounding is performed on the sub-band.

Once all relevant FDX Sub-bands have been addressed for the initializing FDX-capable CM, the CMTS coordinates RBA reconfigurations as needed to restore normal traffic-bearing conditions. The CMTS then assigns an IG, TG (and RBA) to the CM and directs any necessary channel acquisition steps to align the CM's active FDX channels with the assigned RBA. If per the RBA the initializing FDX-capable CM will be transmitting on either one or two active FDX Upstream Channels, the CMTS will coordinate Echo Canceller Training procedures with the CM prior to providing data transmission opportunities on those channel(s).

The FDX-capable CMs use the timing offset from a ranged legacy upstream channel as the initial timing offset for an FDX upstream channel. When adding an FDX upstream channel to an FDX-capable CM, the CMTS ensures that for the new FDX channel it allocates a fine ranging opportunity, receives a fine ranging burst, and responds with a timing adjustment in a ranging response before allocating any other type of transmission to that CM for that upstream FDX channel. After performing fine ranging on an FDX upstream channel being added to an FDX CM, the CMTS probes the CM on that upstream channel

at least once to ensure the power level for that channel is properly set before requesting an FDX CM to transmit any burst that is not ranging or probing.

Any sounding measurement performed by an FDX CM prior to ECT convergence on an RBA is for reference purposes and for sorting out the initial IG. RxMER measurements prior to ECT convergence are not too useful for downstream profile determination. Prior to assigning an FDX CM a TG-ID, the CMTS adds all FDX upstream channels and all FDX downstream channels that the CM is expected to use.

2. FDX Sounding

In FDX, interferences between the bi-directional transmissions need to be mitigated for the signals to be properly received. When one CM transmits upstream to the CMTS, the US signal may leak through the cable plant and becomes interference in the DS direction at the receiving CMs. FDX addresses this by grouping CMs that interfere with each other, CMs in the same group transmit or receive along the same direction at any given frequency and time. CMs from different groups have enough RF isolation to allow simultaneous US and DS transmissions at the same frequency.

2.1. Sounding

An Interference Group (IG) is a group of CMs that can interfere with each other when the downstream and upstream channels they share are used simultaneously. This occurs when the levels of the CM-to-CM co-channel interference (CCI) are above a design threshold when one CM transmits and other CMs receive simultaneously over the same FDX spectrum. IG Discovery includes a test process, known as Sounding, to allow the CMTS to assess the CCI level between any CM pair that may share the same spectrum for FDX operation. During Sounding, the CMTS selects one or more FDX capable CMs as ‘test’ CMs to transmit test signals on designated subcarriers, while directing other CMs as ‘measurer’ CMs to compute and report the received MER (RxMER) on the same set of subcarriers. The CMTS repeats this procedure until the interference levels are tested on all relevant subcarriers and between all CM combinations. A Test CM refers to an FDX-Capable CM that transmits the sounding test signal in an FDX sub-band to allow the CMTS to detect potential co-channel interferences that other FDX-Capable CMs may experience when operate in the DS direction in the same FDX sub-band. A Measurer CM refers to an FDX-Capable CM that measures and reports the RxMER in an FDX sub-band to allow the CMTS to detect the co-channel interference caused by one or multiple Test CMs’ transmitting the sounding test signals in the same FDX sub-band.

For a given FDX sub-band, the CMTS selects one or more FDX-capable CMs as the Test CMs to transmit test signals, while directing other FDX-Capable CMs to measure and report the DS RxMER as the Measurer CMs. The CMTS repeats this procedure until the interference relationships are tested on all relevant frequencies in the FDX sub-band and between all intended Test and Measurer CM pairs.

2.1. Sounding methods

To assess the CM to CM CCI, the CMTS allocates one or multiple sounding opportunities in the FDX spectrum. A Sounding Opportunity can be either a CWT Sounding Opportunity or an OUDP Sounding Opportunity, see figure below.

A CWT (Continuous Wave Tone) Sounding Opportunity is narrow in frequency but lasts longer in time; an OUDP (OFDM Upstream Data Profile) Sounding Opportunity covers the entire FDX sub-band but lasts shorter in time. Regardless of the type of the test signal used for sounding, a Sounding Opportunity

consists of a Test Signal Transmission Opportunity and a Test Signal Interference Region. A Test Signal Transmission Opportunity specifies a minislot region that contains the subcarrier locations for transmitting the test signals and necessary guard subcarriers adjacent to the neighboring regular US transmission region. A Test Signal Interference Region contains a consecutive set of DS subcarriers that encompasses the test signal transmissions in both time and frequency. The CMTS protects the DS transmissions in the Test Signal Interference Region with zero-bit-loaded subcarriers in case of CWT sounding or zero-bit-loaded symbols in case of OUDP sounding.

An FDX CM has the capability to generate multiple CWTs at specific frequency locations. Similar to the sounding routine using the CWT Test Signal as an interference source, there is an alternative approach that allows using the OUDP Test Bursts as an interference source. These OUDP Test Bursts can be used as the test signal for sounding when the measurer CMs involve only FDX CMs. However, both FDX CMs and FDX-L CMs may operate as test CMs.

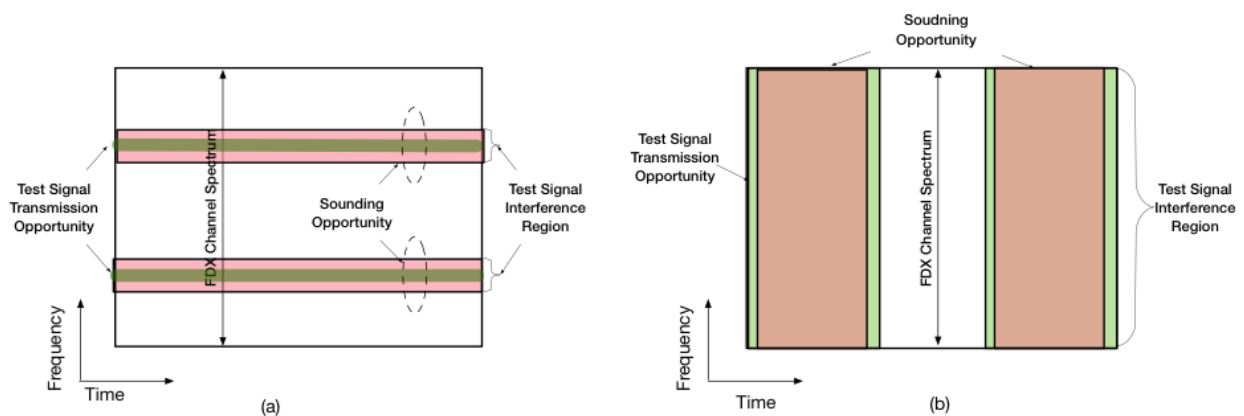


Figure 3 – Sounding Opportunities (a) CWT Sounding (b) OUDP Sounding

2.2. IGs & TGs

The measured interference allows the CMTS to sort CMs into IGs per FDX sub-band, such that for a given IG, the CCI experienced by any CM inside the IG due to the US transmission from at least one other CM in the IG is greater than the desired design limit and the CCI experienced by any CM outside the IG due to the US transmission from any CM inside the IG is less than the desired design limit.

Since the path loss, which determines the interference between a Test CM and a Measurer pair, could vary significantly over frequency, a Test CM may be required to send test signals at multiple subcarrier locations for IG discovery within the FDX sub-band. Consequently, the CCI limit of an IG can be represented as a function of frequency, for example, as a list of threshold values corresponding to different frequency locations in the FDX sub-band under test. The algorithm that determines the IG CCI limit and the test signal frequency locations are CMTS vendor-specific.

Interference Groups (IGs): Identify groups of CMs, that would interfere with each other if they were allowed to transmit and receive at the same time in a sub-band

Transmission Groups (TGs): IGs are grouped together into smaller number of TGs, which is used by the CMTS scheduler. CMs in same TG either transmit or receive on any given sub-band & time. CMs from different TGs have enough isolation to transmit and receive at same time in the same sub-band

2.3. Initial and Periodic Sounding

Initial sounding is the sounding operation the CMTS conducts before an IG has been identified for an FDX-Capable CM in an FDX sub-band. It's the first stage in IG Discovery for the CMTS to establish the initial interference relationship between the sounding CM and other FDX-Capable CMs operating in the sub-band.

Periodic sounding refers to the subsequent sounding operations after the initial sounding. It's the second stage in IG Discovery for the CMTS to monitor the CCI variations over time among the FDX cable CMs operating in a given FDX sub-band. Periodic sounding is also used by the CMTS to incrementally refine CCI estimations with more test samples at different frequencies and time.

2.4. CWT Sounding Overhead

For the CWT sounding, a sounding test opportunity includes a set (one or more) of CWT subcarriers and a few guard subcarriers on both sides, to prevent inter-symbol interference at adjacent data subcarriers. Comparing to the OUDP sounding, a CWT test opportunity occupies much narrower spectrum however lasts longer in time. A CWT tone typically lasts around 800 ms and the corresponding time in the DS is an additional amount of time on each side to make sure the measurement is aligned and complete. It typically takes around 200~300 milliseconds for the RxMER measurement scheme to converge.

In case of the CWT sounding, given a Sounding Opportunity only occupies a fraction of an FDX US/DS channel spectrum, simultaneous US data transmissions (data or CWTs from other CMs) may coexist with the CWTs at frequency locations outside the CWT Transmission Opportunity in the same FDX US channel. A CWT Transmission Opportunity specifies a minimum region of minislots containing the subcarrier locations of one or multiple CWTs for a duration required for the Measurer CMs to measure the RxMER while the CWTs are being transmitted. Given a CWT may cause the inter-carrier interferences (ICI) to adjacent US data subcarriers, a CWT Transmission Opportunity needs to include certain guard subcarriers on each side of the CWTs and a guard time to allow proper ramping at both the start and the end of the CWT transmission.

Table 1 – CWT Sounding Size

Component	Calculation	Value
K	number of symbols per frame	36
Upstream Elementary Period Rate	$(T_{su}) = 1 / \text{Upstream Sampling Rate}$	1/102.4 MHz
Symbol size	$4096 * T_{su}$	40 μ s
Cyclic Prefix Size	$192 * T_{su}$	1.875 μ s
Frame size	$K * (\text{symbol size} + \text{CP size})$	1.508 ms
CWT ramp size	$(128 * 8192 * T_{su} / K) * 2 \text{ ramps}$	14.22 frames
CWT duration	Chosen by CMTS	525 US frames
Total Total CWT Time US	$(\text{CWT ramp size} + \text{CWT duration}) * \text{Frame size}$	812.7 ms
Total CWT Time DS	20 % overhead on either side of CWT to co-ordinate ramp, RxMER measurements etc. $(1.4 * \text{Total CWT Time US})$	1138 ms

The CMTS allocates a CWT Transmission Opportunity for the whole duration of the intended CWT transmissions rounded up to the OFDMA frame boundary, including the headroom to cover the ambiguity of the CWT-REQ propagation delay and processing time. The CMTS also ensures sufficient CWT transmission time to allow all the Measurer CMs participating in sounding to measure and report RxMER. The table above shows the calculations needed to compute the CWT total duration.

A CWT Interference Region defines the DS spectrum in time and frequency that may be impacted by a CWT transmission in sounding. The CMTS uses zero-bit loading at the DS subcarriers corresponding to the CWT frequency locations and includes a number of guard subcarriers on each side of the CWTs to avoid ICI to adjacent DS data subcarriers.

The table below shows the overhead of the Periodic Sounding process using the CWT method. The assumption is the CMTS will ask the CM to transmit a CWT in a sub-set of minislots across the channel. This could be for example one CWT for every minislot in the channel, which equates to 237 CWTs. (There are 237 minislots in a 95 Mhz OFDMA channel.) The CMTS could ask the CM to transmit multiple CWTs within a minislot, and a subset of minislots across the channel. Another example would be to use 50 minislots for one CM's CWT transmission (with say 3-4 CWTs within each minislot). The below table calculates the overhead for Sounding. As one CM uses less of the channel, the remainder of the channel can be used by other CMs to complete the sounding process sooner.

Table 2 – CWT Sounding Overhead

Number Minislots consumed by CWTs from each test CM	Part of channel occupied by each Test CM (out of 237 Minislots)	CWT Sounding Time for Each CM (ms)	CWT Sounding Time for All CMs (Seconds)	Overhead w Sounding Cycle = 3600 secs	Overhead w Sounding Cycle = 10800 secs
237	100%	1138	72.834	2.02%	0.67%
200	84%	960.362	61.463	1.71%	0.57%
150	63%	720.271	46.097	1.28%	0.43%
100	42%	480.181	30.732	0.85%	0.28%
50	21%	240.090	15.366	0.43%	0.14%
Number of CMs 64					

2.5. OUDP Sounding Overhead

The CMTS allocates an OUDP Test Signal Interference Region for the whole duration of the OUDP Test Bursts Transmission Opportunity with an addition of the ambiguity of the propagation delay difference in between farthest and closest CMs on the plant and frames misalignment between two upstream channels of the sub-band and downstream channel recovery time. The CMTS allocates this region fully composed of zero-bit loaded corresponding DS data subcarriers. The allocated recovery time is based on the longest value of the t-ds-reacquisition capability on the MAC domain. The away time, in this case, is a time between the start of the first transmitted upstream symbol to the end of the last transmitted upstream symbol on a sub-band. The below table shows the calculation steps to compute the OUDP sounding duration.

Table 3 – OUDP Sounding Size

Component	Calculation	Value
K	number of symbols per frame	36
US Elementary Period Rate	$(T_{su}) = 1 / \text{Upstream Sampling Rate}$	1/102.4 MHz
Symbol size	$4096 * T_{su}$	40 μ s
Cyclic Prefix Size	$192 * T_{su}$	1.875 μ s
Frame size	$K * (\text{symbol size} + \text{CP size})$	1.508 ms
OUDP duration	Chosen by CMTS	50 US frames
OUDP Time	OUDP frame duration * frame time	75.375 ms
Ambiguity Offset	2 * Frame time	3.015 ms
Recovery Time	2* Ambiguity Offset	6.030 ms
Total OUDP Time (US)	OUDP Time + Ambiguity Offset + Recovery Time	84.420 ms

The CMTS does not grant OUDP Test Burst Transmission Opportunity to more than one test CM, to ensure that the source of the interference introduced by a particular test CM can be uniquely identified. For the OUDP sounding, a sounding test opportunity covers the entire FDX channel width in frequency and lasts about 60 to 80 milliseconds in time. Thus, no spectrum can be used for traffic when the OUDP sounding burst is present on the FDX channel under test.

Table 4 – OUDP Sounding Overhead

Part of channel Occupied by each Test CM	OUDP Sounding Time for Each CM	OUDP Sounding Time for All CMs	Overhead w Sounding Cycle = 3600 secs	Overhead w Sounding Cycle = 10800 secs
100%	84.420 ms	5.40 secs	0.15%	0.05%
Number of CMs = 64				

3. FDX Echo Cancellation

3.1. FDX Echo Cancellation at Node

To ensure proper operation of FDX at the Node, the interference resulting from FDX operation needs to be suppressed, via echo cancellation at the Node. Two types of EC techniques can be implemented in an RPD node to cancel or suppress the echoes. Analog EC cancels out the echoes in the analog domain before the ADC. Conventionally analog EC will take a copy of the DS signal, and manipulate its phase and magnitude to generate a canceling signal that is then added to the receiver path to cancel out the echo. Digital EC cancels out the echoes in the digital domain after ADC. After the echoes pass through the ADC and are converted into bits in digital domain, their magnitude and phase can be computed, and the cancelling signal can be generated from the DS reference signal with the proper magnitude and phase and subtracted from the received signal.

To cancel out the echo, the canceling signal is generated from the reference signal and needs to have the proper magnitude and phase. These EC coefficients are computed over a time period by comparing and

tracking the magnitude and phase difference between the reference and echoes embedded in the received signal. The procedure with which the EC coefficients are computed/tracked is called EC training, and the time period over the EC training period. The FDX node will schedule periods where it can complete EC training.

3.2. FDX Echo Cancellation at CM

Echo cancellation is used to improve FDX CM receiver performance by cancelling Adjacent Leakage Interference (ALI) and Adjacent Channel Interference (ACI) resulting from upstream transmissions. Echo Cancellation is required for each RBA in which there is at least one sub-band in the upstream direction and at least one sub-band in the downstream direction. The CM determines echo canceller training success. Until the Echo Canceller (EC) has converged the RBA sub-band direction set is not usable by the CM for anything but upstream maintenance, sounding, and EC training transmissions.

Prior to the CM's Echo Canceller being trained, the CMTS cannot request a CM to make an RxMER measurement in one sub-band while it is transmitting upstream in another sub-band. Hence, prior to its Echo Canceller being trained, an FDX CM cannot be a test CM in one sub-band while being a measurer CM in another sub-band at that same time.

3.2.1. *Foreground and Background Training*

The FDX CM uses two different methods to perform Echo Canceller Training, Foreground Training and Background Training.

In Foreground Training, the FDX CM transmits at regular power levels in the sub-bands that are in the upstream direction in the RBA sub-band direction set. Upstream bandwidth is dedicated via ECT probe allocations for all upstream channels in upstream sub-bands in the RBA. Foreground training may include zero bit loading (ZBL) on the downstream sub-bands in the RBA.

In Background Training, no upstream bandwidth is consumed. Instead, the FDX CM sends a low-level signal on the sub-band(s) that are downstream direction in the RBA. Background training does not require the use of probe allocations but does require assignment of a training window because the CMTS is responsible for limiting the number of CMs performing background EC training at the same time in order to manage the total emissions on the plant.

The FDX CM may require multiple EC training methods. Foreground training with ZBL is the only method that can be used for the cancellation of both ALI and ACI. Background training can be used for the cancellation of ALI, but not for the cancellation of ACI and foreground training without ZBL can be used for the cancellation of ACI, but not for the cancellation of ALI.

3.2.2. *Initial vs Periodic Training*

There are two phases to Echo Cancellation Training: initial and periodic. Initial EC Training is used to initially train the Echo Canceller for a given RBA sub-band direction set. Once the FDX CM has achieved sufficient convergence on an RBA sub-band direction set, periodic EC training is used to maintain sufficient convergence of the FDX CM's Echo Canceller for that RBA sub-band direction set.

3.3. EC Training & Overhead

For a CM, the recommended EC retraining period due to thermal drift is about 10 ~15 sec or as requested from the CM. The maximum number of symbols needed for training at a given time is estimated at a

maximum up to 128 symbols (upper bound). For background training the number of symbols needed is in the order of a few thousand symbols (~5000).

For initial training the table below calculates the time taken for ECT in various stages. For Initial training, the overhead is not meaningful as it happens one time at FDX initialization. Periodic training, is the main ECT to calculate the overhead for. Foreground training needs ECT probe allocations. Foreground training, if it uses ZBL on the Downstream, will also take away bandwidth from the downstream. Background ECT without ZBL doesn't take away active bandwidth but the CMTS needs to schedule these training windows for the CM.

ECT time is calculated as the Number of Symbols (for ECT) * symbol time * number of channels. For 2 US channels (in an RBA), a 40 μ s symbol time, and a CM which needs 128 symbols to EC train, will have a total EC time of (128*2*40=) 1024 μ s. The below table describes the values for case where the RBA has 2 US & 1 DS or 2 DS & 1 US sub-bands, for foreground training and for background training. The overhead is calculated as the (ECT time for all CMs) / EC periodicity.

Table 5 – EC training Overhead

Num Symbols needed for ECT	Num US Channels	Num DS Channels	EC time US (ms)	EC time DS (ms)	Total ECT time US (ms)	Total ECT time DS (ms)	US OverHead	DS OverHead (if ECT is with DS ZBL)
Symbol time = 40 μs, EC Periodicity = 15 seconds, Number of CMs = 64								
Foreground ECT w ZBL (Initial)								
128	2	1	10.24	5.12	655	328	NA	
128	1	2	5.12	10.24	328	655		
Foreground ECT w ZBL (Periodic)								
32	2	1	2.56	1.28	164	82	1.09%	0.55%
32	1	2	1.28	2.56	82	164	0.55%	1.09%
Background ECT wo ZBL (Initial or Periodic)								
5000	2	1	400	200	25600	12800	NA	
5000	1	2	200	400	12800	25600		

For Periodic training, the Training Expiration Times are the amount of time that the FDX CM is able to maintain EC training sufficient convergence when performing Foreground and Background training. This can be in multiples of the Training Periodicity.

4. FDX Resource Block Allocation

FDX operation is full duplex from the perspective of the CMTS but frequency division duplex (FDD) from the perspective of the CM. The FDX band is divided into sub-bands and the CMTS assigns which sub-band(s) each CM is to use for upstream FDX operation and which sub-band(s) each is to use for downstream FDX operation. This assignment is referred to as a resource block assignment (RBA). The sub-band is the atomic unit of allocation meaning that the entire sub-band is either assigned to be used for downstream traffic or upstream traffic. It is recognized that different CMs will have different bandwidth demand for both the upstream and downstream directions and that this demand can change dynamically. FDX allows resource assignment to be changed dynamically and this is known as Dynamic FDD.




Figure 4 – RBA Example

The CMTS periodically sends the RBA message to describe the current or upcoming RBA mapping. RBAs are associated with TGs and CMs are assigned to TGs. The above figure shows 2 TGs which are assigned different resource block assignments, TG1 (up,down,down) and TG2 (down, down, up).

4.1. Fast and Slow RBA Switching

Switching is defined as a change in the RBA for the TG. RBA switching has been designed to support a range of implementations all the way from static systems to highly dynamic systems. This provides operators with a tool set with which FDX deployments can be tuned to specific needs. RBA switching needs to account for channel acquisition and tracking processes that are dependent upon channel attributes which are time varying. Processes include channel acquisition and echo cancellation tracking. The duration since a sub-band last had an identical direction assignment determines if an RBA switch is ‘fast’ or ‘slow’. The ‘Away Time’ can be defined as the amount of time a CM spends away from a certain direction and in the opposite direction for a particular sub-band, before it returns to the original direction. The CMTS determines if a switch is fast or slow based upon the duration of the away time and CM capabilities. The switching procedure is the same for both fast and slow switching, the difference is the duration a CMTS waits before sending or scheduling PDUs after a sub-band direction change.

When switching a sub-band assignment to the downstream direction, the CMTS waits for at least the downstream switching reacquisition time prior to sending traffic to a CM on a downstream channel in the sub-band. Likewise, when switching a sub-band to the upstream direction, the CMTS re-ranges and allows a CM to re-train echo cancellation prior to providing grants if the away time has been greater than that specified by the CM capabilities.

RBAs can change regularly and quickly. Thus, the rate at which RBAs are sent could possibly become taxing to FDX-L CMs if each message had to be processed by the CMs CPU. To prevent this, RBA messages are of two types a hardware friendly version and a software friendly version. FDX-L CMs are expected to process software-based RBAs and to drop the hardware-based RBAs. This allows a CMTS to send hardware based RBAs at a very high rate without impacting FDX-L CMs. The CMTS will inform FDX-capable CMs which type of RBA message to utilize based on CM capabilities.

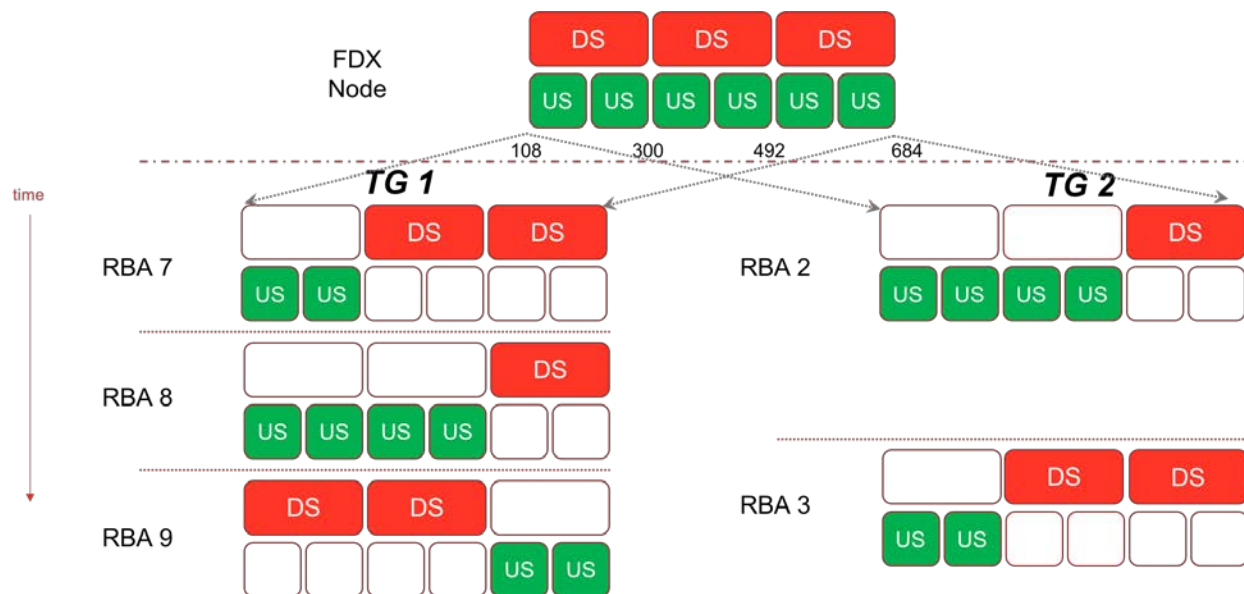


Figure 5 – RBA Switch in different TGs

To avoid CM-to-CM interference, in a given FDX sub-band, the CMTS makes sure all CMs in the same TG (group of IGs) are operating in the same direction at the same time, each TG is assigned an RBA. The figure above shows the RBA sequence changing independently across two TGs.

4.2. RBA Switching Use cases

There are many reasons an CMTS would want to allocate or change RBAs in the FDX band. The way the operator would change the RBA depends on the deployment scenarios and the bandwidth to offer the customers. Some of the reasons why an operator may want to switch RBAs are as follows:

Load Balancing across the plant: E.g., diurnal cycles: There could be more activity from business users during the day, home users in the evening. An operator may choose to run the FDX band to be more symmetrical (with more upstream bandwidth) during the day for business but favoring downstream bandwidth in the evening for home users.

Activity patterns across the plant: At any given time, a particular TG may have many active users, while some other TG have many idle CMs. An operator may choose to set the FDX band to address different activity patterns in specific TGs

Immediate bandwidth demand changes across the plant. At a given time one user doing a large upload is filling most of the upstream sub-bands assigned to that TG, and then as a second user begins an upload, the TG requires more upstream bandwidth. In this case the operator may configure the CMTS to respond to such immediate changes in demand.

The following are the different configurations of RBA settings an operator could use to meet the various use cases.

Static RBA: In this case the operator, needs a certain amount of bandwidth in each direction. As an example, the operator may need US bandwidth equivalent to 2 FDX sub-bands. In this case the operator could set the RBA for the TG (or all TGs) with 2 US sub-bands. e.g. an RBA configuration of 110 (US, US, DS).

Duty Cycle RBA: In this case the operator, needs a certain amount of bandwidth in each direction which fluctuates over a time period. As an example, the operator may need US bandwidth equivalent to 2 FDX sub-bands during the day and DS bandwidth equivalent to 2 FDX sub-bands during the day. In this case the operator could set the RBA for the TG (or all TGs) to switch between a couple of RBAs over some time period e.g. an RBA configuration of 100 (US, DS, DS) \leftrightarrow 110 (US, US, DS) which could change at a rate chosen by the operator (could be on the order of hours or minutes).

Dynamic RBA: In this case the operator, configures the CMTS (or an external application) to drive the RBA configuration and let it decide how to set RBAs based on current demand. In this case the CMTS could set the RBA for the TG (or all TGs) to switch between any possible RBA combinations (could be on the order of 100 ms) as per the demand on the network.

4.3. RBA switching & Overhead

When the CMTS changes RBAs, it coordinates the sending and scheduling of traffic such that all traffic is sent or received by a CM on a channel that is in the CM's TG's RBA and is valid for that CM at the time the traffic is sent or received. When the CM is switching the RBAs there is a period of time which the CM needs to change state from one direction to another. This budget to make changes along with other network parameters will be used by the CMTS to determine when it can schedule or send traffic from or to each CM that is affected by a change in the RBA.




Figure 6 – RBA Switching Time Calculation

The Figures above shows the component in computing the time it takes to change the direction of a sub-band from one direction to the other. In this figure the sub-band is in the downstream direction prior to T-Suspend-ds and switches to upstream after T-rba. The labels are explained below.

- T-rba: The time the new RBA becomes active (from the perspective of the CMTS).
- CM Switch Time (t-switch-xx): The maximum duration it takes for a CM to enact an RBA change from DS to US (t-switch-du), or from US to DS (t-switch-ud):
- t-cm-rba-proc: This is the minimum time that a CMTS allows for a CM to process RBA messages in advance of the time that a CM will begin to enact the RBA.
- Downstream Latency (t-latency-ds): The time it takes for a PDU to travel from the CMTS to the CM. It includes any time-interleaving, D3.1 convergence layer processing, and CIN Latency.
- CIN Jitter (t-jitter-cin): The maximum variance in CIN Latency a PDU could take to traverse the CIN between the CCAP Core and the RPD in the RPHY architecture. (In a CMTS architecture with co-located MAC & PHY, the CIN Jitter will approach zero.)

- Upstream Latency (t-latency-us): The time it takes for a PDU to travel from the CM to the CMTS burst receiver (within the CMTS PHY).
- T-suspend-ds: The CMTS ensures that no PDUs arrive at a transitioning CM after the CM has switched that spectrum from DS to US.

The following numbers will describe the components of the RBA Switch Time.

Table 6 – RBA OverHead

Delay Component	Calculation	Value
DS PHY layer propagation delay	5 μ s per KM \rightarrow 8 km distance	40 μ s
DS Interleaving Delay	(M-1) * Symbol time = (3-1) * 22.5 μ s	45 μ s
DS pipeline delay	3 * DS symbol size = 3 * 22.5 μ s	67.5 μ s
RPD Forwarding Latency	200 μ s	200 μ s
t-latency-ds	Propagation delay + Interleaving delay + Pipeline delay + RPD Fwd latency	352.5 μ s
US PHY layer propagation delay	5 μ s per KM \rightarrow 8 km distance	40 μ s
US pipeline delay	(3 * Frame Size)	405 μ s
t-latency-us	Propagation delay + Pipeline delay	445 μ s
t-switch-du	Depends on the capability of the new FDX silicon. Somewhere in the range of 10 μ s to 1000 μ s	10 μ s, 500 μ s or 1000 μ s
t-jitter-cin	Depends on the CIN network designed by the operator	1ms to 8ms
Total Pause Time for FDX channel	t-latency-us + t-switch-du + t-latency-ds + t-jitter-cin	~Range from 1.8 ms to 9.7 ms

The figure below shows the RBA switch time (pause time) for an FDX channel as a function of the CIN jitter, which is the biggest contributing factor as seen in the table above. Each curve in the figure is for a different CM Switch time (which is a CM capability) from 10 μ s to 1000 μ s.




Figure 7 – CM RBA Switch/Channel Pause Time

The figure below shows an FDX Sub-band changing from upstream to downstream and back to upstream. During each of these RBA changes there is sometime went to CMTS pauses transmission on channel.

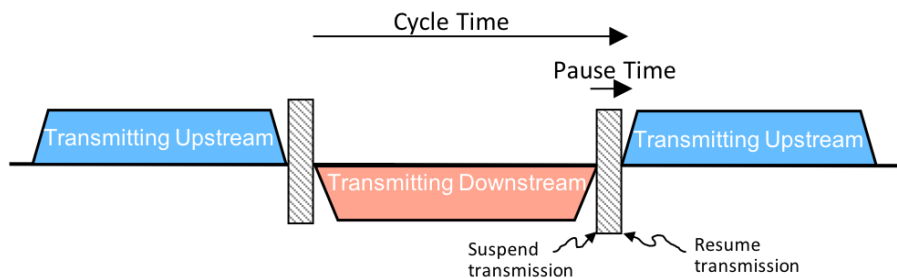


Figure 8 – RBA Overhead

The overhead is defined as the “pause” time during which transmissions are halted divided by the “cycle time” (time data being transmitted + switching time). For a given FDX channel: $\text{Overhead} = \frac{\text{Pause Time}}{\text{cycle time}}$. If the FDX sub-band is in the downstream direction (cycle time) for 100 ms and the switch time is 1.8 ms, the RBA Change Time Overhead is 1.8%. The below graph shows the overhead in RBA switching as the FDX direction cycle time ranges from 100 milliseconds to a 1000 ms. Each curve is for a CM’s RBA switch time from 1.8 ms to ~10 ms. For RBA switch cycles greater than 400 ms the overhead is less than 2%, and as the cycle time is greater than a second the overhead is less than 1%




Figure 9 – FDX Channel Pause/Cycle Time (Overhead)

FDX Scenarios

5. FDX Capacity Basics

5.1. FDX Modulation orders

The following table shows the RxMER (dB) level needed at a CM to be able to receive data traffic at a certain modulation order.

Table 7 – RxMER needed for Modulation orders

Constellation	D3.1 DS MER(db)	FDX DS *	D3.1 US	FDX US
QPSK			11	12.5
8-QAM			14	15.5
16-QAM	15	16.5 (17/17)	17	18.5
32-QAM			20	22
64-QAM	21	22.5 (23/23)	23	25.5
128-QAM	24	25.5 (26/25.5)	26	29
256-QAM	27	28.5 (29/29)	29	32
512-QAM	30.5	31.5 (33/32)	32.5	36
1024-QAM	34	35 (42/36)	35.5	44
2048-QAM	37	40 (NA/44.5)	39	
4096-QAM	41		43	
8192-QAM	46			
16384-QAM	52			

* The FDX DS values are from the tables for the External ACI Test and the values in the (parentheses) are from the Self ACI Test as defined in the [D3.1 PHY Spec].

5.2. FDX TG formation

The FDX architecture provides two-way signal transmission within the same spectral band. This requires a passive architecture without amplifiers. In this case, the fiber node connects to a single series of multiport taps. An example of one coaxial branch of a fiber deep node + 0 architecture is shown figure below. Here is an IG Discovery example using the RxMER measurement data listed in Table below.

Table 8 – CM-to-CM Interference Levels and Interference Groups

MER(db)		Transmit					
Receive		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
	Tap 1	6	39.9	39.9	39.9	39.9	39.9
	Tap 2	39.9	9.8	37.5	37.5	37.5	37.5
	Tap 3	39.9	37.5	9.8	34.5	34.5	34.5
	Tap 4	39.6	37.5	34.5	13.2	31	31
	Tap 5	39.9	37.5	34.5	31	15.8	27
	Tap 6	39.6	37.5	34.5	31	27	15.8

Modulation Order		Transmit					
Receive		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
	Tap 1	0	11.5	11.5	11.5	11.5	11.5
	Tap 2	11.5	0	11	11	11	11
	Tap 3	11.5	11	0	10	10	10
	Tap 4	11.5	11	10	0	9	9
	Tap 5	11.5	11	10	9	0	8
	Tap 6	11.5	11	10	9	8	0

The measured interference allows the CMTS to sort CMs into IGs per FDX sub-band. For a given IG, the CCI experienced by any CM in the IG due to the transmission from any other CM in the IG is greater than the desired design limit, and the CCI experienced by any CM outside the IG due to the transmission from a CM inside the IG is less than the desired design limit. The different colored cells show the potential IG groupings. This translates to 4 different IGs shown in the figure below.




Figure 10 - CM-to-CM Interference Levels and Interference Groups

Since the path loss, which determines the interference between a Test CM and a Measurer pair, could vary significantly over frequency, a Test CM may be required to send test signals at multiple subcarrier locations for IG discovery within the FDX sub-band. Consequently, the CCI limit of an IG can be represented as a function of frequency, for example, as a list of threshold values corresponding to different frequency locations in the FDX sub-band under test.

5.3. FDX Node, FDX CM & FDX-L CM Capacity

The FDX Node supports the FDX sub-bands and the channels in both directions at the same time i.e. it simultaneously receives & transmits in same FDX spectrum. The FDX Node performs echo cancellation to make sure that the transmitted signal does not interfere with the received signal. For the same amount of spectrum, the FDX Node is passing double the number of bits. This is in contrast with the CM that either receives or transmits in the same FDX spectrum at a given time.

The FDX system is essentially Full Duplex from perspective of CMTS and Frequency Division Duplex from perspective of CM. As the RxMER levels obtained are different across each Interference Group, each TG will have a different capacity as per the modulation orders/profiles which can be used by that

TG. The FDX CM capacity at a given time is limited to the RBA setting for its TG. The FDX Node capacity is not per TG, it is per MAC Domain as all the FDX channels are shared by all the TGs.

5.4. CM Channel Support

As discussed in the introduction, the following tables summarize the capabilities of an FDX CM and an FDX-L CM (A software upgraded D3.1 CM)

Table 9 – FDX Channel Support

Channel	Device	OFDM/OFDMA	SC-QAM
FDX			
Downstream	CM	4 OFDM Channels (3 FDX) (5 Optional)	32
Downstream	CMTS	6 OFDM Channels (3 FDX)	32
Upstream	CM	7 OFDMA Channel, (6 FDX, 2 Non-FDX)	4 (8 Optional)
Upstream	CMTS	8 OFDMA Channels (6 FDX, 2 Non-FDX)	4 (8) Optional)
FDX-L (D3.1)			
Downstream	CM	2 OFDM Channels	32
Downstream	CMTS	2 OFDM Channels	32
Upstream	CM	2 OFDMA Channels	8
Upstream	CMTS	2 OFDMA Channels	8

5.5. FDX MER and Capacity assumptions

The following table describes the OFDM/A channel capacities for various bit loading, after considering the overhead (PLC, NCP, FEC, pilots, pilot patterns, frame size etc). This calculation has been described in other papers, the values below assume a somewhat aggressive setting (of cyclic prefixes, pilots, pilot patterns) to get the more bits per sec in both upstream and downstream.

Table 10 – D3.1 Bit Loading and Channel Capacities

Bit Loading	DS Mbps	US Mbps
D3.1	OFDM	OFDMA
8	634	737
9	711.5	815
10	790.5	906
11	863	994
12	953.5	1072
D3.0	SC-QAM	SC-QAM
256 SCQAM	38.81	(not used in examples)

For the calculations in the series of figures (describing RBA configurations) in the next chapter, this paper assumes an average bit loading/modulation order of 12 for the an OFDM channel, a 11 for an FDX DS channel. Similarly, in the Upstream, it assumes an average bit loading/modulation order of 11 to be achievable for the an OFDMA channel, a 10 for an FDX US channel. Based these assumptions the following table shows the calculated bandwidth (Mbps) for the different types of channels used in chapter 6,7.

Table 11 – D3.1/FDX Bit Loading and Channel Capacities

Channel Type	Bandwidth (MBps)
FDX DS-192	1726
FDX DS	863
FDX US	906
D-192-OFDM	1907
D-150-OFDM	1489.8
U-96-OFDMA	994
U-80-OFDMA	828.3
SCQAM-32	1241.92
SCQAM-24	931.44

6. FDX RBA configurations

The following section describes the various RBA configurations which an operator can use, and the bandwidth provided by each configuration.

6.1. 5 Sub-band configurations

The following figure illustrates all the possible sub-band configurations within FDX. The FDX spectrum begins from 108 MHz and extends two 684 MHz. The following figure displays Spectrum in 96 MHz chunks. The downstream sub bands/channels are marked as “FDX DS” or “FDX DS-192” in red, while the upstream sub bands/channels are marked as “FDX US” in Green. The values on the right give the total MBps bandwidth as seen by the FDX CM.

FDX BAND											
108	204	300	396	492	588	684					
FDX							FDX CM (MBps)		FDX Node (MBps)		
							MHz	DS	US	DS	US
	FDX DS						96	863	0	863	906
	FDX US							0	906		
	FDX DS	FDX DS					192	1726	0	1726	1812
	FDX DS	FDX US						863	906		
	FDX US	FDX DS						863	906		
	FDX US	FDX US						0	1812		
	FDX DS	FDX DS	FDX DS				288	2589	0	2589	2718
	FDX DS	FDX DS	FDX US					1726	906		
	FDX DS	FDX US	FDX DS					1726	906		
	FDX DS	FDX US	FDX US					863	1812		
	FDX US	FDX DS	FDX DS					1726	906		
	FDX US	FDX DS	FDX US					863	1812		
	FDX US	FDX US	FDX DS					863	1812		
	FDX US	FDX US	FDX US					0	2718		
	FDX DS-192	FDX DS-192					384	3452	0	3452	3624
	FDX DS-192	FDX US	FDX US					1726	1812		
	FDX US	FDX US	FDX DS-192					1726	1812		
	FDX US	FDX US	FDX US	FDX US				0	3624		
	FDX DS-192	FDX DS-192	FDX DS-192				576	5178	0	5178	5436
	FDX DS-192	FDX DS-192	FDX US	FDX US				3452	1812		
	FDX DS-192	FDX US	FDX US	FDX DS-192				3452	1812		
	FDX DS-192	FDX US	FDX US	FDX US	FDX US			1726	3624		
	FDX US	FDX US	FDX DS-192	FDX DS-192				3452	1812		
	FDX US	FDX US	FDX DS-192	FDX US	FDX US			1726	3624		
	FDX US	FDX US	FDX US	FDX US	FDX DS-192			1726	3624		
	FDX US	FDX US	FDX US	FDX US	FDX US	FDX US		0	5436		
	96	192	288	384		576					

Figure 11 – RBA SubBand Configurations & Capacity

6.1. Static US & DS

The static US/DS case is likely the initial deployment strategy for operators who want to increase upstream bandwidth. The following figure illustrates again the possible sub-band configurations within FDX, this time re-arranged to show the static upstream case and the static downstream case. It also groups the other RBA configurations to show the various levels of bandwidth with the different RBA configurations, in the order of increasing US bandwidth.

Mhz							Mbps	
108	204	300	396	492	588	684	DS	US
	FDX DS						863	0
	FDX DS	FDX DS					1726	0
	FDX DS	FDX DS	FDX DS				2589	0
	FDX DS-192		FDX DS-192				3452	0
	FDX DS-192		FDX DS-192		FDX DS-192		5178	0
	FDX US						0	906
	FDX US	FDX US					0	1812
	FDX US	FDX US	FDX US				0	2718
	FDX US	FDX US	FDX US	FDX US			0	3624
	FDX US	FDX US	FDX US	FDX US	FDX US	FDX US	0	5436
	FDX DS	FDX US					863	906
	FDX US	FDX DS					863	906
	FDX DS	FDX DS	FDX US				1726	906
	FDX DS	FDX US	FDX DS				1726	906
	FDX US	FDX DS	FDX DS				1726	906
	FDX DS	FDX US	FDX US				863	1812
	FDX US	FDX DS	FDX US				863	1812
	FDX US	FDX US	FDX DS				863	1812
	FDX DS-192		FDX US	FDX US			1726	1812
	FDX US	FDX US	FDX DS-192				1726	1812
	FDX DS-192		FDX DS-192		FDX US	FDX US	3452	1812
	FDX DS-192		FDX US	FDX US	FDX DS-192		3452	1812
	FDX US	FDX US	FDX DS-192		FDX DS-192		3452	1812
	FDX DS-192		FDX US	FDX US	FDX US	FDX US	1726	3624
	FDX US	FDX US	FDX DS-192		FDX US	FDX US	1726	3624
	FDX US	FDX US	FDX US	FDX US	FDX DS-192		1726	3624

Figure 12 – Different view of the RBA capacities

The following figures show different examples of the possible US and D bandwidths for different RBA configurations. In each of these configurations the CMTS switches each TG from one RBA to another RBA, for a set amount of time, i.e. the CMTS is maintain a certain duty cycle for each RBA. RBAs which used have a value of $\frac{1}{2}$ or $\frac{1}{3}$ in the 'RBA Fraction' column indicating the CMTS is switching RBAs (at a certain rate) between those RBAs, and RBAs which are not used have a value of 0

Figure 13 – Duty Cycle RBA example 96 MHz FDX Band

Figure 14 – Duty Cycle RBA example 192 MHz FDX Band

Now summing up the effective data rates across the TG's gives (863,2718), which is greater than the node capacity (1726,1812) in the upstream and less in the downstream. This means that the downstream's on the node are being underutilized and the upstreams are being oversubscribed. This means that if a single user is using the peak upstream rate on one TG, another user on a different TG may not be able to use the upstream rate provided by the channel at the same time.

In the above example of a 384 MHz FDX band, we have just 1 TG. For TG1, the CMTS switches RBA's between RBA 00 (down, down) 1/4th of the time, RBA 01 (up, down) 3/8th the time, RBA 11 (up, up) 3/8th of the time, to get an effective bandwidth (1510,2038) for those TG.

Now summing up the effective data rates across the TG's, since there is only one TG here, gives (1510,2038) which is less than the node capacity (3452,3624), in the upstream and in the downstream. This means that both the downstream channels and the upstream channels are being underutilized.

108	204	300	396	492	588	684											
FDX							Node (Mbps)		Effective Data Rate								
							DS	US	DS	US	RBA Fraction	DS	US				
TG1	FDX DS-192		FDX DS-192		FDX DS-192		5178	0	5178	5436	0						
	FDX DS-192		FDX DS-192		FDX US	FDX US	3452	1812			0						
	FDX DS-192		FDX US	FDX US	FDX DS-192		3452	1812			0						
	FDX DS-192		FDX US	FDX US	FDX US	FDX US	1726	3624			0	2589.00	2718.00				
	FDX US	FDX US	FDX DS-192		FDX DS-192		3452	1812			1/2						
	FDX US	FDX US	FDX DS-192		FDX US	FDX US	1726	3624			0						
	FDX US	FDX US	FDX US	FDX US	FDX DS-192		1726	3624			1/2						
	FDX US	FDX US	FDX US	FDX US	FDX US	FDX US	0	5436			0						
TG2	FDX DS-192		FDX DS-192		FDX DS-192		5178	0					0				
	FDX DS-192		FDX DS-192		FDX US	FDX US	3452	1812					1/2				
	FDX DS-192		FDX US	FDX US	FDX DS-192		3452	1812					0				
	FDX DS-192		FDX US	FDX US	FDX US	FDX US	1726	3624					1/2	2589.00	2718.00		
	FDX US	FDX US	FDX DS-192		FDX DS-192		3452	1812					0				
	FDX US	FDX US	FDX DS-192		FDX US	FDX US	1726	3624					0				
	FDX US	FDX US	FDX US	FDX US	FDX DS-192		1726	3624					0				
	FDX US	FDX US	FDX US	FDX US	FDX US	FDX US	0	5436					0				
TG3	FDX DS-192		FDX DS-192		FDX DS-192		5178	0					0				
	FDX DS-192		FDX DS-192		FDX US	FDX US	3452	1812					1/2				
	FDX DS-192		FDX US	FDX US	FDX DS-192		3452	1812					1/2				
	FDX DS-192		FDX US	FDX US	FDX US	FDX US	1726	3624					0	3452.00	1812.00		
	FDX US	FDX US	FDX DS-192		FDX DS-192		3452	1812					0				
	FDX US	FDX US	FDX DS-192		FDX US	FDX US	1726	3624					0				
	FDX US	FDX US	FDX US	FDX US	FDX DS-192		1726	3624					0				
	FDX US	FDX US	FDX US	FDX US	FDX US	FDX US	0	5436					0				
TG4	FDX DS-192		FDX DS-192		FDX DS-192		5178	0					0				
	FDX DS-192		FDX DS-192		FDX US	FDX US	3452	1812					0				
	FDX DS-192		FDX US	FDX US	FDX DS-192		3452	1812					0				
	FDX DS-192		FDX US	FDX US	FDX US	FDX US	1726	3624					0	863.00	4530.00		
	FDX US	FDX US	FDX DS-192		FDX DS-192		3452	1812					0				
	FDX US	FDX US	FDX DS-192		FDX US	FDX US	1726	3624					0				
	FDX US	FDX US	FDX US	FDX US	FDX DS-192		1726	3624					1/2				
	FDX US	FDX US	FDX US	FDX US	FDX US	FDX US	0	5436					1/2				
</																	

Figure 17 – Duty Cycle RBA example 576 MHz FDX Band

In the above example of a 576 MHz FDX band, we have 4 TGs.

For TG1, the CMTS is switching RBA's between RBA 100 (up,down,down) 1/2 of the time and RBA 110 (up,up,down) 1/2 the time, gets an effective bandwidth (2589, 2718) for the TG.

For TG2, the CMTS is switching RBA's between RBA 001 (down,down,up) 1/2 of the time and RBA 011 (down,up, up) 1/2 the time, getting the same effective bandwidth as TG1, (2589, 2718) for the TG 2.

For TG3, the CMTS is switching RBA's between RBA 001 (down,down,up) 1/2 of the time and RBA 010 (down,up, down) 1/2 the time, getting the effective bandwidth of (3452, 1812) for TG3.

For TG4, is more upstream heavy, and it switches RBA's between RBA 110 (up,up,down) 1/2 of the time, RBA 111 (up,up,up) 1/2 the time, to get an effective bandwidth (863,4530) for those TG.

Now summing up the effective data rates across the TG's gives (9493,11778) which is greater than the node capacity (5178,5436), in the upstream and in the downstream. This means that both the downstream channels and the upstream channels are being oversubscribed. This means that if a single user is using the peak rate on one TG, another user on a different TG may not be able to use the rate provided by the channel at the same time.

6.3. Dynamic RBAs and the Need for RBA Management

Giving the examples seen above, one can see that setting RBA's for the set of TGs and the set of CMs in each TG, can quickly become complicated. The RBA setting depends on the service level of an individual user, the aggregate service needs across CM's in a TG, and the aggregate service needs across all TGs.

An operator for initial FDX deployments, can start with static RBA settings (this could be different static RBA's for the different TG's, based on the needs of each TG). A second step would be simple duty cycle RBA settings, to get finer granularity of FDX upstream and downstream bandwidth capabilities.

Once an operator moves past static RBA settings and simple duty cycle RBA settings, the goal would be to set RBAs dynamically based on the needs in the plant. This means tracking the aggregate upstream and downstream usage on a TG basis, and predicting the appropriate RBA setting needed to best meet the service level agreements and performance on the FTX channels

D3.1 introduced the need for dynamic profile management across an OFDM/OFDMA channel. This has become an application external to the CMTS which can create profiles and push them into the CMTS. In a similar fashion one can imagine RBA management application combined with profile management for the FDX channels becoming a necessity to manage bandwidth across the FDX band effectively

7. FDX + D3.1 Channel Capacity

Now that we are familiar with the capacity provided by FDX channels, in this section we look to understand maximum data rates which can be provided by an FDX and an FDX- L CM.

In addition to the FDX band of channels, based on the D3.1 and FDX CM capabilities, the assumption here is as follows. For the downstream, CMs can tune to 32 SC QAMs, and a D3.1 OFDM channel(192Mhz) outside of the FDX band and use OFDMA channels for the Upstream. The Channel capacity scenarios, in the figures below shows the CM using the FDX band and in addition shows the CM being able to use the upstream from 5-85 MHz and non-FDX DS from 684 to1218 MHz.

Now FDX CM cannot receive non-FDX channels (SC-QAM or OFDM) within the FDX band (108-684 MHz). So even if the occupied FDX band is say 108 to 300 MHz, the FDX-CM cannot receive non-FDX channels in the 300-684 Mhz zone. This drives the placement of the 192 MHz OFDM channel, and the 32 SCQAM channels above 684 MHz. For the upstream the CM could receive SCQAM or OFDMA channels, or both using Time and Frequency division multiplexing, but for simplicity the assumption is an OFDMA channel across the usable upstream spectrum. The capacity of the channels is calculated based on the table defined in section 5.5.

7.1. FDX CM Common Use case (Use Case 1)

In this scenario, the operator introduces the FDX band for use with FDX channels. In addition, the operator introduces a (non-FDX) 192 MHz OFDM channel (684-876 zone), and adds 32 SC-QAM channels in the 876-1068 zone. The CM is also receiving an 80 MHz OFDMA upstream in the legacy band.

5	85	108	204	300	396	492	588	684	780	876	972	1068	1164	1218	FDX CM (Mbps)	
															DS	US
FDX																
U-80-OFDMA			FDX DS						DS DS 1.1 OFDM 192		DS DS 0.5 SC-QAM 32*6					
U-80-OFDMA			FDX US						D-192-OFDM		SCQAM-32				401.2	8.28
									D-192-OFDM		SCQAM-32				3149	1734
U-80-OFDMA			FDX DS	FDX DS					D-192-OFDM		SCQAM-32				485.5	8.28
U-80-OFDMA			FDX DS	FDX US					D-192-OFDM		SCQAM-32				401.2	1734
U-80-OFDMA			FDX US	FDX DS					D-192-OFDM		SCQAM-32				401.2	1734
U-80-OFDMA			FDX US	FDX US					D-192-OFDM		SCQAM-32				3149	2640
U-80-OFDMA			FDX DS	FDX DS	FDX DS				D-192-OFDM		SCQAM-32				5738	8.28
U-80-OFDMA			FDX DS	FDX DS	FDX US				D-192-OFDM		SCQAM-32				485.5	1734
U-80-OFDMA			FDX DS	FDX US	FDX DS				D-192-OFDM		SCQAM-32				485.5	1734
U-80-OFDMA			FDX DS	FDX US	FDX US				D-192-OFDM		SCQAM-32				401.2	2640
U-80-OFDMA			FDX US	FDX DS	FDX DS				D-192-OFDM		SCQAM-32				485.5	1734
U-80-OFDMA			FDX US	FDX DS	FDX US				D-192-OFDM		SCQAM-32				401.2	2640
U-80-OFDMA			FDX US	FDX US	FDX DS				D-192-OFDM		SCQAM-32				401.2	2640
U-80-OFDMA			FDX US	FDX US	FDX US				D-192-OFDM		SCQAM-32				3149	3546
U-80-OFDMA			FDX DS-192		FDX DS-192				D-192-OFDM		SCQAM-32				6601	8.28
U-80-OFDMA			FDX DS-192	FDX US	FDX US				D-192-OFDM		SCQAM-32				485.5	2640
U-80-OFDMA			FDX US	FDX US	FDX DS-192				D-192-OFDM		SCQAM-32				485.5	2640
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US			D-192-OFDM		SCQAM-32				3149	4452
U-80-OFDMA			FDX DS-192		FDX DS-192			FDX DS-192	D-192-OFDM		SCQAM-32				8327	8.28
U-80-OFDMA			FDX DS-192	FDX DS-192	FDX DS-192			FDX US	FDX US		SCQAM-32				6601	2640
U-80-OFDMA			FDX DS-192	FDX US	FDX US				FDX DS-192		SCQAM-32				6601	2640
U-80-OFDMA			FDX DS-192	FDX US	FDX US	FDX US	FDX US	FDX US	FDX US		SCQAM-32				485.5	4452
U-80-OFDMA			FDX US	FDX US		FDX DS-192		FDX DS-192	D-192-OFDM		SCQAM-32				6601	2640
U-80-OFDMA			FDX US	FDX US	FDX DS-192	FDX DS-192		FDX US	FDX US		SCQAM-32				485.5	4452
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US			FDX DS-192		SCQAM-32				485.5	4452
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US	FDX US	FDX US	D-192-OFDM		SCQAM-32				3149	6264

Figure 18 – FDX CM Total channel Capacity- Use Case 1

The channel capacity Use Case 1, in the figure above: per the capacity numbers on the right edge of the figure above, as the size of the FDX band increases, the maximum downstream capacity increases, from about 4 Gbps to 8.3 Gbps, and the upstream capacity ranges from 1.7 Gbps to 6.2 Gbps

7.2. Additional OFDM channel for 1 or 2 FDX sub-band (Use case-2)

An FDX CM can receive up to four OFDM channels. In this Use case, for the scenarios in which the FDX bands which are 92 MHz, 192 MHz and 384 MHz, the FDX CM has an extra OFDM receiver which can be put to use. In this scenario, the operator introduces an additional 192 MHz OFDM channel, and moves the SC-QAM channels to the 1068-1218 zone, this also reduces the number of SC-QAM channels from 32 to 24.

5	85	108	204	300	396	492	588	684	780	876	972	1068	1164	1218	FDX CM (Mbps)	
															DS	US
FDX																
U-80-OFDMA			FDX DS						DS DS 1.1 OFDM 192			DS DS 0.5 SC-QAM 24*6				
U-80-OFDMA			FDX US						D-192-OFDM		D-192-OFDM		SCQAM-24		5608	8.28
									D-192-OFDM		D-192-OFDM		SCQAM-24		4745	1734
															0	0
U-80-OFDMA			FDX DS	FDX DS					D-192-OFDM		D-192-OFDM		SCQAM-24		6471	8.28
U-80-OFDMA			FDX DS	FDX US					D-192-OFDM		D-192-OFDM		SCQAM-24		5608	1734
U-80-OFDMA			FDX US	FDX DS					D-192-OFDM		D-192-OFDM		SCQAM-24		5608	1734
U-80-OFDMA			FDX US	FDX US					D-192-OFDM		D-192-OFDM		SCQAM-24		4745	2640
															0	0
U-80-OFDMA			FDX DS	FDX DS	FDX DS				D-192-OFDM			SCQAM-32			5738	8.28
U-80-OFDMA			FDX DS	FDX DS	FDX US				D-192-OFDM			SCQAM-32			4857	1734
U-80-OFDMA			FDX DS	FDX US	FDX DS				D-192-OFDM			SCQAM-32			4857	1734
U-80-OFDMA			FDX DS	FDX US	FDX US				D-192-OFDM			SCQAM-32			4012	2640
U-80-OFDMA			FDX US	FDX DS	FDX DS				D-192-OFDM			SCQAM-32			4857	1734
U-80-OFDMA			FDX US	FDX DS	FDX US				D-192-OFDM			SCQAM-32			4012	2640
U-80-OFDMA			FDX US	FDX US	FDX DS				D-192-OFDM			SCQAM-32			4012	2640
U-80-OFDMA			FDX US	FDX US	FDX US				D-192-OFDM			SCQAM-32			3149	3546
															0	0
U-80-OFDMA			FDX DS-192		FDX DS-192				D-192-OFDM		D-192-OFDM		SCQAM-24		8197.44	8.28
U-80-OFDMA			FDX DS-192	FDX US	FDX US				D-192-OFDM		D-192-OFDM		SCQAM-24		6471	2640
U-80-OFDMA			FDX US	FDX US	FDX DS-192				D-192-OFDM		D-192-OFDM		SCQAM-24		6471	2640
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US			D-192-OFDM		D-192-OFDM		SCQAM-24		4745	4452
															0	0
U-80-OFDMA			FDX DS-192		FDX DS-192		FDX DS-192		D-192-OFDM			SCQAM-32			8327	8.28
U-80-OFDMA			FDX DS-192	FDX DS-192	FDX US	FDX US			D-192-OFDM			SCQAM-32			6601	2640
U-80-OFDMA			FDX DS-192	FDX US	FDX US		FDX DS-192		D-192-OFDM			SCQAM-32			6601	2640
U-80-OFDMA			FDX DS-192	FDX US	FDX US	FDX US	FDX US		D-192-OFDM			SCQAM-32			4857	4452
U-80-OFDMA			FDX US	FDX US		FDX DS-192	FDX DS-192		D-192-OFDM			SCQAM-32			6601	2640
U-80-OFDMA			FDX US	FDX US		FDX DS-192	FDX US	FDX US		D-192-OFDM			SCQAM-32		4857	4452
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US		FDX DS-192		D-192-OFDM			SCQAM-32		4857	4452
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US	FDX US	FDX US		D-192-OFDM			SCQAM-32		3149	6264

Figure 19 – FDX Total CM channel Capacity- Use Case 2

The channel capacity Use Case 2, in the figure above: per the capacity numbers on the right edge of the figure above, this use case is the same for the 288,576 MHz FDX-Band scenarios. But for the 96,192 and 384 MHz Scenarios the maximum downstream capacity increases, from about 4 Gbps to 5.6 Gbps, about 4.8 Gbps to 6.4 Gbps, and about 6.6 Gbps to 8.1 Gbps,

7.3. Additional Optional OFDM channel Support (Use case-3)

An FDX CM can receive up to four OFDM channels and optionally it could support one additional FDX Channel. In this scenario, the assumption is the CM can support 5 OFDM channels. In this scenario, the operator introduces an additional 192 MHz OFDM channel, and moves the SC-QAM channels to the 1068-1218 zone, this also reduces the number of SC-QAM channels from 32 to 24. For the 96, 192, 384 MHz FDX Bands (2 sub-band Use cases) the operator can replace the SC-QAM with a more efficient OFDM Channel (150 MHz from 1068 to 1218 MHz), since an OFDM Tuner will be available.

5	85	108	204	300	396	492	588	684	780	876	972	1068	1164	1218	FDX CM (Mbps)	
FDX															DS	US
U-80-OFDMA			FDX US						D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6167	828
U-80-OFDMA			FDX US						D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	5304	1734
U-80-OFDMA			FDX US	FDX DS					D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	7010	828
U-80-OFDMA			FDX US	FDX DS	FDX US				D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6167	1734
U-80-OFDMA			FDX US	FDX DS	FDX DS				D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6167	1734
U-80-OFDMA			FDX US	FDX DS	FDX DS				D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	5304	2640
U-80-OFDMA			FDX US	FDX DS	FDX DS				D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	7334	828
U-80-OFDMA			FDX US	FDX DS	FDX DS	FDX US			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6471	1734
U-80-OFDMA			FDX US	FDX DS	FDX DS	FDX DS			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6471	1734
U-80-OFDMA			FDX US	FDX DS	FDX DS	FDX DS			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	5608	2640
U-80-OFDMA			FDX US	FDX DS	FDX DS	FDX DS			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6471	1734
U-80-OFDMA			FDX US	FDX DS	FDX DS	FDX DS			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	5608	2640
U-80-OFDMA			FDX US	FDX DS	FDX DS	FDX DS			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	5608	2640
U-80-OFDMA			FDX US	FDX DS	FDX DS	FDX DS			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	4745	3546
U-80-OFDMA			FDX US	FDX DS	FDX DS	FDX DS			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM		
U-80-OFDMA			FDX US-192	FDX US-192	FDX US-192				D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	8756	828
U-80-OFDMA			FDX US-192	FDX US-192	FDX US-192				D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	7010	2640
U-80-OFDMA			FDX US	FDX US	FDX US-192				D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	7010	2640
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	5304	4452
U-80-OFDMA			FDX US-192	FDX US-192	FDX US-192	FDX US-192			D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	9923	828
U-80-OFDMA			FDX US-192	FDX US-192	FDX US-192	FDX US	FDX US		D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	8197	2640
U-80-OFDMA			FDX US-192	FDX US	FDX US	FDX US	FDX US		D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	8197	2640
U-80-OFDMA			FDX US-192	FDX US	FDX US	FDX US	FDX US		D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6471	4452
U-80-OFDMA			FDX US	FDX US	FDX US-192	FDX US-192	FDX US		D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	8197	2640
U-80-OFDMA			FDX US	FDX US	FDX US-192	FDX US	FDX US		D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6471	4452
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US-192	FDX US-192		D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	6471	4452
U-80-OFDMA			FDX US	FDX US	FDX US	FDX US	FDX US		D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	D-192-OFDM	4745	6264

Figure 20 – FDX Total CM channel Capacity- Use Case 3

The channel capacity Use Case 3, in the figure above: per the capacity numbers on the right edge of the figure above, the maximum downstream capacity is from about 6.1 Gbps to 9.9 Gbps, and the upstream capacity stays the same, ranging from 1.7 Gbps to 6.2 Gbps

7.4. FDX-L CM (Use case-4)

An FDX-L CM can receive two OFDM channels and two OFDMA Channels.

In this scenario, the operator uses a 192 MHz OFDM channel outside of the FDX band, and uses the SC-QAM channels within the FDX band 108-684 when possible, for the 96, 192, 288, 384 MHz FDX Bands use cases the operator, for the 576 MHz FDX band use case, the operator can move the SC-QAM to the 876-1068 MHz range.

FDX CMs are purpose built CMs designed with hardware and software capable of supporting FDX functionality. FDX-L CMs are D3.1 CMs with limited capabilities for operating within the FDX Band.

The figure below uses the term FDX-L-H for D3.1 High-split CMs which have been software upgrade to use the FDX Band, and the term FDX-L-M for D3.1 Mid-split CMs which have been software upgrade to use the FDX Band. The D3.1 CM supports a downstream lower band edge of 258 MHz. This would apply to FDX-L-H and FDX-L-M CMs. The D.1 CM optionally supports a downstream lower band edge of 108 MHz when the CM is configured to use an upstream upper band edge of 85 MHz or less. This means an FDX-L-M (D3.1 Mid Split CM) may be able to use (FDX) Downstream channels starting at 108 Mhz.

FDX-L CM	5	85	108	204	300	396	492	588	684	780	876	972	1068	1164	1218	FDX CM (Mbps)	
	FDX															DS	US
FDX-L-M	U-80-OFDMA	*	FDX US					DS D3.0 SC-QAM 32*6		DS D3.1 OFDM 192						4012	828
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM						3149	1900
FDX-L-M	U-80-OFDMA	*	FDX US							D-192-OFDM						2770	828
FDX-L-M	U-80-OFDMA	*	FDX US							D-192-OFDM						2770	828
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		D-192-OFDM				5056	1900
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		D-192-OFDM				5056	1900
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM						4012	828
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM						4012	828
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM						4012	828
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM						4012	828
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		D-192-OFDM				5056	1900
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		D-192-OFDM				5056	1900
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		D-192-OFDM				5056	1900
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		D-192-OFDM				5056	1900
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM						4875	828
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM						4875	828
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		D-192-OFDM				4875	1900
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		D-192-OFDM				5056	1900
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM		SCQAM-32				4875	828
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM		SCQAM-32				4875	828
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM		SCQAM-32				4875	828
FDX-L-M	U-80-OFDMA	*	FDX US					SCQAM-32		D-192-OFDM		SCQAM-32				4875	828
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		SCQAM-32				4875	1900
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		SCQAM-32				4875	1900
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		SCQAM-32				4875	1900
FDX-L-H	U-96-OFDMA		FDX US					SCQAM-32		D-192-OFDM		SCQAM-32				3149	1900

FDX sub-band can be optionally used by FDX-L CM, if receiver available
FDX sub-band cannot be used by FDX-L CM

Figure 21 – FDX-L CM CM channel Capacity- Use Case 4

The channel capacity Use Case 4, in the figure above: per the capacity numbers on the right edge of the figure above, the maximum capacity for FDX-L CMs is similar to D3.1 CMs, but with the benefit that an operator may be able to reuse the FDX capacity for currently deployed D3.1 CMs. The downstream capacity ranges from about 4 to 5 Gbps, and the upstream capacity is at ~0.8 Gbps for FDX-L-M CMs (Mid- Split D3.1 CMs) to 1.9 Gbps for FDX-L-H (High Split D3.1 CMs)

Conclusion

Full duplex DOCSIS is a game changing advance in DOCSIS access network technology. It enables a lot more upstream capacity in the network and also allows the operator to dynamically match the customer demand in both the upstream and downstream direction. Each transmission group within the FDX band, gets assigned a unique resource block assignment, to best match the demand from the CMs. The assignment of RBA's has the potential to be complicated and may need to be managed independently. FDX bandwidth along with the legacy(D3.1) downstream and upstream bandwidth easily enables Multi Gigabit service offerings.

Abbreviations

bps	bits per second
ACI	Adjacent Channel Interference
ALI	Adjacent Leakage Interference
CCI	Co-channel interference
CWT	Continuous Wave Tone
D3.1	DOCSIS 3.1
DS	Downstream
ECT	Echo Cancellation Training
FEC	forward error correction
FDX	Full Duplex DOCSIS 3.1
FDX-L	FDX-Limited (functionality)
HFC	hybrid fiber-coax
Hz	hertz
IG	Interference group
RBA	Resource Block Assignment
RxMER	Receive Modulation Error Ratio
SCTE	Society of Cable Telecommunications Engineers
TG	Transmission Group
US	Upstream
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
OUDP	OFDM Upstream Data Profile
ZBL	Zero Bit Loading

Bibliography & References

Full Duplex DOCSIS PHY Layer Design and Analysis for the Fiber Deep Architecture, Richard Prodan, SCTE 2017

Echo Cancellation Techniques for Supporting Full Duplex DOCSIS, Hang Jin, John Chapman, SCTE 2017

Interference Group Discovery for FDX DOCSIS, Tong Liu, SCTE 2017

Accurately Estimating D3.1 Channel Capacity, Karthik Sundaresan, SCTE 2017

DOCSIS 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I14-180509, Cable Television Laboratories, Inc.

DOCSIS 3.1, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I15-180509, Cable Television Laboratories, Inc

FDX DOCSIS Line Extender

Deploying FDX DOCSIS Beyond N+0

A Technical Paper prepared for SCTE•ISBE by

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

Hang Jin

Cisco Distinguished Engineer, Office of Cable CTO
Cisco Systems
170 W Tasman Dr, San Jose, CA 92677
469-255-2666
hangjin@cisco.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Deployment Model	4
1. FDX DOCSIS Background	4
2. Introducing N+M(T) Nomenclature.....	6
3. FDX Node Functional Model	7
4. FDX LE Functional Model	8
5. FDX LE Basic Deployment Model.....	11
6. FDX LE N+2(8) Deployment Model	12
7. FDX LE Maximum Model	15
Technical Considerations.....	16
8. Echo Cancellation Concepts	16
8.1. Analog EC	17
8.2. Digital EC	17
8.3. Reference Signal for EC	17
9. FDX LE Design Guidelines	18
9.1. FDX LE Specifications	18
9.2. Echo Cancellation Performance Target.....	19
Conclusion.....	19
Abbreviations	20
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 – FDX Spectrum	4
Figure 2 – FDX Operation	5
Figure 3 – N + M(T).....	6
Figure 4 – Simplified FDX Node	7
Figure 5 – Functional FDX Node	8
Figure 6 – Simplified Analog Line Extender.....	8
Figure 7 – Functional FDX Node	9
Figure 8 – FDX Noise Model.....	10
Figure 9 – FDX DOCSIS N+0	11
Figure 10 – FDX DOCSIS N+2(2).....	11
Figure 11 – FDX DOCSIS N+2(8).....	12
Figure 12 – FDX LE Partial Decomposition	16
Figure 13 – FDX LE with Specs.....	18

List of Tables

Title	Page Number
Table 1 – FDX DOCSIS 576 MHz Band Throughput.....	13
Table 2 – FDX DOCSIS US Full Band Throughput	14
Table 3 – FDX DOCSIS DS Full Band Throughput	14
Table 4 – Modulation Order vs Amp Count.....	15

Introduction

The 42 MHz (65 MHz for Europe) upstream return path is running out of bandwidth. The 90 to 100 Mbps it has today is enough to support the 1 Gbps DOCSIS downstream of today, but it is not enough to support the oncoming 10 Gbps DOCSIS downstream of tomorrow.

That means that the entire HFC plan will need to be upgraded to a new return path. There are multiple options available: 85 MHz and 204 MHz using classic FDD and up to 684 MHz using FDX DOCSIS. 85 MHz only offers 400 Mbps of bandwidth which could support a 4 Gbps downstream, assuming a 10:1 ratio in bandwidth, but is still not enough for a full 10 Gbps downstream.

FDX DOCSIS offers the most bandwidth possible in the return path with the least impact on forward path bandwidth. However, the cost of deploying FDX DOCSIS today is high as FDX DOCSIS currently is specified to work on an N+0 HFC plant whereas the 204 MHz does not require N+0. However, 204 MHz does require upgrading every single amp as well as the node. What if FDX DOCSIS could be deployed by also just upgrading the amps and nodes, rather than trenching new fiber for new node locations? This would make FDX DOCSIS be much closer to cost parity with 204 MHz. In fact, it could save operators billions of dollars. [1]

This white paper will show how FDX DOCSIS can be taken beyond N+0. It will do so by describing a deployment model that is scalable and allows the cable operator to trade-off cost versus performance.

Deployment Model

1. FDX DOCSIS Background

Full Duplex (FDX) is a new option in DOCSIS 3.1 that was introduced via an ECN in DOCSIS 3.1 [2] [3] in October 2017. FDX DOCSIS is the result of highly innovative work in the cable industry [4][5][6][7][8][9][10].




Figure 1 – FDX Spectrum

The fundamental premise of FDX DOCSIS is to share a common frequency spectrum between the Hybrid Fiber Coax (HFC) forward path and the reverse path. Note that when referring to the HFC fiber or coax

paths, the terms “forward path” and “reverse path” are used. In DOCSIS, there is a “downstream” and an “upstream” path. With DOCSIS and HFC technologies being blended with technologies like Remote PHY, these terms tend to get used interchangeably.

The frequency spectrum sharing is shown in Figure 1. The spectrum plan for FDX DOCSIS is:

- 5 to 42 MHz: legacy upstream spectrum for DOCSIS 3.0/2.0/1.1 ATDMA
- 42 to 85 MHz: new upstream spectrum for ATMDA and/or OFDMA
- 85 to 108 MHz: cross-over band. Also used for OOB.
- 108 MHz to 684 MHz: shared downstream and upstream spectrum
- 684 MHz to 1218 MHz: downstream spectrum for DOCSIS and legacy MPEG Video

It is this shared spectrum that is unique for FDX DOCSIS. The current HFC plant is built on the principle of what is called Frequency Division Duplex (FDD), where the direction of the HFC plant is defined at a particular frequency. All optical nodes, coax amplifiers and line extenders (LE) today are built on those principles and contain diplexers that enforce this frequency division. FDX DOCSIS basically uses an echo cancellation mechanism instead of a diplexer.




Figure 2 – FDX Operation

Here is a brief recap of how FDX DOCSIS works.

1. We echo cancel at the CMTS PHY.
 - This is a PHY operation that works beneath the DOCSIS 3.1 protocol.
 - The CMTS PHY is located in the Remote PHY Device (RPD) in the node
 - The RPD receives the upstream signal that also has the downstream signal combined with it. The downstream signal is effectively now noise in the upstream path. The downstream signal is higher in power than the upstream signal, so there is a negative signal-to-noise (SNR) ratio.
 - The RPD contains an echo canceller (EC) that subtracts the downstream signal from the upstream signal. The EC effectively attenuates the noise generated by the downstream

signal and creates enough positive SNR that 2K modulation will work.

2. We measure and sort CMs into interference groups (IG) and IGs into transmission groups (TG).
 - This is a MAC operation.
 - CMs use an EC to eliminate self-interference but cannot use an EC to eliminate interference from neighbors.
 - We measure the attenuation between all CMs and sort them into IGs. CMs within an IG can hear each other; CMs in different IGs cannot hear each other.
 - There are usually too many IGs to separately schedule, so we combine them into two or three TGs for scheduling purposes.
3. We use FDD within a TG so that those CMs do not interfere with each other.
 - Each TG is like a small DOCSIS 3.1 MAC domain with its own FDD frequency plan.
4. We overlap TGs in frequency and time so that 100% of the spectrum and 100% of the timeline are used for both DS and US.
 - Each TG is like a DOCSIS 3.1 island with its own unique FDD frequency plan
 - The sum of all the separate FDD frequency plan is an FDX frequency plan.
 - In the example in Figure 2, there are two TGs, each with the DS and US frequency spectrums swapped. When added together at the CMTS, it appears as if the entire DS and US spectrum is simultaneously used.

FDX DOCSIS was originally specified to work on a node plus zero amplifier (N+0) system where the node was redesigned to accommodate FDX DOCSIS. If FDX DOCSIS could work in say an N+2 system, then there is the potential to dramatically lower installation costs as there will be less additional nodes and less fiber pulled to those nodes.

In this white paper, we will look at redefining the HFC amplifier and HFC line extender (LE). In the HFC world, an HFC LE is a single output amplifier, where as an HFC amplifier has two or three outputs. Although this white paper will focus on the HFC LE use case, all these principles are extensible to an HFC amplifier.

2. Introducing N+M(T) Nomenclature




Figure 3 – N + M(T)

Before proceeding into the details of an FDX LE, we need to look at the HFC plant that the LE will be located in. Classically, an HFC plant is described by how many amplifiers exist in a chain from the node to a home. Thus, an N+2 HFC plant had a depth of two amplifiers from the node to the home.

For FDX DOCSIS, we will soon see that it is important to know not only the depth of the amplifier chain, but the total number of amplifiers as well. An example is shown in Figure 3. Both systems are N+2 in that there are two amplifiers in any one downstream path. However, in the second example, the downstream path has been split and there is actually a total of six amplifiers, although only two in any one downstream path. The system on the right uses more amplifiers to be able to pass more homes but maintains the downstream signal quality by keeping the path length to two amplifiers.

While the downstream path is similar, from an upstream path viewpoint, it is quite different. There is a classic phenomenon in HFC called noise funneling where every extra upstream path adds more noise, and the total noise that is at the upstream path in the node is related to the sum of all the noise on all the paths.

One of the design criteria for FDX LE deployment will be to limit the total number of amplifiers and thus the total number of paths associated with those amplifiers. A proposed nomenclature to represent both the amplifier depth and the total number of amplifiers is:

$$N + M(T)$$

Where:

N = Node

M = amplifier depth from node to home

T = total number of amplifier outputs from a node to all homes serviced by that node.

So in the example in Figure 3, the system on the left is N+2(8) since there is two amps in any chain, but eight amp outputs in total, while the system on the right is N+2(24) since it is also two amps in any chain, but has twenty-four amps per node.

3. FDX Node Functional Model

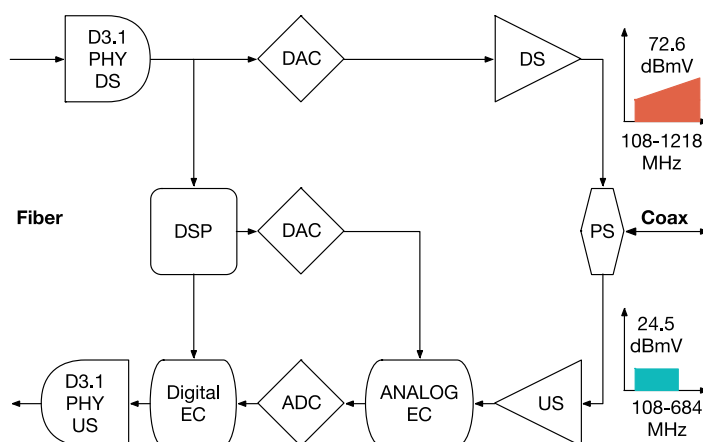


Figure 4 – Simplified FDX Node

A simplified diagram of an FDX node is shown in Figure 4. This diagram is meant to include the signal path that is comprised of the Ethernet over Optics, the RPD and the HFC node analog electronics.

The downstream path consists of the DOCSIS 3.1 modulation followed by a digital to analog converter and an amplifier path to a power splitter. Note the use of a power splitter instead of a diplexer. A diplexer would enforce an FDD frequency plan whereas a power splitter permits both the downstream and upstream spectrum to coexist.

The upstream path is received and buffered. At this point, the upstream path contains both the upstream and downstream signals. The downstream signal power is higher than the upstream signal power. This has an impact on the dynamic range of the ADC. To address this practical constraint, an analog EC is used to lower the interfering downstream signal level to be equal to or lower than the upstream signal level. The signal is then digitized and a digital echo canceller removes most of the rest of the interfering downstream signal such that the upstream signal can be recovered by the DOCSIS 3.1 upstream receiver.




Figure 5 – Functional FDX Node

Even the simple model of Figure 4 can get overwhelming and contains more information that is needed for basic understanding and deployment analysis. Figure 5 represents a truly barebones functional model of an FDX Node. There is a downstream path, an upstream path, and an echo canceller separates the upstream spectrum from the downstream spectrum.

Now with this functional model of an FDX node, let's move on to an FDX LE.

4. FDX LE Functional Model

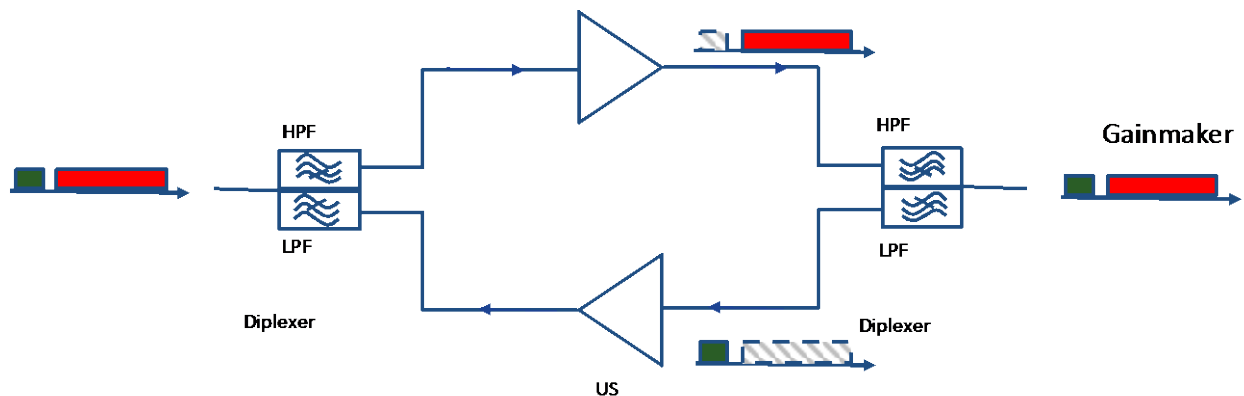


Figure 6 – Simplified Analog Line Extender

A simplified diagram of a classic HFC amplifier is shown in Figure 6. There is an amplified forward path and an amplified reverse path. A diplexer separates both paths on both ports. The amplifier is FDD based so that the forward and reverse paths do not interfere with each other and can be separated with passive filters.

What we have learned in an FDX node is that an echo canceller along with a power splitter can convert a bi-directional coax path into two separate forward and reverse amplified paths. If that is true, could two echo canceller circuits be placed back-to-back to create an FDX amplifier? To keep the diagram simple and to focus on the outcome, Figure 7 shows a functional model of an FDX LE.




Figure 7 – Functional FDX Node

Does this work? In theory, if the echo cancellation was perfect, it should work just fine. In practice, what is required?

We can see in Figure 7 that the downstream signal leaks into the upstream path and results in noise that is attenuated by the reverse path echo canceller. By placing a forward path echo canceller on the northbound port of the amp, some of the upstream signal is reflected into the downstream path and results in noise that is attenuated by the downstream echo canceller. Note that some of noise in the downstream path could actually originate from the downstream path; that is downstream signal power that has been circulated and attenuated by both the forward and reverse path echo cancellers.

But what if the echo cancellers are not perfect? What if that energy kept circulating around the forward and reverse paths? If that happened, the whole system would become unstable. So, the echo cancellation performance must be good enough to ensure that this does not happen. As we will see in the section on Technical Considerations, that means that the echo cancellation level in each direction has to be on the order of 50 dB or more.

Assuming this can be built and that it works, what is the impact on performance? The answer is that the EC just becomes another noise source that needs to be managed and accounted for.




Figure 8 – FDX Noise Model

A simple noise source model is shown in Figure 8. This works for both the forward and reverse paths.

Each amplifier in the signal path contributes a noise floor which is determined by the thermal noise of the amplifier and is measured by the noise factor (NF). NF is one of the reasons that amplifier chains are limited to no more than 5 amplifiers in order to keep the noise accumulated noise floor low. Note that the thermal noise in each amplifier is not correlated, so the noise does not directly add in terms of dB.

Next is ingress noise. In the upstream, this tends to be lower frequency noise that can be across frequency or impulse noise and tend to come from homes. In the downstream, the ingress noise tends to be LTE energy (~700 to 800 MHz). This energy is very dependent on noise source location and bad connector locations, so the noise again is not always additive.

The EC noise in the upstream is always from the downstream path. It is the same signal, just delayed by different amounts at each EC. So, the energy is more additive, although the delay of each echo will impact that additivity.

For the forward path, the signal originates in the node. In the reverse path, the signal originates in the cable modem (CM). There is a fundamental difference in these paths when it comes to noise accumulation. The HFC forward path is a series of splitters. Thus the noise in any one path from the node to the home is contributed by that singular path. The HFC reverse path is a series of combiners. Thus the noise in the combined upstream path that arrives at the node is the combination of all noise sources in the upstream from all upstream paths. This is referred to as noise funneling.

That means that an N+2(8) system as shown in Figure 3 would have two amplifier noise sources in any one downstream path but eight amplifier noise sources in the singular upstream path. The N+2(24) system shown in Figure 3 would also have two amplifier noise sources in any one downstream path but 24 amplifier noise sources in the singular upstream path.

The exact noise models, their frequency dependencies, delay dependencies and additive properties will be left for future papers. For now, let's assume that the thermal and ingress noise problems are already accommodated by the current network designs and that EC noise is additive. For the sake of example, let's assign a value of 1.5 dB of noise contribution per EC.

Here's the thing. What if we allow the system to get noisier in order to lower the cost of a system deployment? More noise can be managed by lower the modulation order which will lower throughput. Lower modulation orders mean lower throughput. Lower throughput results in lower cost. What is the

right throughput versus cost ratio? That might be different for every operator and may differ from year to year.

We now have a scenario where cost can be traded off against performance.

5. FDX LE Basic Deployment Model



Figure 9 – FDX DOCSIS N+0

In the functional diagram Figure 9 we have an N+0 system where a node has an RPD with an FDX EC in the upstream direction. The shared spectrum is passed between the node and the CM over a passive networks of splitters and combiners known as taps (not shown). The CM has a single EC that separates the spectrums back into downstream and upstream paths. The total number of ECs in the path are two and there are two open loop EC paths.

This system has been carefully specified and designed to work in the DOCSIS 3.1 specifications. The performance targets are 4K modulation in the downstream and 2K modulation in the upstream.




Figure 10 – FDX DOCSIS N+2(2)

In the functional diagram of Figure 10, two FDX line extenders have been inserted between the node and the CM. We are using the N+M(T) nomenclature introduced earlier to state that the amplifier depth is two and the total number of amplifiers is also 2. The total number of ECs in the path has been increased to six and there are now two closed loop EC paths and two open loop EC paths.

In this simple example, if we assume the ECs are doing their job, the forward and reverse paths will be intact and working. However, the noise floor will be elevated. In this basic example, we can see three EC noise contributions in the downstream path and three EC noise contributions in the upstream path. If each noise contribution is 1.5 dB, then the total noise contribution is 4.5 dB. If we subtract out the noise contributions at each end that are already included in the FDX N+0 design, the increased noise floor is 3 dB in each direction.

A 3 dB increase in noise floor can be managed by reducing the modulation order by one. *Thus, when compared to an FX N+0 system that is expected to work at 4K modulation in the downstream and 2K modulation in the upstream, this FDX N+2(2) system should be expected to work at 2K modulation in the downstream and 1K modulation in the upstream.*

That is a nice cost versus performance trade-off. How does this scale? What more trade-offs can be made?

6. FDX LE N+2(8) Deployment Model




Figure 11 – FDX DOCSIS N+2(8)

Let's now scale this example to a full four port node with two line extenders on each port for a total of eight amplifiers. The N+2(8) system is shown in Figure 11. The node itself is a splitting/combining system. The downstream spectrum is derived from a fiber input and split out to four outputs and the four return paths are combined into a single optical return path.

There may be some new examples to a node split/combine model when an RPD is present. If the RPD is a 1x1 configuration, containing one downstream port and one upstream port, then the split/combine model is the same as the analog optical node. If the RPD is a 2x2 or a 4x4 where node ports are truly independent of each other, then the node will perform as a segmented node and the four ports will not be combined. This analysis will assume a non-segmented node with a 1x1 RPD.

Doing the math for the upstream, there are eight amps and hence eight ECs dumping FDX related noise into the upstream path in addition to the one EC in the RPD in the node that is already accounted for. Eight noise sources at 1.5 dB each, assuming direct addition, would be 12 dB of increased noise floor. Dividing that by 3 dB per modulation level says that the modulation level must be taken down by four

orders. The starting modulation for N+0 systems is 2K QAM which has an order of 11. Subtracting four results in a modulation order of 7 which is 128K QAM. This is shown in formula 1.

$$\text{US Modulation Order} = 11 - \text{Ceiling}[\text{Total Amps} * \text{Amp Noise} / 3 \text{ dB}] \quad \text{formula \#1}$$

$$\text{US Modulation Order} = 11 - \text{Ceiling}[8 * 1.5 / 3] = 7 \quad (128 \text{ QAM})$$

Doing the math for the downstream, we only care about amplifier depth and the starting point is 4K QAM which has a modulation order of 12.

$$\text{DS Modulation Order} = 12 - \text{Ceiling}[\text{Amp depth} * \text{Amp Noise} / 3 \text{ dB}] \quad \text{formula \#2}$$

$$\text{DS Modulation Order} = 12 - \text{Ceiling}[2 * 1.5 / 3] = 11 \quad (2K \text{ QAM})$$

To translate these modulation order to throughputs, we have created a series of tables. Table 1 shows the throughput in the FDX DOCSIS band from 108 MHz to 684 MHz. The calculations in the table assume 20% total overhead between raw bits and packets and can be used for both the upstream and downstream direction. Keep in mind that for N+0 FDX operation today, the starting point is a modulation order of 11 for the upstream (2K QAM) and 12 for the downstream (4K QAM).

Table 2 is for a full upstream band and adds in 400 Mbps, 100 Mbps for 5 to 42 MHz and 300 Mbps for 42 to 85 MHz. Table 3 is for a full downstream band and adds in 4.91 Gbps, 3.68 Gbps for two more 192 MHz OFDM channels running at 4K QAM and 1.23 Gbps for 32 channels of 256-QAM, both with 20% overhead.

Table 1 – FDX DOCSIS 576 MHz Band Throughput

FDX Band Gbps		FDX Frequency Band, 108 MHz to ...						MHz
		204	300	396	492	588	684	
M Order	QAM	1	2	3	4	5	6	96 MHz Blocks
12	4096	0.92	1.84	2.76	3.69	4.61	5.53	
11	2048	0.84	1.69	2.53	3.38	4.22	5.07	
10	1024	0.77	1.54	2.30	3.07	3.84	4.61	
9	512	0.69	1.38	2.07	2.76	3.46	4.15	
8	256	0.61	1.23	1.84	2.46	3.07	3.69	
7	128	0.54	1.08	1.61	2.15	2.69	3.23	
6	64	0.46	0.92	1.38	1.84	2.30	2.76	
5	32	0.38	0.77	1.15	1.54	1.92	2.30	
4	16	0.31	0.61	0.92	1.23	1.54	1.84	
3	8	0.23	0.46	0.69	0.92	1.15	1.38	
2	4	0.15	0.31	0.46	0.61	0.77	0.92	

Overhead per OFDM Channel: 20%

Base Throughput added: 0 Mbps

Table 2 – FDX DOCSIS US Full Band Throughput

Total US Gbps		FDX Frequency Band, 108 MHz to ...						MHz
		204	300	396	492	588	684	
M Order	QAM	1	2	3	4	5	6	96 MHz Blocks
11	2048	1.24	2.09	2.93	3.78	4.62	5.47	
10	1024	1.17	1.94	2.70	3.47	4.24	5.01	
9	512	1.09	1.78	2.47	3.16	3.86	4.55	
8	256	1.01	1.63	2.24	2.86	3.47	4.09	
7	128	0.94	1.48	2.01	2.55	3.09	3.63	
6	64	0.86	1.32	1.78	2.24	2.70	3.16	
5	32	0.78	1.17	1.55	1.94	2.32	2.70	
4	16	0.71	1.01	1.32	1.63	1.94	2.24	
3	8	0.63	0.86	1.09	1.32	1.55	1.78	
2	4	0.55	0.71	0.86	1.01	1.17	1.32	

Overhead per OFDM Channel: 20%

Base Throughput added: 400 Mbps 20-85 MHz

Table 3 – FDX DOCSIS DS Full Band Throughput

Total DS Gbps		FDX Frequency Band, 108 MHz to ...						MHz
		204	300	396	492	588	684	
M Order	QAM	1	2	3	4	5	6	96 MHz Blocks
12	4096	5.84	6.76	7.68	8.60	9.52	10.44	
11	2048	5.76	6.60	7.45	8.29	9.14	9.98	
10	1024	5.68	6.45	7.22	7.99	8.75	9.52	
9	512	5.61	6.30	6.99	7.68	8.37	9.06	
8	256	5.53	6.14	6.76	7.37	7.99	8.60	
7	128	5.45	5.99	6.53	7.06	7.60	8.14	
6	64	5.37	5.84	6.30	6.76	7.22	7.68	
5	32	5.30	5.68	6.07	6.45	6.83	7.22	
4	16	5.22	5.53	5.84	6.14	6.45	6.76	
3	8	5.14	5.37	5.61	5.84	6.07	6.30	
2	4	5.07	5.22	5.37	5.53	5.68	5.84	

Overhead per OFDM Channel: 20%

Base Throughput added: 4914 Mbps (2 OFDM ch, 32 QAM ch)

In the N+2(8) example, the upstream could:

- reduce modulation order from 11 to 7 in the FDX band, so from 2K QAM to 128K QAM
- in the FDX band, it would reduce from 5.07 Gbps to 3.23 Gbps, a 36% reduction
- in full band, it would reduce from 5.47 Gbps to 3.63 Gbps, a 34% reduction

and the downstream could:

- reduce the modulation order from 12 to 11 in the FDX band, so 4K QAM to 2K QAM
- in the FDX Band, it would reduce from 5.53 Gbps to 5.07 Gbps, a 9.2% reduction
- in the DS full band, it would reduce from 10.44 Gbps to 9.98 Gbps, a 4.4% reduction

Here is where the cable operator can make trade-offs between cost and performance. If the service offering for the DOCSIS Service Group only needed say 5 Gbps down and 3 Gbps up, then this extra bandwidth was not needed and money can be saved by deploying a N+2(8) configuration rather than a N+0 configuration. Conversely, if the top performance is needed, then an N+0 network would have to be deployed.

Bear in mind that that these are theoretical numbers. The amplifier noise figure for FDX is purely a budget number and has not been measured. It could be better or worse. It has also not been proven or disproven if the noise is strictly additive. There is also unknown margin in the N+0 design that has not been accounted for. And one of the most interesting parts is that since DOCSIS 3.1 manages modulation profiles based upon measured SNR, DOCSIS 3.1 will automatically manage and optimize this for us.

7. FDX LE Maximum Model

This is always the next question – what is the maximum number of line extenders that could be deployed?

The design limit is obviously the upstream as it degrades faster since it is a function of the total number of amps instead of the number of amps in cascade. The answer in theory would be to run the FDX upstream spectrum at its lowest modulation level which is QPSK and reverse calculate the number of amps.

Also, it might be prudent to add an extra 3 dB of margin just to allow for SNR room when the network changes. The FDX network is constantly cancelling echoes. Those echoes can change due to physical conditions such as temperature, wind and vibration. So, when the plant changes and there are a large number of amps each changing at once, or at slightly different times, does this cause brief periods of increased noise floor?

Table 4 – Modulation Order vs Amp Count

Mod Order	Modulation	Throughput	#LE	#LE + 3dB
11	2K	5.47 Gbps	0	0
10	1K	5.01 Gbps	2	0
9	512	4.55 Gbps	4	2
8	256	4.09 Gbps	6	4
7	128	3.63 Gbps	8	6
6	64	3.16 Gbps	10	8
5	32	2.70 Gbps	12	10
4	16	2.24 Gbps	14	12
3	8	1.78 Gbps	16	14
2	QPSK	1.32 Gbps	18	16

The results are shown in Table 4. The approach uses the basic theory of the paper. It assumes a reference FDX N+0 system at 2K QAM as a baseline and then drops one order of modulation for every two LEs added to the system. The far right-hand column allows an additional 3 dB of system noise margin.

So, in theory, if 1.32 Gbps was sufficient performance for the deployment time period under review, then 16 to 18 amplifiers could be deployed. That could then support an N+2(16) system but would not support the N+2(24) system that was shown in Figure 3. Now if the system in Figure 3 was a three-port node with six amps per port, that would fit without the extra 3 dB of noise margin. Or, if N+2(16) was a four-port node with four LE (or amp equivalents) per port total, that could very well work.

When extending this concept to a two or three port line amp instead of a port extender, the question to ask is how many ECs are in a two or three port amplifier. Typically, the EC will be associated with the single north bound port and the two or three southbound ports are just separate amplified paths. In that case, the two or three port amp is equivalent to a LE load. If, however, the amp was designed with the EC on each southbound port, then there would be two to three ECs per amp which would then be equivalent to two or three LE loads. The net result is to count the EC noise sources, not the amplifier outputs.

Technical Considerations

8. Echo Cancellation Concepts




Figure 12 – FDX LE Partial Decomposition

Full Duplex transmission and echo canceller technology has been around in other technical venues such as telephony for a long time [11]. If we choose to re-use those concepts and definitions, there are a few worth noting.

Echo Return Loss (ERL)

$$\text{ERL (dB)} = \text{Original Signal Level (dBm)} - \text{Echo Signal Level (dBm)}$$

This is a measure in dB of the signal that is reflected back when compared to the signal being sent out. The correct use would be from a reference point in the downstream path to a reference point in the upstream prior to the EC.

Echo Return Loss Enhancement (ERLE)

ERLE is the additional echo attenuation provided by the echo canceller.

These values are shown in Figure 12. The point of having two separate values is the ERL embodies the reflections from the power splitter, the shell housing and the outside plant. ERLE is focused on the EC itself. The downstream signal will get attenuated by the combination of ERL and ERLE to create an upstream noise floor. As explained above, back-to-back echo cancellations need be implemented in FDX LE to cancel or suppress the echoes on both input and output ports. There are two types of EC techniques.

8.1. Analog EC

An analog EC is used to cancel out the echoes in the analog domain before the ADC. Conventionally, the analog EC will take a copy of the transmitted signal and manipulate its phase and magnitude to generate a canceling signal that has the same magnitude but 180 degrees out-of-phase from the echo. This canceling signal is then added to the receiver path to cancel out the echo. As there will be multiple echoes coming from multiple sources (FDX LE output connectors, taps, etc.), multiple cancelling signals need be generated, one for each echo. All these need to be done in the analog domain.

The analog EC used in FDX LE is actually a hybrid solution. The cancelling is still in the analog domain before the ADC to enable the benefits of analog EC, but the cancelling signal is generated in digital domain first and then converted into analog domain through a digital-to-analog converter (DAC). All the delays and magnitudes are computed and set in the digital domain through EC digital signal precessing (DSP).

8.2. Digital EC

Digital EC cancels out the echoes in the digital domain after ADC. After the echoes pass through the ADC and are converted into bits in digital domain, their magnitude and phase can be computed, and the cancelling signal can be generated from the transmitted reference signal with the proper magnitude and phase and subtracted from the received signal. Unlike analog EC which must be implemented in time domain, digital EC can be implemented in either time domain or frequency domain or combination of both.

8.3. Reference Signal for EC

The cancelling signal is generated from the transmitted signal with the proper magnitude and phase computed from the echoes embedded in the received signal. The transmitted signal used to generate the cancelling signal is called the EC reference signal. The theoretic base of the EC (both analog and digital EC) is that all the reflections are true copies of the same transmitted signal, just with various magnitudes and phases, depending on how the echoes are generated.

The same EC algorithm could be implemented for both input and output ports (back-to-back EC). One just needs to keep in mind that the reference signal used is different. For the input port, the transmitted signal is the US, and the received signal is the DS. Thus the reference signal for the input port EC will be the US signal. On the contrary, for the output port, the transmitted signal is the DS, the received signal is US, thus the reference for the output port EC is the DS signal.

9. FDX LE Design Guidelines




Figure 13 – FDX LE with Specs

There is no design specification for an FDX LE yet as this is a new device, so this white paper will propose some rough guidelines to help drive feasibility. The final values may vary.

9.1. FDX LE Specifications

Forward Path

- The downstream path frequency response is 88 MHz to 1218 MHz with 19 dB tilt
- The input power level is 0 dB and flat.
- If the LE had the same output power level as a FDX node, it would be 72.6 dBmV.
- Power could be dropped by 3.5 dB from 1002 to 1218 MHz if the output total composite power (TCP) is reduced to 71 dBmV
- Gain is 48 dB

Reverse Path

- US Frequency response
 - 10 MHz to 85 MHz, flat
 - 108 MHz to 684 MHz, with 6 dB up-tilt
 - 85 to 108 is notched out of the upstream path with DSP
- TCP output is 55 dBmV
 - Assuming to be the similar but less than a CM upstream output
- Gain is 27 dB

9.2. Echo Cancellation Performance Target

One of the design criteria for FDX LE is that the total close loop echo suppression must be greater the total close loop gain, as expressed as follows:

$$\text{ERL1} + \text{ERL2} + \text{ERLE1} + \text{ERLE2} > \text{G1} + \text{G2} + \text{A} \quad \text{formula \#3}$$

Where ERL1 and ERL2 are, respectively, the Echo Return Losses of input and output ports, ERLE1 and ERLE2 are, respectively, the EC gains of input and output ports. G1 and G2 are, respectively, the forward and reverse path gains. A is the design margin. A needs to be >10dB.

If we design the gains of FDX LE as the same as a legacy LE, where the forward gain is <= 50dB and the reverse gain is <=30dB; And set A=10, we have:

$$\text{ERL1} + \text{ERL2} + \text{ERLE1} + \text{ERLE2} > 90 \text{ dB} \quad \text{formula \#4}$$

Then assume ERL1 and ERL2 are around 15dB (reflection from input/output connectors, taps), we have:

$$\text{ERLE1} + \text{ERLE2} > 60 \text{ dB} \quad \text{formula \#5}$$

With the state-of-art ADC/DAC, it is expected that each echo cancellation can achieve 50dB echo suppression (ERLE1/ERLE2 >= 50dB). Thus, 60 dB closed-loop echo suppression can be readily achieved.

Using the numbers in the above example and inserting into formula #3 and solving for A,

$$\text{A} = -15 \text{ dB} - 15 \text{ dB} - 50 \text{ dB} - 50 \text{ dB} + 48 \text{ dB} + 27 \text{ dB} = 55 \text{ dB} \quad \text{formula \#6}$$

This is much higher than the 10 dB we were looking for, so there appears to be plenty of design margin.

In practice, for this to work in a field deployment, the performance of each echo canceller must be:

- across the entire spectrum of FDX operation and for every subcarrier within that spectrum,
- across operating temperature, so -20 C to + 85 C (internal node temperatures),
- across device variations,
- across all deployment time,
- when switching EC coefficients.

Conclusion

In this white paper, we have introduced a technology and a deployment methodology to deploy FDX DOCSIS in networks with amplifier depths greater than N+0. We showed that a reverse path echo canceller that is in an FDX node today could be combined with a forward path echo canceller to create an FDX line extender (single port) or FDX amplifier (multi-port).

As each EC restores a forward or reverse path, it also introduces noise into that path. The existence of multiple ECs can mean degraded performance, since the noise in the downstream is the sum of all noise in a particular path and the noise in the upstream is the sum of all ECs in the upstream. That noise could be accommodated by adjusting the performance of DOCSIS 3.1 channels within the FDX band. Thus, the FDX LE and the FDX amp allows an operator to tradeoff deployment costs by having more amplifier stages past a node against throughput performance of that node.

We introduced a new terminology for measuring the depth and size of an HFC plant, N+M(T), where N referred to a node, D to the depth of the amplifiers, and T for the total number of amplifiers attached to the node.

A conservative example was given where an N+2(8) system. Assuming each LE inserted 1.5 dB of additional noise, and an extra 3 dB of operating margin, then the upstream throughput for FDX might be reduced from about 5 Gbps to about 2.5 Gbps which is still much higher than the 100 Mbps of today.

This example may be a worst case scenario since:

- The FDX LE may contribute much less than 1.5 dB of noise (these have not been built yet)
- The noise may not be strictly additive.
- There is design margin in the current FDX designs that has not been used in these calculations
- DOCSIS 3.1 Profile Management Application (PMA) will optimize the downstream and upstream performance.

The example could be a best case scenario as well as the HFC plant is not always a stable environment. When the HFC plant moves due to wind and temperature changes, the echoes may change as well which may then require a change in the EC coefficients. With more ECs, there are more coefficient changes in more places. Thus, there is a living, breathing HFC plant to be managed.

What is now obvious is that it is possible to extend beyond N+0 to at least N+1 and some N+2 HFC systems, and maybe even beyond that. This will lower the cost of an FDX deployment and allow a cost/performance trade-off to be made by cable operators on their network design.

Abbreviations

ADC	analog to digital converter
CM	cable modem
CMTS	cable modem termination system
DAC	digital to analog converter
dB	decibel
DOCSIS	Data Over Cable System Interface Specification
DS	downstream
EC	echo canceller
ERL	echo return loss
ERLE	Echo return loss enhancement
FDD	frequency division duplex
FDX	full duplex
Gbps	gigabits per second
HFC	hybrid fiber-coax
HPF	high pass filter
Hz	Hertz
IG	interference groups (used in FDX DOCSIS)
ISBE	International Society of Broadband Experts
LE	line extender (one port HFC amplifier)
LPF	low pass filter
Mbps	Megabits per second

N+0	node plus zero amplifiers
N+M(T)	node plus amplifiers with a depth of M and a total of T
NF	noise factor
OFDM	orthogonal frequency division multiplexing
PMA	Profile management application (used in DOCSIS 3.1)
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying
RPD	Remote PHY Device
SCTE	Society of Cable Telecommunications Engineers
SNR	signal to noise ratio
TCP	total composite power
TG	transmission groups (used in FDX DOCSIS)
US	upstream

Bibliography & References

- [1] *HFC Evolution – The Best Path Forward*, by Nader Foroughi, Shaw Communications, SCTE Expo 2018 Fall Technical Forum, October 22, 2018
- [2] *DOCSIS 3.1 Physical Layer Specification*, CM-SP-PHYv3.1-I15-180926, CableLabs, September 26, 2018, <https://apps.cablelabs.com/specification/>
- [3] *DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv3.1-116-180926, CableLabs, September 9, 2018, <https://apps.cablelabs.com/specification/>
- [4] *Full Duplex DOCSIS*, by John T. Chapman & Hang Jin, Cisco Systems, INTX 2016 Spring Technical Forum, May 16, 2016, <https://www.nctatechnicalpapers.com/Paper/2016/11-Jin>
- [5] *Interference-Aware Spectrum Resource Scheduling for FDX DOCSIS*, by Tong Liu, John T Chapman and Hang Jin, Cisco Systems, SCTE Journal, August 11, 2016
- [6] *Interference Group Discovery for FDX DOCSIS*, by Tong Liu, Cisco Systems, SCTE Expo 2017 Fall Technical Forum, October 17, 2017, <https://www.nctatechnicalpapers.com/Paper/2017/2017-interference-group-discovery-for-fdx-docsis>
- [7] *Echo Cancellation Techniques for Supporting Full Duplex DOCSIS*, by Hang Jin & John Chapman, Cisco Systems, SCTE Expo 2017 Fall Technical Forum, October 17, 2017, <https://www.nctatechnicalpapers.com/Paper/2017/2017-echo-cancellation-techniques-for-supporting-full-duplex-docsis>
- [8] *Characterization of Spectrum Resource Scheduling in FDX DOCSIS*, by Tong Liu, Cisco Systems, SCTE Expo 2018 Fall Technical Forum, October 22, 2018
- [9] *Full Duplex DOCSIS Technology over HFC Networks*, by Belal Hamzeh, CableLabs, INTX 2016 Spring Technical Forum, May 16, 2016, <https://www.nctatechnicalpapers.com/Paper/2016/2016-full-duplex-docsis-technology-over-hfc-networks>

[10] *Full Duplex DOCSIS PHY Layer Design and Analysis for the Fiber Deep Architecture*, Richard S Prodan, Broadcom, SCTE Expo 2017 Fall Technical Forum, October 17, 2017, <https://www.nctatechnicalpapers.com/Paper/2017/2017-full-duplex-docsis-phy-layer-design-and-analysis-for-the-fiber-deep-architecture>

[11] *Recommendation ITU-T G.168 Digital network echo cancellers*, ITU-T, April 2015, https://www.itu.int/rec/dologin_pub.asp?lang=e&id=T-REC-G.168-201504-I!!PDF-E&type=items

Guaranteeing Seamless 4K OTT Content Delivery

A Technical Paper prepared for SCTE•ISBE by

Jos Delbar

Director Product Management Managed Services
Technicolor
Prins Boudewijnlaan 47, B-2650 Edegem, Belgium
+32 3 443 64 16
Jos.Delbar@technicolor.com

Bart Vercammen

CTO Customer Premises Equipment BU
Technicolor
Prins Boudewijnlaan 47, B-2650 Edegem, Belgium
+32 3 443 6 519
Bart.Vercammen@technicolor.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Guaranteeing Seamless 4K OTT Content Delivery	4
1. A burning platform	4
2. Market Overview	5
2.1. Streaming Video Players.....	5
2.2. UHD Screens and Content.....	7
2.3. UHD Video Streams.....	7
2.4. Wi-Fi for Multimedia Access.....	8
3. Wi-Fi Performance Requirements For A Seamless UHD Experience	9
3.1. Quantifying and Qualifying Wi-Fi Performance and Performance Losses	9
3.2. Throughput	11
3.3. Sustained Throughput Vs. Peak Throughput.....	12
3.4. Duration of Temporary Bandwidth Constraints	13
4. Diagnosing Video Quality Of Experience Issues Based On Wi-Fi Metrics	14
4.1. Levels of Inference	14
4.2. Detection of a Wi-Fi Video Client Device	15
4.3. Wi-Fi Performance Monitoring	17
4.4. Diagnosis.....	17
5. Service Assurance for Video Over Wi-Fi	18
5.1. Proactive Assurance	18
5.1.1. WMM.....	18
5.1.2. Airtime (un)fairness.....	19
5.1.3. Band Steering	19
5.1.4. Channel Planning.....	20
5.1.5. Client Steering.....	22
5.2. Reactive Assurance	23
Conclusion.....	24
Abbreviations	25

List of Figures

Title	Page Number
Figure 1 - Evolution of cord cutting in USA	5
Figure 2 - Ownership of streaming video players in USA	6
Figure 3 - Penetration of OTT media players in US households	6
Figure 4 - UHD TV shipments as percentage of total TV shipments worldwide	7
Figure 5 - Growth of Internet traffic, fixed and wireless	9
Figure 6 - Example Wi-Fi link performance graph	10
Figure 7 - Poor Wi-Fi connection prevents upscaling to UHD	12
Figure 8 - Buffering at start of new video playback (excellent conditions).....	13
Figure 9 - Buffering at start of new video playback (good conditions)	13
Figure 10 - Codec recovery.....	14

Figure 11 - The effect of Wi-Fi quality on QoE.....	15
Figure 12 - MAC address structure including OUI	16
Figure 13 - Distribution of number of root causes (out of physics, interference, saturation) affecting Wi-Fi performance losses	18
Figure 14 - Percentage of dual-band capable Wi-Fi stations that do not connect to 5GHz	20
Figure 15: 2.4 GHz Wi-Fi channels.....	21
Figure 16 - Interference patterns between channels 1 and 13	21
Figure 17 - Households with optimal Wi-Fi QoE	22
Figure 18 - Incidence of radio path issues for OTT media players associated with four different models of Wi-Fi AP	24

List of Tables

Title	Page Number
Table 1 - Network throughput requirements (in Mbps) by video quality	11

Introduction

Today OTT is becoming a common way of delivering video to the home, with Wi-Fi as the transport medium once inside. Wi-Fi quality of experience (QoE) is far from acceptable in today's deployments, and streaming solutions try to overcome this disadvantage by using adaptive streaming algorithms and packet prioritization techniques. However, when a consumer chooses to stream 4k content from an OTT service, he or she will not be satisfied with SD video quality due to bad Wi-Fi.

Today's full home coverage solutions offer a combination of RRM/SON capabilities, extenders and roaming solutions. Although these solutions vastly increase the Wi-Fi QoE for the subscriber, they are unable to guarantee a 4k video service delivery over Wi-Fi.

Technicolor is convinced the desired experience for a subscriber lies in a dynamic, self-adapting home network. Dynamic is the key word, as it ensures the Wi-Fi network does not impact non-video applications when there is no video content active. All services need to blend seamlessly, without creating an impression that a consumer needs to sacrifice. Furthermore, the network needs to be able to deal with environmental changes, without noticeable impact for the subscriber.

A dynamic, service-aware system that can monitor and guarantee 4k OTT content delivery must be capable of the following:

- 1) Detection: Dynamic identification of video service flows and required bandwidth.
- 2) Monitoring: Through continuous monitoring of the system, indicate whether video service quality was adequate at any moment in time.
- 3) Proactive care: The system will proactively steer non-video devices to other bands or access points, to reduce their airtime consumption and safeguard the 4k video services.
- 4) Reactive care: In cases where Wi-Fi issues are so impactful that 4k video quality cannot be guaranteed, the system must indicate root causes and potential cures for the future.

Guaranteeing Seamless 4K OTT Content Delivery

1. A Burning Platform

Over-the-top (OTT) video streaming is quickly becoming a common way of delivering video to the home, with as many as 13% of US citizens having fully abandoned traditional cable and satellite TV services today (see Figure 1) and with many more using one or more OTT video streaming services alongside traditional TV. The adoption of ultra-high definition (UHD) quality content is also moving faster for OTT compared to traditional TV. All in all, OTT video is both a significant and a demanding internet service.




Figure 1 - Evolution of cord cutting in USA¹

Streaming video service providers obviously want their subscribers to enjoy an optimal quality of experience (QoE). In achieving this they are at the mercy of the quality of their subscribers' broadband and in-home network connections. Video streaming solutions try to overcome poor network performance issues by applying adaptive streaming algorithms² and packet prioritization techniques. However, when a consumer pays a premium in order to stream UHD content, he or she will not be satisfied with standard definition (SD) video quality due to bad broadband or bad Wi-Fi.

Consumers easily recognize video quality issues related to network performance. Consumers expect their internet service provider (ISP) to provide the required bandwidth (CAPEX investment) and technical support (OPEX cost) and will, almost instinctively, raise any issues first with the ISP before addressing their streaming video service provider, if ever. QoE issues have a negative impact on an ISP's net promoter score (NPS) and therefore the ISP wants to assure that their network services meet the requirements for an optimal OTT video QoE.

More and more network connections are wireless connections and we know that in-home Wi-Fi is already a pain point today. Combine that with the increased bandwidth requirements of UHD content and we see a big risk for ISPs who fail to guarantee seamless UHD content delivery in the home. On the flip side, ISPs who do succeed in this will surely benefit.

In this paper we explain how to achieve this with a solution based exclusively on Wi-Fi metrics. In other words, we will not rely on application level metrics, because these are often unavailable to ISPs.

2. Market Overview

In this chapter we look at the penetration rate of OTT video in terms of media players and the availability of UHD screens and content. We then look at the use of Wi-Fi connectivity by OTT media players.

2.1. Streaming Video Players

Providers of streaming video services rely on video client hardware and/or software to get their content in front of the consumer. The most widely available clients are websites and apps running on desktop PCs,

¹ eMarketer, July 2018; cord cutters are defined as individuals (millions) of age 18+ who no longer have access to traditional pay TV services

² Deepthi Nandakumar, Sagar Kotecha, Kavitha Sampath, Pradeep Ramachandran, Tom Vaughan. "Efficient Multi-Bitrate HEVC Encoding for Adaptive Streaming". IBC, 2016

laptops and smartphones. As such, almost every consumer with a broadband connection can access streaming video services. For an optimal viewing experience, however, the preferred video clients are ISP managed set-top-boxes (STB), over-the-top (OTT) media players, games consoles and smart TVs because these allow media to be consumed on large screens in the comfortable space of the living room.




Figure 2 - Ownership of streaming video players in USA³

OTT media players are penetrating consumer households at a steady rate. According to research from Parks Associates, ownership of media players has risen from about 6% of U.S. broadband households in 2010 to almost 40% at the beginning of 2018 (see Figure 2). Four vendors Roku (Roku TV), Amazon (Fire TV), Google (Chromecast) and Apple (Apple TV) together hold about 90% of this market. ComScore paints a similar picture of the streaming video services and their OTT media players adopted by US households (see Figure 3).

Besides dedicated video client hardware, streaming video services can also rely on video client software, usually in the form of apps, deployed to set-top-boxes and smart TVs. The same research from Parks Associates teaches us that more than half of U.S. broadband households own a smart TV and, of those households, almost half own an OTT media player in addition to the smart TV.




Figure 3 - Penetration of OTT media players in US households⁴

³ Parks Associates, May 2018

⁴ comScore, April 2017

2.2. UHD Screens and Content

The success of the UHD market is determined by the availability of UHD screens and the availability of UHD content. According to Parks Associates, UHD purchases represented 30% of US flat-panel screen purchases in 2017. According to information from Statista, a similar percentage of almost 30% is predicted for total TV shipments worldwide (see Figure 4).




Figure 4 - UHD TV shipments as percentage of total TV shipments worldwide⁵

Futuresource Consulting, who forecasts 35% of global TV sales in 2017 will be UHD, puts the worldwide household penetration of UHD TVs to 8%.⁶

There is a gap in the availability of screens versus the availability of content and the ability to deliver it. While more and more new video content is being produced in UHD, the majority share of existing content in content libraries is older content produced in lower quality formats. Both Futuresource Consulting and The Guardian see ISP own broadcast services—which still represents the primary video delivery method available to most consumers—lag behind in UHD content delivery versus streaming video services.⁷

2.3. UHD Video Streams

ITU-T approved the initial release of the H.265 video codec standard, enabling image resolution of up to 8192 by 4320 pixels, on April 13, 2013. Without any compression, transport of images at the maximum resolution would require a bitrate close to 50 Gbps.

Thanks to work in ITU-T and the MPEG forum, a next-generation compression technique has been defined with a higher efficiency than its predecessor (H.264). The new codec enables UHD video streams

⁵ Statista, December 2014

⁶ Futuresource Consulting, “4K UHD Content is Now Abundantly Available, but it's Not Reaching Devices”, November 2017

⁷ The Guardian, “2018 will be the year 4K TV goes big, but HDR still lags behind”, December 2017

to be delivered over widely available broadband and home connectivity solutions at speeds of a few tens of megabits. This is the H.265 video codec (also known as HEVC or the VP9N alternative), the successor to the H.264 codec (also known as MPEG 40 AVC). Specifically designed for optimal performance at ultra-high resolutions and high frame rates, it enables even higher accuracy for displaying motion images (e.g. sports, large screen movies).

ITU-T H.265 defines resolutions up to 8k. However, the industry is commonly adopting the 4k video format first.

ITU-T H.265 defines the 5.1 "main tier" video codec up to 40 Mbps⁸ which means that a product supporting this profile must be able to digest the 40 Mbps video transport stream, at a resolution of 4096x2160. It can be argued that the real target rate will be lower because the UHD resolution which has been mainly adopted today for TV is not based on the highest resolution.

Which actual bitrate will be used finally is unpredictable. Given the vast number of possible permutations (frame rate, chromatic subsampling, resolution, etc.) a long list of options will be available. The goal of this paper is not to predict the rate, but to define requirements for a Wi-Fi system that can handle the required throughput, which is why worst-case bitrates are used to assess the impact and define the solution architecture.

2.4. Wi-Fi for Multimedia Access

Although the 802.11 wireless local area network (WLAN) specification was not designed to transport video, its widespread popularity as a LAN interface has led many suppliers to seek to provide multimedia access via Wi-Fi. They have encountered numerous challenges. The more the WLAN standard has evolved, the more features have been added to help attain the multimedia distribution target.

With the release of 802.11a/g, Wi-Fi technology was able to transport non-real-time audio (e.g. MP3, WAV, FLAC streaming) reasonably well. However, IPTV streaming remained challenging due to an absence of multicast support and insufficient PHY layer techniques to support stable high bandwidths.

As 802.11n was adopted, some companies, including Technicolor, started to look at Wi-Fi products capable of transporting high-end IPTV applications. This was achieved through proprietary techniques (e.g. HW acceleration for Ethernet frames in an embedded SoC) and optional items in the 802.11n standard (e.g. explicit transmit beam-forming, LDPC, etc.).

In addition to IPTV distribution, many local content distribution systems were developed using either core or optional parts of the 802.11n standard. These systems benefitted greatly from the higher PHY rates and improved RF stability introduced with MIMO. The biggest drawback of all these systems remained broad scale interoperability. This was to be expected: many optional/proprietary features were enabled on top of the limited set of 802.11n core features for which the Wi-Fi Alliance provided an interoperability certification test plan.

The introduction of 802.11ac in 2013 and the accompanying 802.11ac certified certificate by the Wi-Fi Alliance addressed these issues. Not only did 802.11ac drastically improve interoperability, it also introduced new features. Thanks to the introduction of QAM 256 and 80 MHz/160 MHz modulation, the standard enables large bandwidth boosts and forward error correction. This created rich new opportunities for multimedia distribution over Wi-Fi.

⁸ ITU-T H.265 Annex A.4/Table A.1

Today, more and more network connections are wireless connections, be it Wi-Fi, mobile or other wireless technologies like Bluetooth. According to Cisco research, wireless devices accounted for 60% of all Internet traffic in 2015 and this share will rise to 78% by 2020 (see Figure 5). This trend is of particular concern to customer experience, knowing that getting connected through Wi-Fi is a major pain point today. Technicolor's own research shows that 1 out of 2 consumers experience Wi-Fi issues at home. This finding is corroborated by ISPs who systematically rank Wi-Fi-related issues (configuration, coverage, compatibility) at the top of their list of customer support tickets.




Figure 5 - Growth of Internet traffic, fixed and wireless⁹

3. Wi-Fi Performance Requirements For A Seamless UHD Experience

Robust streaming of UHD quality video content requires a sufficiently fast and stable network connection all the way from the video distribution platform (in the network) to the video client (in the home). In this paper we focus on LAN performance and in particular on Wi-Fi performance, which is often the weakest link, and we assume that WAN performance is sufficient.

Before proceeding with requirements, we first describe how to quantify and how to qualify Wi-Fi performance.

3.1. Quantifying and Qualifying Wi-Fi Performance and Performance Losses

In the next paragraphs we will show several figures which represent the performance of a Wi-Fi link between a Wi-Fi access point (AP) and a Wi-Fi station (STA). In these figures, the actual performance of the Wi-Fi link is plotted vs. its maximum theoretical performance. The maximum theoretical performance, or the maximum PHY rate, is determined by the Wi-Fi link configuration which includes the chosen Wi-Fi technology (e.g. 802.11n or 802.11ac), frequency band (e.g. 2.4GHz or 5GHz), bandwidth (e.g. 20MHz, 40MHz or 80MHz), MIMO configuration (e.g. 1x1, 2x2, 3x3 or 4x4) and short guard interval (SGI). The Y-axis of the figures represents the maximum PHY rate. The X-axis represents time.

In practice, the maximum theoretical performance of a Wi-Fi link is never reached due to a combination of factors. These factors can be qualified into three categories of issues:

⁹ Cisco Visual Networking Index Global IP Traffic Forecast, 2015

1. **Physics issues** arise from poor Wi-Fi coverage (STA too far away from the AP) stemming from long range and/or from constructions which heavily degrade the Wi-Fi signal (reinforced concrete walls, insulation, metal doors ...). Physics issues are what most people think of first when it comes to Wi-Fi issues. Whenever Wi-Fi link issues are experienced, end users instinctively check the “connection bars” on their device which indicate the signal strength of the Wi-Fi connection. Received signal strength (RSSI), a derivative of signal to noise ratio (SNR), relates directly to the maximum throughput¹⁰ or link capacity that can be achieved over a link. The further a STA is moved from an AP, the lower the achievable link capacity gets.
2. **Interference issues** arise from a destructive use of the shared Wi-Fi medium. Interference can be caused by a non-Wi-Fi RF interferer transmitting in a Wi-Fi frequency band or by two faraway Wi-Fi access points (so-called hidden nodes) generating collisions at the location of a Wi-Fi station because the access points are too far away to coexist and share the medium in a proper way. Interference issues are often underestimated due to the complexity of detecting and diagnosing them as opposed to physics issues. Distinguishing between interference seen on the side of the AP (near-end interference) and on the side of the STA (far-end interference) is important when determining the root cause.
3. **Saturation issues** arise from overutilization of the shared medium by other stations belonging to the same network. This can stem from one station that completely saturates the available Wi-Fi medium with P2P downloads or from the sum of many stations.

In the figures, we use a color scheme to quantify how much of the maximum theoretical performance is lost due to each of these factors:

1. Blue: physics issues
2. Red and orange: far-end and near-end interference issues
3. Yellow: saturation issues

The actual performance of the Wi-Fi link is shown in two shades of green, where dark green represents traffic to/from the station and light green is the available (unused) link capacity.

Figure 6 is an example of such a Wi-Fi link performance graph.




Figure 6 - Example Wi-Fi link performance graph¹¹

¹⁰ SNR defines the physical layer or modulation rate which is related to the actual throughput experienced by an end user. The relation between both can be expressed by a simple approximation of [throughput = PHY rate x MAC efficiency], whereby the MAC efficiency factor is a representation (in %) of the throughput loss due to framing overhead between the physical layer and the MAC layer. As such, it should be clear that the throughput degrades when the physical layer rate degrades.

¹¹ Technicolor, 2017

3.2. Throughput

When talking about performance, throughput is the most commonly used metric. In the following table, we compare the network throughput requirements for different qualities of video content as specified by three popular video streaming providers: Netflix, Amazon's Prime Video and Google's YouTube. Netflix and Prime Video refer to "Internet speed requirements" while YouTube refers to "video bitrates". We will demonstrate later that neither definition fits the bill entirely.

Table 1 - Network throughput requirements (in Mbps) by video quality

Quality	Netflix ¹²	Prime Video ¹³	YouTube (SDR) ¹⁴
SD (1k)	3	Not specified	5-7.5
HD (2k)	5	Not specified	8-12
UHD (4k)	25	15	35-68

Prime Video does not specify Internet speed requirements for SD and HD quality content. We assume this is because Amazon considers that the vast majority of Internet subscriptions can support these bitrates.

Prime Video and YouTube distinguish between SDR and HDR content in their requirements. While Prime Video specifies the same Internet speed requirements for SDR and HDR content alike—which we assume is either a simplification or a mistake—YouTube sets the requirements for HDR roughly 25% above those for SDR content.

Table 1 shows that the network throughput requirements for streaming UHD quality video are 4 to 5 times higher than those for HD quality video. This may come as a surprise to some consumers because UHD could be considered as "one" step up from HD and 4k could be considered as "twice" 2k. In reality, the resolution of UHD is 3840x2160=8.3M pixels compared to 1920x1080=2.1M pixels for HD, hence the factor 4 to 5 increase of bitrate for UHD vs. HD.

Going from SDR to HDR does not increase the number of pixels but rather increases the color depth from 8 to 10 bits per pixel, hence the 25% increase of bitrate for HDR vs. SDR.

As video encoding technology advances, we may expect a reduction of network throughput requirements for a given video content quality. The uptake of such optimizations is slowed by interoperability concerns. Video streaming providers are incentivized to use more commonly supported video codecs so that they can reach wider audiences with video clients based on older hardware and software platforms. Therefore, we expect these network throughput requirements to remain stable in the next years. As an example, the encoding profiles used by Netflix¹⁵ include both older and newer technologies: VC1, H.264/AVC Baseline, H.264/AVC Main and HEVC.

When the base network throughput requirement for streaming UHD quality video content is not met, then the adaptive streaming algorithm will automatically scale down to lower qualities. Figure 7 shows a real-life example of this scenario. An OTT device (LG 4k TV) is connected to the Internet via a poor Wi-Fi link. Even though the specific Wi-Fi configuration used here could yield a maximum theoretical

¹² "Internet Connection Speed Recommendations", <https://help.netflix.com/en/node/306>, June 2018

¹³ "Prime Video Quality & Formats", <https://www.amazon.com/gp/help/customer/display.html?nodeId=201648150>, June 2018

¹⁴ "Recommended upload encoding settings", <https://support.google.com/youtube/answer/1722171>, June 2018

¹⁵ "High Quality Video Encoding at Scale", <https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-d159db052746>, December 2015

throughput of 144Mbps (11n, 2x2 MIMO, 2.4GHz band, 20MHz bandwidth), poor coverage (the blue in the figure) and to a lesser extent far end interference (the red in the figure) are restricting the actual throughput on this Wi-Fi link to roughly 15Mbps. When playing UHD content (Netflix, The Crown, season 2, episode 5) on this link, the adaptive stream never scales beyond HD quality.




Figure 7 - Poor Wi-Fi connection prevents upscaling to UHD¹⁶

In this example we obviously cannot speak of a seamless UHD experience. The consumer who has invested in a 4k TV, a faster broadband connection and/or a UHD video streaming plan (some providers like Netflix charge extra for access to UHD quality content) will be unhappy. Sustained unhappiness will affect the profitability of Internet service providers and video streaming providers alike. The ISP will see OPEX increase as subscribers call for customer service. Both the ISP and the video streaming provider will see ARPU decrease as subscribers give up on premium broadband and video plans and will see increasing rates of churn.

Figure 7 is a good example of why the definition of “Internet speed requirements” maintained by Netflix and Prime Video misses the mark. The broadband connection is shared by multiple devices and applications in the home. When using Wi-Fi to connect the video client, a variety of factors can lead to an actual speed which is well below the Internet speed. A much better definition would be “speed requirements for the connection to your video client”.

3.3. Sustained Throughput Vs. Peak Throughput

In their Internet speed requirements, video streaming platforms neglect to mention the impact of peak throughput on the UHD experience. Why is peak throughput important? Video streaming clients always strive to maintain a buffer of several seconds of content in order to compensate for temporary drops in network throughput or even temporary loss of the network connection. When launching a new video, the client will hold off before playback until its buffer has been filled above a certain threshold. The duration of this delay depends on the peak network throughput: with a higher throughput the buffer will be filled faster, and vice versa. With a higher peak throughput, playback will commence sooner, which improves the UHD experience. This effect is not only noticeable when commencing a video, but also when skipping through a video forward and backward.

Figure 8 shows an example of buffering under ideal conditions. Our OTT device (LG 4k TV) is now connected to the Internet via an excellent Wi-Fi link. The specific Wi-Fi configuration used here yields a maximum theoretical throughput of 867Mbps (11ac, 2x2 MIMO, 5GHz band, 80MHz bandwidth). Some performance losses due to coverage (the blue in the figure) result in an actual throughput on this Wi-Fi link of roughly 650Mbps, which is far beyond the requirements of UHD video. When playing UHD

¹⁶ Technicolor, June 2018

content (Netflix, The Crown, season 2, episode 1) on this link, video playback commences promptly and immediately in UHD quality.




Figure 8 - Buffering at start of new video playback (excellent conditions)¹⁷

In this example, the data rate at the time of buffering climbs above 100Mbps before converging to a steady rate of 20Mbps.

Figure 9 shows another example of buffering under seemingly ideal conditions. Our OTT device (LG 4k TV) is connected to the Internet via a good Wi-Fi link. The specific Wi-Fi configuration used here yields a maximum theoretical throughput of 144Mbps (11n, 2x2 MIMO, 2.4GHz band, 20MHz bandwidth). Some occasional coverage issues (the blue in the figure) result in an actual throughput on this Wi-Fi link varying between 100Mbps and 120Mbps, which is still well beyond the requirements of UHD video. However, when playing UHD content (Netflix, The Crown, season 2, episode 2) on this link, it takes a few seconds longer to commence video playback compared to the previous example. What's more, initially the video codec commences in HD quality before scaling up to UHD quality (and staying there for the duration of the video).




Figure 9 - Buffering at start of new video playback (good conditions)¹⁸

In this example, the data rate at the time of buffering climbs to about 60Mbps before converging to a steady rate of 20Mbps.

What we learn from these two examples is that, in order to guarantee a truly seamless UHD experience, network throughput requirements should be set higher than what the video streaming platforms specify.

3.4. Duration of Temporary Bandwidth Constraints

Just like adaptive streaming algorithms will scale down to lower qualities when placed under bandwidth constraints, you would expect them to scale back up once the bandwidth constraints have been lifted. This

¹⁷ Technicolor, June 2018

¹⁸ Technicolor, June 2018

is indeed the case after short periods of constrained bandwidth, but when the bandwidth is constrained for a longer period of time then an adaptive streaming algorithm may “give up” and never scale back up to the best quality. Figure 10 shows an example of this scenario. Our OTT device (LG 4k TV) is still connected to the Internet via a poor Wi-Fi link with the same configuration as in Figure 7. The difference is that for the first 45 minutes a combination of poor coverage and far end interference leads to extreme bandwidth constraints, which causes the adaptive streaming algorithm to scale down to SD quality. Even when the interference is removed, the codec remains at SD quality until the end of the episode. Video quality scales up to HD quality only upon starting the next episode.




Figure 10 - Codec recovery¹⁹

What this example teaches us is that, while short periods of bandwidth constraints will be transparent to the consumer as long as the buffer is not depleted, longer periods will reduce the video experience. First of all, the codec will scale down to a lower quality level. If the bandwidth constraints last for too long, the codec will not scale back up even when the constraints are lifted.

4. Diagnosing Video Quality Of Experience Issues Based On Wi-Fi Metrics

4.1. Levels of Inference

The most accurate way to assess the end user quality of experience of streaming UHD quality video content over Wi-Fi is by asking the user for his or her opinion. User experience is subjective by nature and nothing trumps getting the user’s opinion, which will be determined by everything from personal taste, prior experience and visual acuity to screen size, viewing distance and Wi-Fi performance. The goal we have set out in this paper, however, is to assess QoE in an automated fashion based on Wi-Fi quality metrics. Figure 11 shows some different levels of quality inference where user opinion and Wi-Fi metrics occupy opposite ends of the scale.

¹⁹ Technicolor, June 2018




Figure 11 - The effect of Wi-Fi quality on QoE²⁰

Why not focus on application metrics? Application-level diagnostics, in our application (i.e. video) running in the media player hardware and/or software, are well-placed to predict end user QoE. Video codec statistics such as the average frame rate, the resolution selected by the adaptive streaming algorithm and the buffer health (the number of video frames that have been buffered in order to overcome temporary connection issues) clearly indicate the robustness of the UHD video stream. Unfortunately, while these statistics are commonly available to the streaming video service, they are not commonly available to the broadband service provider, perhaps with the exception of joint offerings such as an integrated streaming video app on a set-top-box. So, in general, the broadband service provider needs a more accessible and independent basis for gauging the end user QoE than application metrics.

Network metrics, and Wi-Fi quality metrics in particular, can provide that basis. The ISP controls the broadband connection and monitoring broadband throughput, demonstrated earlier to be a key requirement for streaming video, yields useful data points to infer end user QoE. Broadband throughput is often inconclusive, however, because the broadband connection is just one part of the chain end-to-end and almost always an intermediate part (unless the media player is combined with the broadband access terminator) and often not the weakest part of the chain compared to the in-home Wi-Fi connection between the broadband access terminator and the OTT media player. It is therefore essential for the broadband service provider to monitor Wi-Fi performance in the home.

4.2. Detection of a Wi-Fi Video Client Device

A challenge in monitoring and optimizing the Wi-Fi performance of an OTT media player lies in identifying the Wi-Fi station which represents the media player. An OTT hardware device is, by nature, not provisioned by the service provider and therefore may resemble any other user device in the LAN. OTT traffic is also, by nature, not marked or classified in any particular way, unlike ISP managed traffic such as broadcast video or VoIP. When the OTT video client comes in the form of an app running on a multi-purpose hardware device, such as a tablet or a games console, it becomes even more difficult to recognize.

Several techniques of varying efficiency and effectiveness are available for identifying an OTT media player device and/or an OTT video stream:

²⁰ Diego Da Hora, Karel Van Doorselaer, Koen Van Oost, Renata Teixeira. “Predicting the effect of home Wi-Fi quality on QoE”. IEEE International Conference on Computer Communications, April 2018.

- **Deep packet inspection (DPI):** A technique which identifies network devices and streams by inspecting network packets that are forwarded through a network node (e.g. the broadband router) and matching their payload against a database of fingerprints. OTT video devices can be identified by, for example, inspecting the User-Agent header in HTTP requests which often contains operating system and browser identifiers which can be mapped to specific devices like a Chromecast. OTT video streams can be identified by, for example, inspecting the public certificate used to secure the connection between client and server which can be mapped to specific services like Netflix. DPI is very effective but comes with steep hardware and software requirements which make it costlier to deploy than other techniques.
- **Traffic pattern analysis:** By analyzing the network traffic pattern of a device, the type of device or the type of service used by the device can be deduced. This technique requires access to sufficiently rich live or historical data and may not work well for devices which are used for more than one type of service such as a PC.
- **DHCP options:** When a network device requests an IP address using DHCP, it can include so-called DHCP options in its DHCP request to pass specific requests and extra information to the DHCP server. One of these options is the Vendor Class which is used to convey information about the vendor that manufactured the hardware on which the DHCP client is running. Typically, DHCP options can be used to identify an OTT device vendor such as Apple but not a specific OTT device type such as an Apple TV. On the flip side, because DHCP is so widely used, relying on DHCP options is cost-effective.
- **MAC OUI:** Every Wi-Fi station is uniquely identified at the data link layer by its MAC address. Part of this address (see Figure 12) is reserved for the so-called Organizationally Unique Identifier (OUI) which represents the vendor of the device as registered in the public and global OUI database managed by IEEE. As with DHCP options, MAC OUI can be used to identify a vendor but not a specific device type. MAC OUI is also cost-effective.

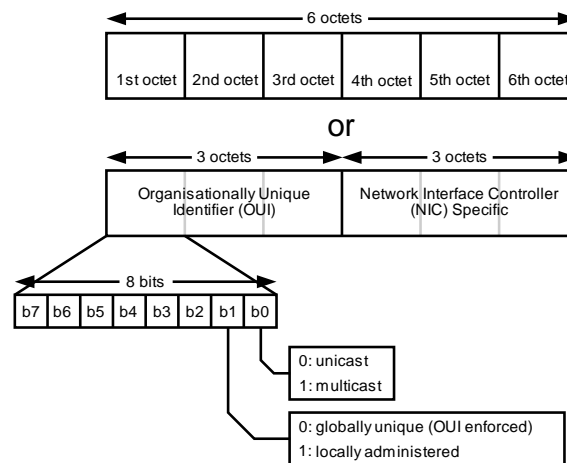


Figure 12 - MAC address structure including OUI²¹

- **User-defined rules:** An end user can manually identify or classify a device as being an OTT device using a web user interface or app provided by the service provider. This technique can be extremely accurate, but a drawback is that it requires interaction with the end user which can be perceived as an annoyance.

²¹ Wikimedia, 2007

Regardless of which techniques are used to identify the OTT media player, the result is that we know which Wi-Fi station's performance to monitor.

4.3. Wi-Fi Performance Monitoring

Once we have identified the Wi-Fi station, we can monitor its Wi-Fi performance. From the practical experiments described in chapter 3, we distill three Wi-Fi performance metrics which must be monitored in order to assess OTT video QoE.

- Must be able to measure Wi-Fi link capacity, because this determines the generally obtainable video quality level and influences the video start and skip delay
- Must be able to measure available Wi-Fi link capacity for a given Wi-Fi station, not just for the Wi-Fi network as a whole, because the home network will be shared with other Wi-Fi devices
- Must be able to measure true Wi-Fi link capacity, because in reality the maximum PHY rate is never attained due to physics and far-end interference issues and the trained PHY rate is never attained due to near-end interference issues and normal sharing of the Wi-Fi medium
- Must be able to perform measurements at a high sample rate, because transient effects will cause QoE issues depending on video buffering settings

4.4. Diagnosis

As elaborated in chapter 3.1, Wi-Fi performance losses can be attributed to several factors like physics, interference and saturation issues. In order to recommend the appropriate course of action for mitigating performance losses, it is imperative to make the correct diagnosis first. Physics issues can be addressed by repositioning the Wi-Fi station or by installing additional Wi-Fi access points to improve coverage. Interference issues can be addressed by using a different channel within the same frequency band or by moving the Wi-Fi station to another band. Saturation issues can be addressed by increasing the Wi-Fi medium's overall capacity or by throttling other stations sharing the medium.

A Wi-Fi performance problem will often have more than one root cause and therefore more than one recommended course of action. Individual actions should be prioritized based on the contribution of every factor to the overall loss of performance. With the lessons learned from several large Wi-Fi deployments by Technicolor, we know that in 92% of the cases Wi-Fi performance losses are not linked to only one root cause. In 80% of the diagnosed cases there are two root causes.




Figure 13 - Distribution of number of root causes (out of physics, interference, saturation) affecting Wi-Fi performance losses²²

In the next chapter, we review various techniques to improve Wi-Fi experience in case of issues and we apply these techniques specifically to OTT video scenarios.

5. Service Assurance for Video Over Wi-Fi

Streaming video service providers want their subscribers to enjoy an optimal quality of experience. In achieving this they are at the mercy of the quality of their subscribers' broadband and in-home network connections, however. Consumers understand this too and recognize easily video quality issues related to network performance. They will, almost instinctively, raise these issues first with their internet service provider before addressing their streaming video service provider, if ever. OTT video quality of experience issues have a negative impact on an ISP's OPEX and NPS and therefore the ISP wants to assure their network services meet the requirements for an optimal streaming video quality of experience.

We can distinguish between proactive assurance and reactive assurance. Proactive assurance includes all measures for identifying and correcting or avoiding outright Wi-Fi performance losses before they manifest as QoE issues for the end user. This form of assurance is obviously preferred. Reactive assurance is when Wi-Fi performance issues are identified and corrected only after being reported by the end user.

5.1. Proactive Assurance

5.1.1. WMM

The Wi-Fi Multimedia (WMM) standard defines basic Quality of Service (QoS) mechanisms for Wi-Fi traffic. By tagging certain traffic as voice or video, this traffic will be transmitted with priority versus traffic tagged as best effort and background. WMM is often used for IPTV distribution over Wi-Fi. Because IPTV streams are well-defined end-to-end from all the way from the access network to the STB, it is trivial to tag the traffic. A challenge for applying WMM to OTT video streams follows from chapter 4.2: in order to tag traffic, one needs to recognize it, and this is not obvious for an OTT media player device and/or an OTT video stream.

²² Technicolor, July 2018

A drawback of WMM is that it is unfair to other Wi-Fi traffic: an OTT media player with weak Wi-Fi signal and low PHY rate receiving high priority traffic can starve all other traffic. By solving one issue for the end user, another is created. Therefore, WMM works best in Wi-Fi networks where there is enough airtime available to transmit all data within the home network.

5.1.2. Airtime (un)fairness

Airtime fairness is a feature that attempts to assign Wi-Fi airtime more fairly between Wi-Fi stations in the home network. The benefits of airtime fairness are most apparent when considering two different Wi-Fi stations, one slow and one fast. When transmitting an equal amount of data, without airtime fairness, the slow station would consume more Wi-Fi airtime (because it takes longer to transmit) than the fast station. An OTT media player that requires a lot of airtime to receive an UHD stream could experience starvation from other stations. With airtime fairness, the slow station will be throttled to give the fast station more airtime.

A drawback of airtime fairness is that being fair is not always best for the user experience. Being fair might cause an UHD quality video to scale down to HD quality while a large file is being downloaded to another PC, which is likely not what the end user wants. Some implementations of airtime fairness allow specific priorities to be set—making things a little more unfair on Wi-Fi level—in which case the challenge is again identifying the OTT media player.

5.1.3. Band Steering

WMM and airtime fairness are techniques that try to optimize the situation within a given Wi-Fi link capacity. Band steering is a feature that attempts to move dual-band capable Wi-Fi stations from the slower 2.4GHz frequency band to the faster 5GHz frequency band, thereby increasing the link capacity for the station. While a connection to 2.4GHz does not necessarily imply that the user experience is bad, it is well understood that a station capable of moving to 5GHz at a specific location will benefit from performing said action. The typical dual-band capable station employs an 802.11ac WLAN radio hence it can use 80MHz modulation rates on the 5GHz band whereas the 2.4GHz band generally only allows the use of 20MHz channels in 802.11n mode.

Many Wi-Fi stations will, under good conditions and when given an equal choice, prefer to connect to 5GHz already. However, we see that Wi-Fi stations will sometimes refrain from choosing 5GHz as preferred operational frequency band and can end up being stuck on 2.4GHz, potentially leading to a degraded user experience as the (maximum) link capacity is limited.

Certain conditions must be met before steering a Wi-Fi station to another band. First, we must confirm that the station is dual-band capable. If the OTT media player is not known beforehand, this capability can be derived from the 802.11 probe request²³ data. Second, we must confirm that the Wi-Fi station is able to connect to the access point on the other frequency band. This is achieved the most easily by assigning the same SSID and credentials to the 2.4GHz and to the 5GHz access point. If either of these conditions would not be met, then a band steering action would risk disconnecting the OTT media player from the network.

²³ Wi-Fi stations send 802.11 probe requests to known APs (and to broadcast MAC addresses) when they are (active) scanning the available channels. These probe request can be captured by an AP as they indicate the STAs ability to operate on a specific channel.




Figure 14 - Percentage of dual-band capable Wi-Fi stations that do not connect to 5GHz²⁴

Figure 14 shows out of a large population of Wi-Fi households how many dual-band capable Wi-Fi stations actually connect to 5GHz. 42% of the stations does not connect to 5GHz even though they are capable of doing so. The reason for this is two-fold. On one hand, some of the households use a different SSID for each band and most consumers connect first to the 2.4GHz SSID and do not bother to configure the 5GHz SSID. On the other hand, even with the same SSID, certain Wi-Fi stations are “sticky” and refrain from connecting to 5GHz. This clearly demonstrates the need for unifying the Wi-Fi configuration in the home and for band steering.

5.1.4. Channel Planning

Channel planning²⁵ aims to select better Wi-Fi channels (within each frequency band, so on 2.4GHz and on 5GHz) to improve overall Wi-Fi link capacity within the constraints of the Wi-Fi environment (e.g. neighboring Wi-Fi networks).

The 2.4GHz frequency band is almost universally supported by Wi-Fi stations for historical reasons and is still widely used today. The band is divided in 13 channels (see Figure 15), of which only the first 11 are permitted in the United States. Only three of these channels are non-overlapping (1, 6, 11), meaning that in a dense Wi-Fi environment only three nearby access points can transmit without interfering with each other. Any additional access points trying to use 2.4GHz in the same location will cause interference. It is easy to understand why the 2.4GHz frequency band is so congested and why channel planning is needed to make the most out of it.

²⁴ Technicolor, July 2018

²⁵ Olivier Jeunen, Patrick Bosch, Michiel Van Herwegen, Karel van Doorselaer, Nick Godman, Steven Latré. “Data-driven Frequency Planning”.




Figure 15 - 2.4 GHz Wi-Fi channels²⁶

The 5GHz frequency band supports more than 20 non-overlapping channels when using a 20MHz bandwidth. Even with the bandwidth set to 80MHz there are still more choices available than in the 2.4GHz frequency band.

The impact of a channel change can be severe as we can see in the example of Figure 16. Until shortly before 7:00 PM a Wi-Fi station is suffering quite badly from interference (red and orange colors). This combined with the fact that the station is not very close to the access point (blue color) leaves almost no available link capacity remaining for user traffic. Just before 7:00 PM the access point is switched to channel 13 and the interference all but disappears.




Figure 16 - Interference patterns between channels 1 and 13²⁷

Most Wi-Fi access points employ a feature called Automatic Channel Selection (ACS) that periodically (e.g. several times a day) scans the regulatory allowed list of channels and reconfigures the access point to use the best channel seen at that moment. The periodic scan interval is not set very aggressive because the scans are service interrupting, except for access points with a dedicated scanning radio. This is a compromise knowing that Wi-Fi interference issues are often intermittent. A residential area will be largely empty during work hours and more vibrant during evenings and weekends, hence the Wi-Fi from those households will exhibit different patterns during the day. The same applies to office buildings but in the reverse. ACS will not catch on to these interference patterns because it only monitors periodically during very short intervals.

A limitation of most ACS implementations is that they are unable to recognize far-end interference. This implies that ACS would not have picked up the interference on channel 1 in the example of Figure 16. If the access point had remained on channel 1, obtaining a seamless UHD video experience would have been impossible. A good channel planning solution must be capable of distinguishing between near-end interference and far-end interference.

²⁶ Wikimedia, 2009

²⁷ Technicolor, 2017

The impact of channel planning on in-home Wi-Fi quality of experience is significant. Figure 17 shows aggregated results from five Wi-Fi deployments by Technicolor, of which three ISPs are using channel planning and two ISPs are not. Despite the mix of geographical regions and Wi-Fi access point products and configurations, it is easy to see who benefits from channel planning by looking at the percentage of households having an optimal Wi-Fi QoE: roughly 2 out of 3 households, compared to roughly 1 out of 2 households when no channel planning is used.




Figure 17 - Households with optimal Wi-Fi QoE²⁸

Wi-Fi interference is location-dependent, which means that different channels will yield different performances for different stations throughout the home. A channel planning solution can be configured to give more weight to certain Wi-Fi stations such as an OTT media player.

To summarize, a good channel planning solution can complement traditional ACS implementations by recognizing far-end interference and by reacting to intermittent interference issues. Coincidentally, these considerations align well with the requirements for a Wi-Fi performance monitoring solution capable of assessing video QoE listed in chapter 4.3.

5.1.5. Client Steering

It is becoming more and more common to extend Wi-Fi coverage in the home by deploying multiple Wi-Fi access points (a.k.a. Wi-Fi mesh networks). However, just having multiple access points does not guarantee that Wi-Fi stations will make the most efficient use of them. While the intelligence of roaming behavior, as initiated by the stations themselves, improves with every new generation of devices, the reality is that most devices in a consumer household are not updated nor upgraded very often. Next to that, a Wi-Fi station does not always have the means to assess the full environment and exploit it, when compared to the information that can be extracted by combining the view of several access points in the home. Last but not least, many Wi-Fi devices were simply not anticipated and designed to be nomadic. As a result, similar to the behavior seen with band steering, Wi-Fi stations can be “sticky” and refrain from connecting to the right access point in a multi-AP deployment. Client steering is a feature which addresses this challenge.

We identify several use cases within the realm of client steering in a multi-AP deployment:

²⁸ Technicolor, July 2018

1. **Signal strength-based roaming** is the most commonly supported use case where Wi-Fi stations that do not move autonomously are roamed proactively before the signal degrades to the point that QoE issues arise,
2. **Interference-based roaming** is a more advanced use case where Wi-Fi stations experiencing performance loss due to interference are roamed proactively to Wi-Fi access points on other channels or frequency bands,
3. **Load-based roaming** is a use case where Wi-Fi stations suffering from oversubscription of the Wi-Fi medium are roamed proactively in such a way that the overall load in the home Wi-Fi network is balanced.

The actual roaming action can be triggered in several ways. Per IEEE 802.11-2016, two mechanisms are defined. The most straightforward one is not really a roaming mechanism at heart, but rather a general-purpose disconnection mechanism that has existed ever since the first version of IEEE 802.11. This mechanism simply implies that an AP wishing to terminate an STA connection sends an IEEE 802.11 disassociation or a deauthentication frame to the target STA, typically combined with blocking future reassociation by applying an access control list. A second, more elegant roaming mechanism was introduced by adoption of the 802.11v substandard. In this case, a proper roaming request is sent from the AP to the target STA, allowing for a smoother transition. Ultimately, even when using 802.11v, every roaming action still carries a small risk of service interruption on application level and therefore roaming should be handled with care.

All of these roaming use cases apply to OTT video players. When the Wi-Fi performance requirements for a seamless UHD experience are not met, client steering can be applied to improve the link capacity and link stability. The fact that OTT video codecs use a generous buffer means that roaming should generally happen transparently for the end user. Nevertheless, a more failsafe approach exists which is to roam other Wi-Fi stations away from an access point in order to improve the Wi-Fi link for the OTT video player which stays behind. Also, it may be desirable to allow only the 802.11v roaming mechanisms for an OTT video player.

5.2. Reactive Assurance

When all proactive assurance measures fail, we must resort to reactive assurance, which implies that the end user is aware of the issue. It is still preferred to inform the end user proactively, rather than wait for the end user to call the helpdesk. This highlights the importance of having a Wi-Fi monitoring system which is capable of identifying or, better yet, predicting video QoE issues and alerting the ISP.

The most common scenario where the end user needs to be involved is when there is a Wi-Fi coverage issue. When a Wi-Fi station is truly too far away, some kind of manual action must be taken such as:

- Moving the station closer to an AP
- Installing an additional AP (a.k.a. a Wi-Fi extender)
- Moving an AP closer to the station (less practical when the main AP in the home is combined with the broadband access terminator)
- Upgrading an AP to a better performing AP
- Switching to a wired connection

The results of a case study into the incidence of radio path issues with different brands of OTT media players are shown in Figure 18. The Apple TV, Chromecast and Roku are found in significant quantities in Wi-Fi households using four different types of Wi-Fi access points. The Apple TV tends to exhibit less radio path issues than the Chromecast and the Roku, perhaps due to better Wi-Fi antenna design. More

interestingly, the incidence of radio path issues is much smaller with AP C and AP D than with AP A and AP B. The explanation is that AP C and AP D have a better Wi-Fi antenna design, an increased MIMO and a higher Wi-Fi power output. This case study demonstrates that moving a Wi-Fi station or installing a Wi-Fi extender is not always required in order to fix a Wi-Fi coverage issue.




Figure 18 - Incidence of radio path issues for OTT media players associated with four different models of Wi-Fi²⁹ AP

Conclusion

In this paper, we demonstrated that guaranteeing a seamless UHD OTT video streaming over Wi-Fi experience is achievable by deploying a dynamic, self-adapting home Wi-Fi network. The right solution relies on four key capabilities:

1. **Detection** of the OTT media player or OTT video stream,
2. **Monitoring** of the Wi-Fi link to the OTT device and of the whole home Wi-Fi network, in particular the accurate assessment of link capacity performed at a high sample rate and the accurate diagnosis of issues,
3. **Proactive assurance** to mitigate issues before they become apparent to the end user,
4. **Reactive assurance** to resolve those remaining issues that mandate end user involvement.

This enables ISPs to assure that their network services meet the requirements for an optimal OTT video QoE and elevate subscriber NPS.

²⁹ Technicolor, 2017

Abbreviations

ACS	Automatic Channel Selection
AP	Access Point
CAPEX	Capital Expense
DHCP	Dynamic Host Configuration Protocol
DPI	Deep Packet Inspection
HD	High Definition
HDR	High Dynamic Range
HTTP	Hyper Text Transfer Protocol
IEEE	Institute of Electrical and Electronics Engineers
IPTV	Internet Protocol Television
ISP	Internet Service Provider
ITU	International Telecommunication Union
LAN	Local Area Network
MAC	Media Access Control
MDU	Multi-Dwelling Unit
MIMO	Multiple Input Multiple Output
MPEG	Moving Pictures Experts Group
NPS	Net Promoter Score
OPEX	Operational Expense
OTT	Over the Top
OUI	Organizationally Unique Identifier
P2P	Peer to Peer
PHY	Physical layer
QoE	Quality of Experience
QoS	Quality of Service
RF	Radio Frequency
RRM-SON	Radio Resource Management – Self Optimizing Networks
RSSI	Received Signal Strength Indication
SD	Standard Definition
SDR	Standard Dynamic Range
SGI	Short Guard Interval
SNR	Signal to Noise Ratio
SP	Service Provider
SSID	Service Set Identifier
STA	Station
STB	Set Top Box
UHD	Ultra-High Definition
WAN	Wide Area Network
WLAN	Wireless Local Area Network
WMM	Wi-Fi Multimedia

Harvesting Unlicensed and Shared Spectrum: Opportunities and Challenges

A Technical Paper prepared for SCTE•ISBE by Fontech

Narayan Menon
CTO & EVP Engineering
Fontech
28 Devine Ave, Syosset, NY 11791, USA
(516) 343-0027
narayan.menon@fon.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Growing Scope for Unlicensed Spectrum Usage.....	3
1. Unlicensed Spectrum as it Exists Today (in 4G).....	3
2. Greater Scope for Unlicensed Spectrum Use in a 5G World – More Bands, a Readymade Standard.....	4
3. 5G Technology Is Unlicensed Spectrum Friendly from the Outset.....	5
4. Spectrum Sharing Paradigms	5
Unlicensed Spectrum Use Cases	6
1. Next-Generation Home Wi-Fi System.....	6
2. Community Wi-Fi Hotspot Network	7
3. Licensed – Unlicensed Aggregation.....	7
4. LTE or 5G Operation Standalone in Unlicensed or Shared Spectrum	7
5. Licensed Assisted Unlicensed Access.....	7
6. Vertical Spectrum Sharing – CBRS	8
Unlicensed Spectrum - Performance Challenges	9
Addressing Performance Challenges	10
1. Setup	11
2. Dynamic Channel Management.....	11
3. Band Steering for Bandwidth Optimization	12
4. Client Steering (Device Mobility) for Coverage Optimization	13
Conclusions.....	14
Abbreviations	14
Bibliography & References.....	15

List of Figures

Title	Page Number
Figure 1: Existing US Spectrum Allocation	4
Figure 2: US Spectrum Allocation in a 5G World.....	4
Figure 3: Next-Generation Home Wi-Fi	6
Figure 4: LTE/5G in Licensed + Unlicensed Spectrum.....	8
Figure 5: CBRS Framework and Use Cases	9
Figure 6: Addressing Performance Issues at Two Levels	10
Figure 7: Reduced Congestion Spikes with RRM.....	11
Figure 8: Dramatically Reduced Number of Congestion Events with RRM.....	12
Figure 9: Bandwidth Optimization with Band Steering.....	13
Figure 10: Inter-AP Client Steering to Optimize Coverage	13

Introduction

The inexorable rise in wireless capacity demand continues to place strains on bandwidth, as we move into a 5G world. On one hand, more and more users are actively using the system, and demanding continuously improved quality of experience. Adding to this is an increasingly diverse array of applications and services, with their own Quality of Service (QoS) requirements. Add in a diversity of terminal devices, and the resultant capacity demand can be expected to push or exceed bandwidth supply even in a 5G world.

Traditionally, wireless technologies have expanded capacity in three ways:

- By adding more spectrum, i.e. increasing supply;
- Via improved spectral efficiency, making better use of the spectrum that is available. This has been done via higher-order modulation schemes (packing more bits into the link), spatial multiplexing (Multiple Input Multiple Output (MIMO) and Multi-User MIMO schemes) and by use of higher spectrum bands (“fatter pipes”);
- And via spectrum reuse: dividing the network into smaller and smaller cells, and reusing wireless capacity (the same channels) over and over across a given geographical area. This provides a capacity multiplier effect, and has historically been a major contributor to capacity growth.

Spectrum sharing in unlicensed bands can significantly add to the available pool of wireless bandwidth. Swathes of spectrum are available today in unlicensed bands, governments are releasing new, lightly used bands for spectrum sharing, and new unlicensed bands are being allocated for use in 5G.

This paper discusses opportunities for use of unlicensed bands to augment capacity, several use cases that unlicensed spectrum can support, key challenges that need to be surmounted and solutions to address them.

Growing Scope for Unlicensed Spectrum Usage

Spectrum sharing in unlicensed bands occurs today. Wi-Fi is a good example – multiple Wi-Fi networks and devices coexist in the 2.4 and 5 GHz band today, obeying a well-defined access etiquette. However, unlicensed spectrum usage will grow further as we move into a 5G future. This will be driven by several factors.

1. Unlicensed Spectrum as it Exists Today (in 4G)

Figure 1 below depicts key spectrum bands allocated (in the U.S.), as they exist today. Several unlicensed bands exist – notable among these are the Industrial Scientific and Medical (ISM) bands in 2.4 GHz and 5 GHz. The ISM bands have traditionally been occupied by unlicensed technologies – predominantly by Wi-Fi, although the 2.4 GHz band also hosts technologies such as ZigBee and Bluetooth. Unlicensed LTE operation variants, such as LTE Licensed Assisted Access (LTE-LAA) and MulteFire, are set to operate in the 5 GHz band, in addition to Wi-Fi.

The 60 GHz band supports ultra-high bandwidth, short-range communications, and is considered part of millimeter-wave (mmWave) spectrum. 802.11ad technology (also branded as WiGig) operates in this band today; WirelessHD is another standard that is designed to operate in this band. Both technologies

target short-range video products that connect high-definition television (HDTV) sets, Digital Video Recorders (DVRs), set-top boxes, gaming stations and other devices capable of stream uncompressed video over short ranges.

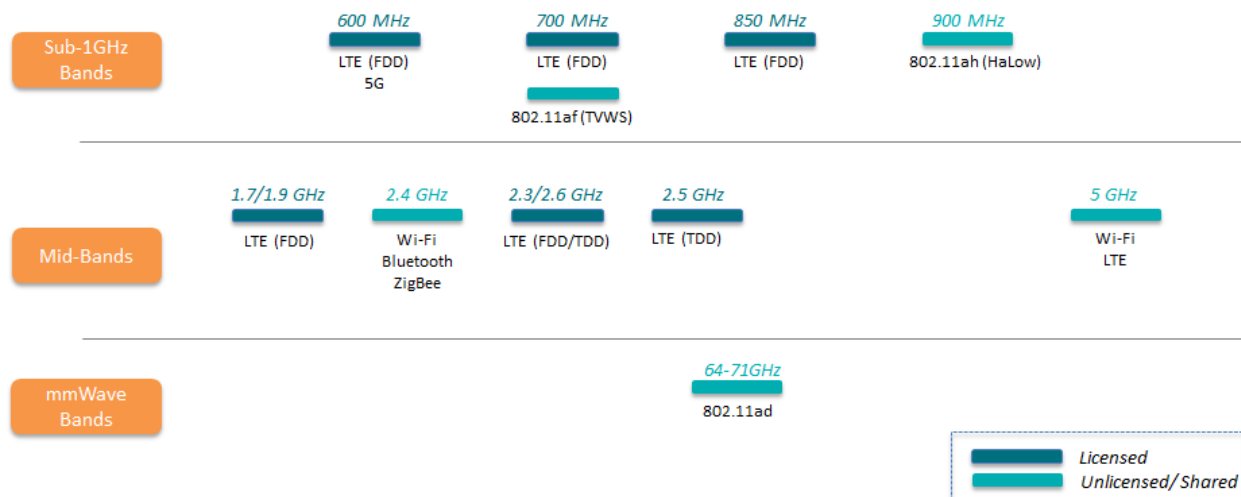


Figure 1: Existing US Spectrum Allocation

2. Greater Scope for Unlicensed Spectrum Use in a 5G World – More Bands, a Readymade Standard

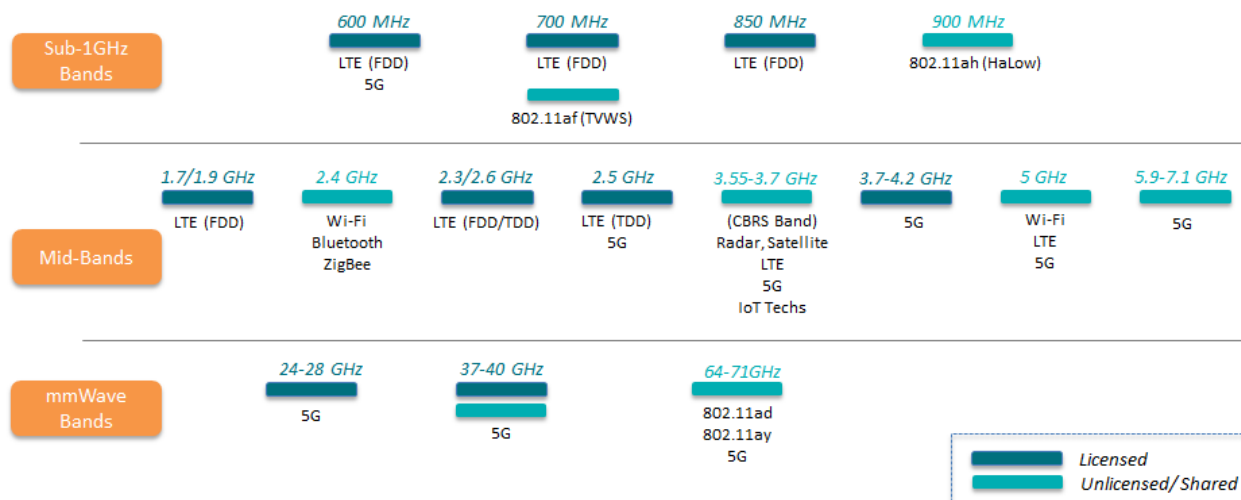


Figure 2: US Spectrum Allocation in a 5G World

As can be seen from Figure 2, the 5G landscape ushers in several additional unlicensed and lightly licensed bands that augment spectrum capacity significantly.

Several additional bands have been allocated for 5G in the U.S. In addition to newly opened-up licensed bands, unlicensed or shared spectrum is available in both mid-bands (between 1 and 6 GHz) and mmWave bands (above 24 GHz). Included in the mid-band range is the 3.5 GHz band opened up by the Federal Communications Commission (FCC) for spectrum sharing – also known as Citizens Broadband

Radio Service (or CBRS). CBRS opens up to 150 MHz of spectrum for lightly licensed and opportunistic use by different technologies – LTE, 5G, IoT technologies and others.

It is worth noting that the 5G New Radio (NR) standard has been designed from the outset to operate in most of these unlicensed bands. We can expect to see 5G networks and devices that are ready from the get-go to leverage unlicensed spectrum more expansively – as opposed to today, where LTE devices and infrastructure have had to be adapted to operate in the 5 GHz and CBRS bands, i.e. as an afterthought.

Figure 2 also indicates that multiple wireless technologies will need to coexist in many of these bands. The 3.5 GHz (CBRS) band is likely to be shared by several technologies. The same is the case with the mmWave 60 GHz band – where 802.11ad, 802.11ay (Wi-Fi's own mmWave technology) and 5G will coexist.

3. 5G Technology Is Unlicensed Spectrum Friendly from the Outset

5G NR incorporates capabilities that lend themselves to efficient unlicensed operation, and should enable 5G to coexist harmoniously with other technologies in the same bands:

- Support of a variety of unlicensed and shared bands by design.
- Band-agility: 5G radios will be able to switch bands nimbly to support specific deployment scenarios and service types.
- 5G NR supports flexible numerology, i.e. physical layer parameters are reconfigurable on the fly to support different operational scenarios and services. For example:
 - Switching from band to band, and reconfiguring to enable optimal operation in the new band;
 - Moving between indoor and outdoor scenarios;
 - Supporting different service types, e.g. applications with latency-critical requirements, high-bandwidth needs, networks with large numbers of devices (e.g. IoT) etc. 5G allows services with different numerologies to be multiplexed within the same carrier (or portion of a carrier).
- 5G NR provides millisecond or sub-millisecond latencies and Gigabit throughputs, which means that a 5G device can switch in and out of a channel in short time, freeing up the channel for use by other devices.
- Network slicing capabilities that allow sharable bandwidth pools to be divvied up and allocated to devices and services in an equitable fashion.

4. Spectrum Sharing Paradigms

Sharing of unlicensed spectrum can be done in multiple ways. Vertical sharing involves a multi-tier structure, where upper tiers have priority over, and are protected from, lower tiers. Examples of vertical sharing models are CBRS and Licensed Shared Access (LSA), where an incumbent system has highest priority access to resources, and lower tiers defer to layers above them (including the incumbent) while trying to access the system. Access to these systems is typically arbitrated by a central database entity.

With horizontal sharing, all systems sharing the spectrum have the same access priority, and access is typically governed by a well-defined access etiquette. ISM band operation is a good example of horizontal sharing. Multiple Wi-Fi and LTE/5G systems can coexist in 5 GHz, without a central coordinating entity. Hybrid models can incorporate elements of horizontal and vertical sharing.

Unlicensed Spectrum Use Cases

Several interesting use cases exist today, and we can expect to see many more in a 5G landscape. Use cases entail standalone unlicensed operation, as well as integrated licensed / unlicensed scenarios.

1. Next-Generation Home Wi-Fi System



Figure 3: Next-Generation Home Wi-Fi

Next-generation Wi-Fi technologies (802.11ax and 802.11ay) will form the cornerstones of home connectivity and coverage in a 5G world. 802.11ax is Wi-Fi's next jump forward. Apart from providing gigabit bandwidths, 802.11ax is designed to enable more efficient use of bandwidth, and provide higher capacity and interference resistance in dense Wi-Fi deployments. 802.11ax can operate in both 2.4 and 5 GHz bands. It incorporates cellular-like Orthogonal Frequency-Division Multiple Access (OFDMA) and scheduling features, as well as Multi-User MIMO, which enable multiple users to simultaneously use a Wi-Fi channel – supporting a larger number of simultaneous users.

802.11ay is Wi-Fi's next-generation mmWave technology, designed to operate in the 60 GHz band. 802.11ay is ideal for ultra-high bandwidth, short-range communication between devices that carry high-definition video. 802.11ay can be used as an in-room technology in different parts of a home, whereas 802.11ax can provide umbrella coverage across the home and the broadband conduit in and out of the home. 802.11ay's Fast Session Transfer capability can switch a connection between 60 GHz and 2.4 / 5 GHz bands, as the user moves around in the house.

2. Community Wi-Fi Hotspot Network

Community Wi-Fi allows a service provider to craft a crowd-sourced public hotspot network, leveraging unused capacity on existing Wi-Fi infrastructure (e.g. residential and enterprise routers). A service provider can also use this excess capacity to offer retail and roaming services to subscribers of partner operators.

Essentially, each component router in the hotspot network emits two signals – a public signal (public Service Set Identifier or SSID) and a private SSID. The residential user accesses their system using the private SSID, while a roamer uses the public SSID. The two signals are firewalled from one another, so that the respective data connections are secure. Roaming users are only allowed to use the Wi-Fi network capacity that is currently not used by the residential user.

Community Wi-Fi deploys hotspot networks multiple times faster than traditional hotspots, and is expected to evolve to leverage next-generation Wi-Fi and 5G technologies to boost speeds and support greater densification.

3. Licensed – Unlicensed Aggregation

This approach allows a service provider to combine LTE or 5G technology with Wi-Fi to augment capacity. Using LTE/5G and Wi-Fi radios simultaneously on the device and infrastructure ends, it is possible to split a single bearer (or traffic flow) across cellular and Wi-Fi links, based on policy and channel conditions. This allows applications to use both cellular and Wi-Fi links simultaneously, and dynamically move traffic between the two links based on changing radio conditions - providing significant performance gains.

This capability has been standardized by the cellular standards body 3GPP (3rd-Generation Partnership Project) as LTE-LWA (LTE-WLAN Aggregation). A 3GPP study on 5G operation in unlicensed spectrum is looking to do the same with 5G.

4. LTE or 5G Operation Standalone in Unlicensed or Shared Spectrum

This use case entails operating LTE or 5G standalone in an unlicensed or shared spectrum band. Flavors include LTE operating solely in the 5 GHz band (also known as MulteFire), LTE operating in the CBRS (3.5 GHz) band, and 5G NR in any of the allocated unlicensed bands allocated to 5G. This approach combines the performance benefits of LTE with the simplicity of Wi-Fi-like deployments.

This model provides particular value to service providers without cellular licenses, e.g. cable operators can deploy LTE small cells to provide outdoor coverage and capacity. It is also suitable for private LTE or 5G networks deployed by enterprises, venues and industrial IoT providers, and can also be used for neutral host deployments that serve multiple operators.

5. Licensed Assisted Unlicensed Access

LTE-LAA (LTE Licensed Assisted Access) augments a service provider's licensed band LTE service by also utilizing the unlicensed 5 GHz band. An anchor licensed carrier (carrying control and data traffic) is combined with an unlicensed carrier that carries data. This approach boosts capacity for an operator with access to licensed spectrum. Unlike the aggregation case discussed previously, the LTE radio actually

operates in unlicensed spectrum, and hence must be able to coexist harmoniously with Wi-Fi in the same band. The 3GPP study on 5G operation in unlicensed spectrum is looking to do the same with 5G.

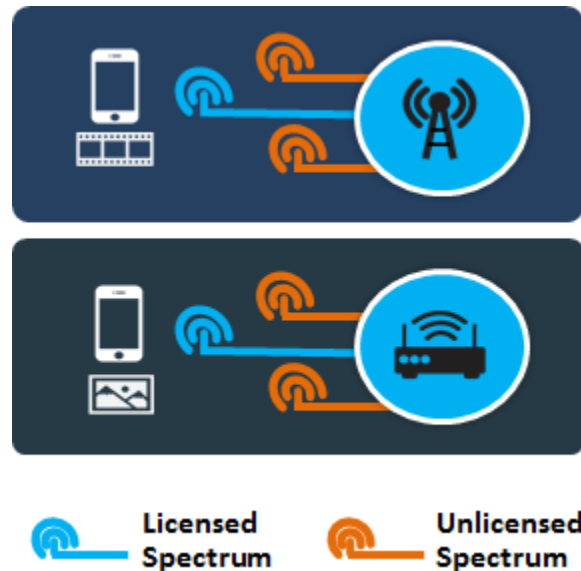


Figure 4: LTE/5G in Licensed + Unlicensed Spectrum

6. Vertical Spectrum Sharing – CBRS

The three-tier Citizens Broadband Radio Service (CBRS) is a well-defined initiative opened up by the FCC to allow lightly licensed and unlicensed use of lightly used spectrum in the 3.55 – 3.7 GHz band. The bandwidth available for sharing is significant – a total of 150 MHz of spectrum is usable when free.

CBRS is a vertical spectrum sharing framework with three access tiers:

- Tier 1 is occupied by incumbents, i.e. military radar, fixed satellite systems
- Tier 2 (Priority Access Layer or PAL) is licensed, and 70 MHz of spectrum is allocated to PAL layers. Up to seven 10 MHz channels can be allocated per “census tract” (about the size of a small town) and awarded to the highest bidders for three-year periods. Tier 2 users can access channels that are not in use by Tier 1 users
- Tier 3 is General Authorized Access (GAA), where no license is required – and is open to opportunistic use. GAA users can access channels that are not in use by Tier 1 or PAL users.

Enforcement of access rules and arbitration of spectrum access is done by a Spectrum Access System (SAS) database. The fact that PAL licenses will be granted for smaller geographical areas (census tracts) than traditional cellular licenses makes these licenses much more affordable for smaller players. This, in fact, is one of the FCC’s goals with CBRS – to stimulate wireless innovation and competition, and encourage smaller players.

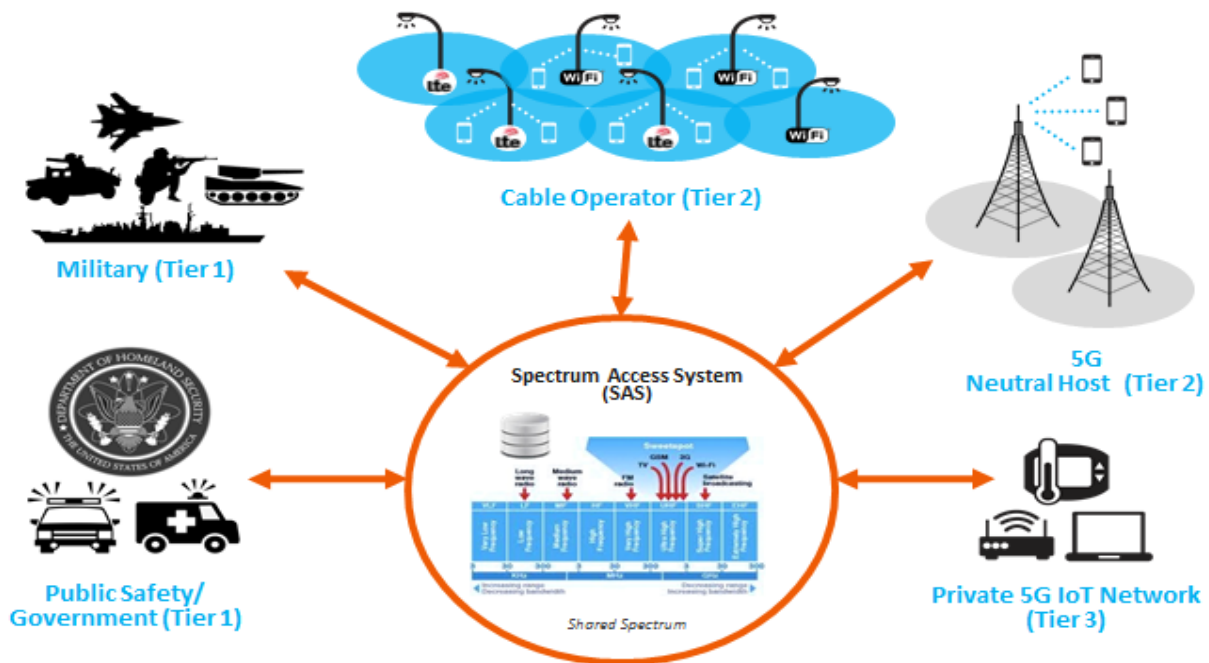


Figure 5: CBRS Framework and Use Cases

Figure 5 shows exemplary use cases that CBRS can support. Players without licensed spectrum like cable MSOs are good examples of Tier 2 licensees. A cable co. can leverage PAL licenses to provide outdoor service using small cells, as an example. The propagation characteristics of 3.5 GHz make it highly suitable for small cell operation. Enterprises and venue owners can acquire PAL licenses and create private LTE or 5G networks. Industrial IoT is a good example of a Tier 3 service. It is worth noting that PAL users can add to their Tier 2 allocation by accessing the system at Tier 3 when possible.

Unlicensed Spectrum - Performance Challenges

As multiple technologies start to use unlicensed bands and deployments densify, congestion and interference become significant issues. This is driven by multiple factors:

- Multiple technologies occupying the same spectrum bands. We already see this within the 2.4 GHz band, where Wi-Fi, Bluetooth and ZigBee operate. The 5 GHz band is set to see LTE operate in it, in addition to Wi-Fi, with perhaps 5G NR in the future as well. The CBRS band is likely to see diverse technologies having to coexist in it.
- Dense deployments will add to the problem. In an area with a large number of devices deployed close to one another, there can be severe contention for radio resources.
- Multiple service providers delivering service in the same areas add to the densification issue.

What this brings about is congestion and potential interference. When a large number of devices, and multiple technologies, contend for the same radio resources, channel access wait times go up, and performance metrics get affected. Latency and jitter go up, throughput reduces, and QoS gets affected.

The other performance variable that needs to be managed is coverage. While the higher frequency bands (e.g. mmWave) support high bandwidths, it is typically at the expense of propagation range. Poor

coverage (as a result of a device experiencing low signal strength) results in low throughputs, high error rates and high latencies.

These issues are observable with Wi-Fi deployments today, and will affect other technologies attempting to coexist in various unlicensed bands.

Addressing Performance Challenges

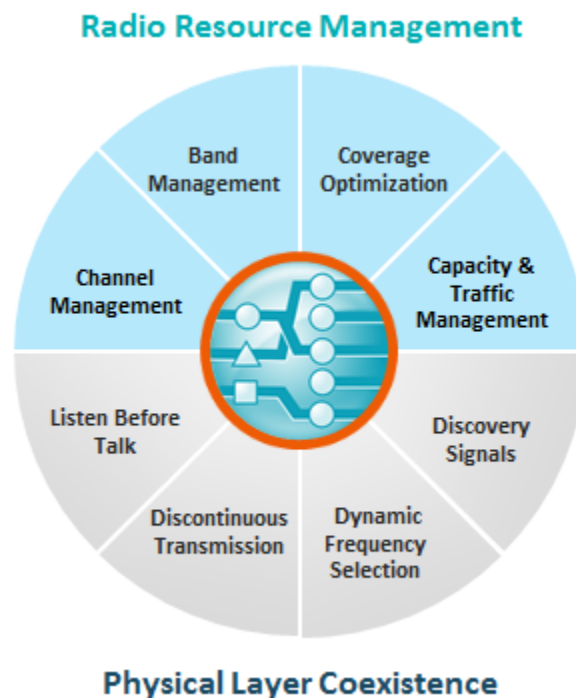


Figure 6: Addressing Performance Issues at Two Levels

Issues have to be addressed at two levels – via physical layer coexistence mechanisms as well as Radio Resource Management (RRM) schemes. With Wi-Fi, physical layer mechanisms like Listen Before Talk (LBT) have been used to prevent devices from stomping on one another via simultaneous transmissions in the same channel. LBT, Discontinuous Transmission and other Physical layer techniques have been adopted by LTE-LAA, which aims to operate in unlicensed bands.

At Fontech, our observations with Wi-Fi have indicated the critical need for Radio Resource Management (RRM) to resolve congestion and coverage issues. Techniques for intelligent allocation of channels, steering of client devices to a less congested band, power control and client steering become critical in making the system work well in contentious environments. RRM plays a proactive, preventative role – it preempts congestion and coverage from becoming issues, via smart allocation of radio resources. RRM is complementary to physical layer schemes, which arbitrate access to resolve issues with the use of the resources.

The remainder of the document describes the results of tests conducted with Wi-Fi that characterized the positive impacts of RRM in resolving performance issues. Three categories of RRM functions are

described here: dynamic channel management, band steering and client steering. While these results are for Wi-Fi, they are relevant to multi-technology coexistence scenarios in different unlicensed bands.

1. Setup

The tests were conducted in collaboration with a prominent North American cable operator, in a college dormitory environment. This was a field trial with a live Wi-Fi network, with Fontech's cloud-based RRM solution managing 100 dorms and 75 Wi-Fi Access Points (APs) – a dense deployment scenario. The system was run for two weeks with RRM functionality switched off, and for the subsequent weeks with RRM enabled – to obtain a clear performance comparison. The Wi-Fi network experienced real-life usage by students through the entire period.

2. Dynamic Channel Management

Dynamic channel management algorithms mitigate channel congestion and interference. They detect developing congestion; if congestion builds up to a configurable threshold level, the algorithms select a less congested channel and switch the AP and its clients over to the selected channel. The goal here is to maintain good service quality, by not allowing key service quality metrics to degrade to poor levels.




Figure 7: Reduced Congestion Spikes with RRM

Figure 7 shows the first set of results. The monitoring system counted the number of congestion spikes (i.e. the number of times the channel usage level exceeded 80% for any Wi-Fi channel) per week. (Channel usage indicates the percentage of time a channel has been busy processing traffic). When RRM was off (the first two weeks), the number of spikes was very high – reaching over 1600 in Week 1. When RRM was enabled, congestion reduced appreciably – the number of spikes dropped by more than 50%. When RRM thresholds were further optimized (prior to Week 6), the number of spikes dropped to a fraction of what it had been without RRM.

Figure 8 depicts the number of occurrences of persistent congestion within the system. With RRM disabled, this number was high. Week 1 had 1,133 events – each of these was a congestion event that warranted a channel change. However, with RRM disabled, no channel change occurred, congestion conditions persisted and the events kept recurring. When RRM was enabled (Week 3), the number of congestion events dropped to negligible levels. This was because when channels got congested, RRM changed channels quickly and alleviated congestion, preventing recurrences of these events. Only 2 APs experienced any kind of congestion during the latter period.




Figure 8: Dramatically Reduced Number of Congestion Events with RRM

3. Band Steering for Bandwidth Optimization

Band steering moves associated clients to a different radio/frequency band to improve quality of experience. When an AP's radio is overloaded, band steering functionality moves client devices to another radio to mitigate the congestion scenario. Band steering can move client devices from mmWave bands to lower-frequency spectrum if coverage is an issue. Band steering can also steer clients with good coverage (strong signal strengths) automatically to the 5 GHz band to provide them much higher bandwidth. Figure 9 below illustrates the effects of this bandwidth optimization capability.




Figure 9: Bandwidth Optimization with Band Steering

The right-hand side of Figure 9 shows average traffic distribution between 2.4 and 5 GHz bands, during the period with RRM enabled. Over 80% of traffic stayed on 5 GHz. The algorithm keeps clients on 5 GHz, as long as signal strength is strong. The left side of the figure shows a “with / without RRM” comparison. The percentage of data on 5 GHz increased significantly after RRM was enabled. Our measurements indicated that overall throughput jumped by over 100% as a result of clients being steered to 5 GHz.

4. Client Steering (Device Mobility) for Coverage Optimization

Setup: 5 Apartments Set Up with 2 APs in Each




Figure 10: Inter-AP Client Steering to Optimize Coverage

The goal of client steering is to steer a client with poor coverage (e.g. low signal strength) to a better AP or radio. If the client is on 2.4 GHz and drifts out to a region with poor coverage, client steering steers it to a better AP (in a multi-AP deployment). If the client is on 5 GHz and starts to experience poor

coverage, it gets steered to the same AP's 2.4 GHz radio, or to another AP (if the network includes multiple APs). The aim is to provide the device a consistently good QoS as it moves around.

Within the dorm field trial, 5 two-level apartments were set up with 2 APs in each. Figure 10 shows the average time (%) spent by client devices on the two APs in each apartment. Clients moved around appreciably between APs in each apartment. Over 1,000 client steers were recorded. An average signal strength improvement of 18 dB was achieved per steering event; Signal-to-Noise ratio (SNR) improved by 17 on average per steering event. The data indicates that client steering improved coverage dramatically.

Conclusions

Unlicensed spectrum usage is expected to proliferate appreciably in the near future. In a 5G world, unlicensed spectrum will not just be the domain of unlicensed technologies like Wi-Fi. We can expect to see diverse technologies – licensed and unlicensed – operating in this space. A lot more spectrum will be available for unlicensed use – new bands allocated to 5G, as well as shared bands such as CBRS.

In addition, 5G NR technology has been designed from the outset to be spectrum sharing friendly. 5G NR is highly band-agile, and supports flexible reconfigurability to optimize itself for diverse bands, deployment scenarios and service types. Unlicensed operation is likely to be the building block for several interesting uses cases – next-generation Home and Community Wi-Fi, Licensed / Unlicensed integration, private LTE / 5G networks, neutral host models and vertical spectrum sharing scenarios.

Considering all factors – more bands, unlicensed-friendly technology, use cases – it is easy to see that unlicensed spectrum use is likely to grow significantly in the coming years. But this does not come without issues. Coexistence of multiple technologies in these bands, high usage and dense deployments will create performance issues – congestion, interference etc. – unless access to the spectrum is coordinated cohesively.

Hence, the use of RRM techniques will become critical. Optimization schemes such as dynamic channel management, band steering and client steering can help mitigate congestion and coverage issues, and optimize service quality. This paper illustrates impacts RRM can have in this regard.

Abbreviations

3GPP	3 rd Generation Partnership Project
5G	5 th Generation
AP	Access Point
CBRS	Citizens Broadband Radio Service
DVR	Digital Video Recorder
FCC	Federal Communications Commission
GAA	General Authorized Access
HDTV	High-definition television
IoT	Internet of Things
ISM	Industrial, Scientific and Medical (band)

LBT	Listen Before Talk
LTE	Long Term Evolution
LTE-LAA	LTE Licensed Assisted Access
LTE-LWA	LTE WLAN Aggregation
LSA	Licensed Shared Access
MIMO	Multiple Input Multiple Output
mmWave	Millimeter-wave
MSO	Multiple System Operator
NR	New Radio
PAL	Priority Access Layer
QoS	Quality of Service
RRM	Radio Resource Management
RSSI	Received Signal Strength Indicator
SAS	Spectrum Access System
SCTE	Society of Cable Telecommunications Engineers
SNR	Signal-to-Noise Ratio
SON	Self-Organizing Networks

Bibliography & References

3GPP Draft Technical Report 38.889: Study on NR-Based Access to Unlicensed Spectrum

3GPP Technical Report 36.889: Feasibility Study on Licensed-Assisted Access to Unlicensed Spectrum

FCC: 3.5 GHz Band / Citizens Broadband Radio Service: Report and Order and Second Further Notice of Proposed Rulemaking

HFC Evolution

The Best Path Forward

An Operational Practice prepared for SCTE•ISBE by

Nader Foroughi
Network Engineer
Shaw Communications
2728 Hopewell Place, NE, Calgary AB
+1 403 648 5937
nader.foroughi@sjrb.ca

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	4
1. Scope	4
2. Analysis	5
2.1. Utilization and Capacity Analysis	5
2.1.1. Assumptions:.....	5
2.1.1.1. Plant Bandwidth:	5
2.1.1.2. Spectrum:.....	5
2.1.1.3. Available Capacity vs Peak Utilization:	5
2.1.1.4. OFDM Modulation Order:.....	5
2.1.1.5. DS Compound Annual Growth Rate:	6
2.1.1.6. Remote PHY (RPHY):.....	6
2.1.1.7. Full Duplex DOCSIS:	6
2.1.2. Results: 6	
2.2. Net Present Value Analysis:.....	11
2.2.1. Assumptions:.....	13
2.2.1.1. Plant Design:.....	13
2.2.1.2. Node and Amplifier Costs:	13
2.2.1.3. Net Present Value (NPV) Rate:	13
2.2.2. Results: 13	
2.2.2.1. Updated Capacity Trend:	17
2.2.2.2. A More Optimistic Approach:	19
Conclusion.....	20
Abbreviations	21

List of Figures

Title	Page Number
Figure 1 – Node DS Capacity –100% Buffer Case.....	7
Figure 2 – Node DS Capacity – 50% Buffer Case	8
Figure 3 – Current Shaw Spectrum Plan	8
Figure 4 – Projected Shaw N+2 Spectrum Plan	9
Figure 5 – Projected Shaw N+2 FDX Spectrum Plan	9
Figure 6 – Projected Shaw N+0 Spectrum Plan	10
Figure 7 – Projected Shaw ESD Spectrum Plan.....	10
Figure 8 – VC819C	12
Figure 9 – VC819C N+0 Design	14
Figure 10 – VC819C N+2 Design	15
Figure 11 – VC819C N+2 to N+0 Design.....	16
Figure 12 – Updated Capacity Trend – 100% Buffer	18
Figure 13 – Updated Capacity Trend – 50% Buffer	18

Figure 14 - Updated Capacity Trend – Optimistic Case	20
--	----

List of Tables

Title	Page Number
Table 1 – 1k QAM Bit Rate	6
Table 2 – 256 QAM Bit Rate	6
Table 3 – Node Parameters	11
Table 4 – Build Cost for Each Plant Design.....	17
Table 5 – Capacity Trend Summary Table	19
Table 6 – NPV Analysis – 100% Buffer Case	19
Table 7 – NPV Analysis – 50% Buffer Case	19
Table 8 - NPV Analysis – Optimistic Case.....	20

Introduction

The majority of MSOs have an N+X outside plant architecture. Assuming that Fibre to the Premises (FTTP) is the final state of the plant, there are varieties of approaches being considered by each provider to increase bandwidth (BW) in the meantime to compete with fibre-based services. Some are considering a leap directly to a fully passive state (N+0), whereas others are considering reducing amplifier cascades gradually, with a passive state in mind. At Shaw we are contemplating an initial move to N+2, meaning the plant is going to be split directly to an N+2 state from its current architecture.

In this paper an analysis has been carried out to evaluate the potential advantages and disadvantages of going directly to a passive (N+0) architecture versus reducing amplifier cascades to a mid-point (N+2) prior to going to N+0, with a long-term goal of FTTP in mind, in both cases. This is assuming that Full Duplex DOCSIS (FDX) will be developed in a cascaded environment in the near future.

Due to the fact that business as usual (BAU) node splits, based on plant congestion, are not scalable, they have been excluded from the comparison.

Based on the current downstream capacity offerings and projected future growth, the difference in capacity between an N+0 and N+2 plant has been evaluated while taking into consideration the various new technologies that will be deployed in the near future.

Furthermore, a net present value analysis has been provided for the transition from N+2 to N+0. At its current state, 75% of Shaw's plant consists of nodes with a longest cascade of 5 or less amplifiers. Depending on when the transition to N+2 or N+0 is projected to occur, a relative estimate for the net present value of the costs has been provided, based on the sample plant selected.

Based on the analysis shown in this paper, assuming that Full Duplex DOCSIS (FDX) is developed in a low-cascade architecture such as N+2, the results show that moving to N+2 → N+2 FDX → N+0 FDX has a lower total cost of ownership (TCO), in comparison to moving directly to N+0 FDX.

Content

1. Scope

In order to quantify the differences between various deployment strategies, two major categories have been considered in this paper:

1. Downstream plant capacity and peak utilization analysis
2. Overall cost and net present value analysis

These are both based on the assumption that the ultimate state of the plant will be FTTP and all other deployments strategies are in-between stages to increase plant capacity to be able to compete with fibre based services.

2. Analysis

2.1. Utilization and Capacity Analysis

In order to provide a utilization and capacity estimate, the various plant stages should be elaborated on.

Currently Shaw's plant is primarily N+X. In this paper the primary focus is on the 75th percentile of the largest number of amplifiers in cascade, which is N+5. Assuming this is the current state and FTTP being the final stage of the plant, the in-between stages have been considered to be:

N+5 → N+2 → N+2 FDX → N+0 FDX → N+0 Extended Spectrum DOCSIS → FTTP

This paper analyses the feasibility of moving to an N+2 FDX environment prior to moving to N+0 FDX. The assumptions for a cascaded environment FDX plant has been described in section 2.1.1.7.

Prior to outlining the details of these analyses, the assumptions for this analysis have been outlined below:

2.1.1. Assumptions:

2.1.1.1. Plant Bandwidth:

The following assumptions have been made regarding the capacity of N+X and N+0 plant

1. Maximum plant BW for N+2 has been assumed to be 1GHz
2. Maximum plant BW for N+0 has been assumed to be 1.2GHz

Although some of the amplifiers in the plant today are not 1GHz, the assumption has been made that the 750MHz and 860MHz amplifiers are going to be swapped out for 1GHz versions to increase plant BW and take full advantage of DOCSIS 3.1 and Orthogonal Frequency Division Multiplexing (OFDM) carriers.

2.1.1.2. Spectrum:

The capacity analysis in this paper considers the end state of the spectrum for each architecture. This means that IP TV is assumed to have been deployed.

2.1.1.3. Available Capacity vs Peak Utilization:

In order to satisfy the peak utilization, the overall available spectrum capacity requirement has been analysed in the two scenarios below:

- Double the peak utilized amount (worst case scenario) or 100% buffer case
- 50% more than the peak utilized amount (realistic case) or 50% buffer case

2.1.1.4. OFDM Modulation Order:

In order to have a basis for capacity calculations, all OFDM carrier modulation orders are assumed to be 1024QAM. This is based on the plant characterization tests and field observations that we carried out as a part of D3.1 deployment. Anything above this, namely 4096QAM, is considered "extra capacity" and will not be used in the capacity calculations demonstrated in the below sections. This is due to the fact that 4096QAM may not be achievable in certain portions of the plant.

The effective throughputs calculated for the 1024QAM OFDM carriers have been shown below. Note that this is a conservative estimate based on field observations:

Table 1 – 1k QAM Bit Rate

Modulation Rate	Effective Throughput (Bits/s/Hz)
1024QAM	7

2.1.1.5. DS Compound Annual Growth Rate:

At Shaw we have experienced a 36% CAGR.

2.1.1.6. Remote PHY (RPHY):

For the purpose of plant progression in this paper, an assumption has been made that the in-between stage of the plant consists of an RPHY node and amplifiers. The final stage of the plant being N+0, will have FDX nodes.

Further to the assumption mentioned above, the N+2 plant has been broken down into two categories:

- An N+2 plant where RPHY nodes and amplifiers will be deployed, but no FDX will be available to them
- A case has also been assumed for FDX deployable in an N+X environment (N+2 in this case). This has been further explained below.

2.1.1.7. Full Duplex DOCSIS:

The capacities for FDX have been based on the assumption that it will be fully deployed in the FDX band (108MHz – 684MHz), by the time plant reaches an N+2 state.

Please note:

- Currently FDX is only being discussed in an N+0 environment
- An assumption has been made for FDX deployable in a cascaded environment. Considering the complexities that would be present in this type of plant, the achievable modulation order has been assumed to be 256QAM.

Table 2 – 256 QAM Bit Rate

Modulation Rate	Effective Throughput (Bits/s/Hz)
256QAM	5

2.1.2. Results:

DOCSIS technology has advanced significantly in recent years, improving the spectral efficiency in the DS and US. This analysis has been based on peak utilization and the available spectrum to be able to provide the capacity and BW to satisfy the peak utilization.

The peak utilization per node is gathered at Shaw bi-monthly. The method of calculating the peak utilization is:

Peak Utilization

$$= \text{Number of available carriers} \times \text{Capacity of carriers (Mbps)} \\ \times \text{maximum \% utilized Sampled Every 5 minutes}$$

Based on the historical data available, the 75th percentile of the peak utilization for the past 5 years was calculated. The results produced a starting point of 508.66 Mbps in peak utilization. Referring to assumptions section 2.1.1.3, the capacity requirement for each scenario can be calculated as:

- 100% Buffer Case: 1017.3 Mbps
- 50% Buffer Case: 763 Mbps

Based on this, assuming a 36% CAGR, the graphs below can be produced to estimate the capacity required in the future:




Figure 1 – Node DS Capacity –100% Buffer Case




Figure 2 – Node DS Capacity – 50% Buffer Case

To be able to estimate when Shaw needs to reach each capacity point, the potential overall spectrum capacity in the various available stages needs to be estimated. The current Shaw spectrum plan has been shown below:




Figure 3 – Current Shaw Spectrum Plan

Transitioning to N+2 plant, the spectrum has been projected to look as below:




Figure 4 – Projected Shaw N+2 Spectrum Plan

Based on this and the D3.1 OFDM carriers starting at 648MHz, the capacity for N+2 plant can be estimated as:

$$\text{Shaw } N + 2 \text{ Capacity} = \text{D3.0 Capacity} + \text{D3.1 Capacity} = 32 \times 36\text{Mbps} + 2465\text{Mbps} \cong 3.6\text{Gbps}$$

DS/US Capacity: 10/1

DS Tier: Gigabit Services

Referring to the assumptions section 2.1.1.7, if FDX is to be developed in any N+X environment, the spectrum has been assumed to look as demonstrated below:

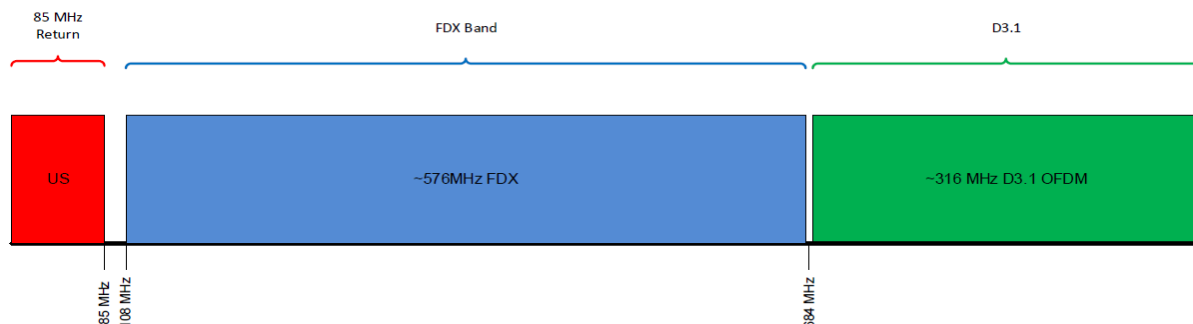


Figure 5 – Projected Shaw N+2 FDX Spectrum Plan

Based on the assumption of 256QAM achievable in the FDX DS band, as an end state, the capacity can be estimated as:

$$\text{Shaw } N + 2 \text{ FDX Capacity} = \text{FDX} + \text{D3.1} = 2880\text{Mbps} + 2212\text{Mbps} \cong 5 \text{ Gbps}$$

DS/US Capacity: 1.5/1

Tier:

- Gigabit symmetrical services

or

- Multi-gigabit DS & gigabit Upstream (US) Services

Transitioning to N+0, the spectrum has been assumed to look as below:




Figure 6 – Projected Shaw N+0 Spectrum Plan

Based on this, the capacity can be estimated as:

$$\text{Shaw } N + 0 \text{ Capacity} = \text{FDX} + \text{D3.1} = 4032\text{Mbps} + 3738\text{Mbps} \cong 7.8\text{Gbps}$$

DS/US Capacity: 2/1

Tier:

- Gigabit Symmetrical Services

or

- Multi-gigabit DS & gigabit US Services

Assuming no drastic changes will occur in the DS peak utilization and Shaw deploys extended spectrum DOCSIS (ESD), the spectrum will look as below:




Figure 7 – Projected Shaw ESD Spectrum Plan

Based on this, the capacity can be estimated as:

$$\text{Shaw Extended Spectrum } N + 0 \text{ Capacity} = \text{FDX} + \text{D3.1} = 4032\text{Mbps} + 7812\text{Mbps} \cong 12\text{Gbps}$$

DS/US: 1/1 (assuming static FDX US)

Tier: Multi-gigabit Symmetrical Service

Note that Figure 7 is not evaluated in this paper. It is simply inserted as a point of discussion for future architectures, as ESD gains more traction in the industry.

The capacities calculated above will be analyzed in detail in section 2.2.2.1 where they will be inserted in the capacity trends discussed earlier in this paper.

2.2. Net Present Value Analysis:

In this section the details of the net present value analysis, based on the progression of the plant has been provided.

The sample plant (VC-819C) that was selected for this paper consists of the parameters below:

Table 3 – Node Parameters

	Homes Passed	Trunk Amp.	Distr. Amp.	2 Way Tap	4 Way Tap	8 Way Tap	2 Way Splitter	3 Way Splitter	4 Way Splitter	Directional Coupler
VC819C	350	3	18	0	54	13	4	3	2	4

The node above was selected based on the 75th percentile of the largest number of amplifiers in cascade, in the top 3 biggest regions in Shaw. The reason why homes passed was not considered for the selection of this node is due to the fact that focusing on the number of amplifiers provides a more challenging environment for the node to be split down to N+2 and/or N+0, in comparison to focusing on homes passed (HP), due to the density factor. In other words, focusing on the number of amplifiers in cascade provides a reasonable-worst-case scenario.

The map for the selected node has been shown below:



Figure 8 – VC819C

The progression of this plant was planned in such a way to:

1. Go to N+0 from the current state
2. Go to N+2 from the current state, then N+0

The assumptions below were taken into consideration for this analysis:

2.2.1. Assumptions:

2.2.1.1. Plant Design:

- The N+2 plan was designed with N+0 in mind, this means that some of the N+2 node locations overlap with future N+0 locations. This lines up with Shaw's current node split strategy.
- In the N+2 design, the nodes are not optimized for reach. This means that the nodes were not centralized, and plant turn-arounds were avoided.
- The N+0 plan was designed to optimize the node location with minimal coaxial work required. This means existing amplifier locations were used to accommodate future node locations.

2.2.1.2. Node and Amplifier Costs:

Given that this analysis assumes the plant to be in its final stage for each of the items mentioned above, the nodes and amplifiers are assumed to be RPHY/FDX capable nodes and amplifiers. The cost of the node and amplifiers have been outlined below:

- FDX Node: Approximately twice as expensive as an optical node
- FDX Amplifiers: Approximately the same cost as an optical node

Note: Head-end/Hub costs have not been included in any of the estimates.

2.2.1.3. Net Present Value (NPV) Rate:

For the NPV analysis, the discount rate has been assumed to be 8%.

2.2.2. Results:

The plant design for each case has been demonstrated below:

N+0 Design:



Figure 9 – VC819C N+0 Design

Each coloured circle in the figure above is a future node location, with its boundary highlighted.

Total Node count: 9

N+2 Design




Figure 10 – VC819C N+2 Design

Total Node count: 3

N+2 to N+0 Design:



Figure 11 – VC819C N+2 to N+0 Design

Total Node count: 10

The build cost for each case has been shown in Table 4. Note that the N+2 design has been split to two categories. Non FDX N+2 and FDX N+2. The reason behind this is the fact that the FDX amplifiers will be more expensive than the non-FDX (RPHY) versions, as mentioned in the assumptions above.

The table below outlines the overall build cost for each case:

Table 4 – Build Cost for Each Plant Design

Design	Build Cost (1000\$)
N+2 (non-FDX)	77
N+2 (FDX)	120
N+2 to N+0 (FDX)	242
N+0 (FDX)	298

Assuming a fixed yearly budget, from the table above shows that:

- N+2 non-FDX can be reached 3.5 times faster in comparison to N+0
- N+2 FDX can be reached 2.5 times faster in comparison to N+0

In order to carry out a net present value analysis, the dates where the projected capacity trends meet the required capacity demands, for each case, should be estimated.

2.2.2.1. Updated Capacity Trend:

Referring back to Figure 1 and Figure 2, they can be updated respectively, based on the resulting number of nodes in each design scenario, in the section above. The starting points for the 100% and the 50% buffer cases can be calculated as:

- 100% Buffer Case:
 - $N + 2 \text{ Starting point} = \frac{1017.3}{3} = 339.1 \text{ Mbps}$
 - $N + 0 \text{ Starting point} = \frac{1017.3}{9} = 113.03 \text{ Mbps}$
- 50% Buffer Case:
 - $N + 0 \text{ Starting point} = \frac{763}{3} = 245.33 \text{ Mbps}$
 - $N + 2 \text{ Starting point} = \frac{763}{9} = 84.77 \text{ Mbps}$

These can be entered into figures 1&2, as demonstrated below:




Figure 12 – Updated Capacity Trend – 100% Buffer




Figure 13 – Updated Capacity Trend – 50% Buffer

From Figure 12 and Figure 13, the dates where capacity trends will reach N+2 FDX for each buffer case can be estimated. These have been demonstrated in the table below:

Table 5 – Capacity Trend Summary Table

Date		
Case	100% Buffer	50% Buffer
N+2 FDX	2026	2027
N+0 FDX	2029	2032

The dates above can be utilized to estimate the net present value of each scenario:

Table 6 – NPV Analysis – 100% Buffer Case

	NPV	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6
N+2 (FDX)	\$281,000	\$77,000	\$43,000	0	0	0	\$242,000
N+0 (FDX)	\$298,000	\$298,000	0	0	0	0	0

Table 7 – NPV Analysis – 50% Buffer Case

	NPV	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7
N+2 (FDX)	\$269,000	\$77,000	\$43,000	0	0	0	0	\$242,000
N+0 (FDX)	\$298,000	\$298,000	0	0	0	0	0	0

In the tables above the \$242,000 for moving to N+0 has been placed in ‘year 6’ and ‘year 7’ for the 100% and 50% buffer cases, since in Figure 12 and Figure 13, the capacity trend will reach the N+2 FDX limit by 2026 and 2027 respectively. Also note that the \$120,000 for moving to N+2 FDX has been broken down to an initial \$77,000 to move to N+2 and an additional \$43,000 to be spent the year after, since the technology is estimated to be available at that time.

It can be seen from both Tables 6 and 7 that N+2 FDX has a lower TCO by roughly:

- \$17,000 in the 100% buffer case
- \$30,000 in the 50% buffer case

2.2.2.2. A More Optimistic Approach:

So far, the parameters in this paper have been selected conservatively. Switching to a more optimistic approach, with the same method used above, the results can vary significantly. To demonstrate this, consider the parameters below, instead of the ones used so far:

- CAGR = 30%
- Buffer = 50%
- DS FDX N+X modulation order = 1k QAM
- Spectrum BW: 1.2GHz and 1GHz have been shown

As a result, the capacity trend has been shown below:




Figure 14 - Updated Capacity Trend – Optimistic Case

Based on Figure 14, the NPV table will be:

Table 8 - NPV Analysis – Optimistic Case

	NPV
N+2 (FDX)	\$213,000
N+0 (FDX)	\$298,000

It can be seen that adjusting the DS modulation order to 1kQAM in conjunction with adjusting the CAGR to 30%, as it's been decreasing, can reduce the overall cost of N+2 FDX by \$68,000 in comparison to the 100% buffer and by \$56,000 in comparison to the 50% buffer conservative cases.

Conclusion

According to the analysis carried out on the node, selected based on the 75th percentile of the largest number of amplifiers in cascade in Shaw's plant, and assuming that FDX is developed in an N+2 environment, moving to an N+2 FDX mid-point has a lower TCO, in comparison to moving directly to N+0 FDX. Furthermore, assuming a fixed yearly budget, N+2 can be reached 3.5 times and N+2 FDX 2.5 times faster, in comparison to N+0 FDX.

This was shown to be valid in the reasonable worst-case scenarios demonstrated, in which case, moving to N+2 FDX also provides Shaw with an additional 1.5Gbps in DS capacity, in comparison to N+2. This enables the capability of offering gigabit symmetrical services. Furthermore, moving to N+2 FDX secures Shaw's spectrum capacity until 2027, assuming we maintain a 50% buffer above the peak utilized speed and a 36% CAGR. This is prior to having to move to N+0 FDX to gain an additional 2.8Gbps in DS capacity.

Abbreviations

BW	bandwidth
BAU	business as usual
CAGR	compound annual growth rate
DS	downstream
DOCSIS	data over cable service interface specification
ESD	extended spectrum DOCSIS
FDX	full duplex DOCSIS
FTTP	fibre to the premises
HP	homes passed
NPV	net present value
OFDM	orthogonal frequency division multiplexing
RPHY	remote PHY
TCO	total cost of ownership
US	upstream

How An MSO Can Leverage SD-WAN To Grow Its Enterprise Revenue

Impact of Rapid Penetration of Overlay Connectivity and Pull-Through Services

A Technical Paper prepared for SCTE•ISBE by

Narayan Raman

Principal and DMTS
Bell Labs Computing/Nokia
600 Mountain Avenue, Murray Hill, NJ 07974, U.S.A.
narayan.raman@bell-labs-consulting.com

Yadhav Krishnan

Senior Consultant
Bell Labs Computing/Nokia
601 Data Drive, Plano TX 75075, U.S.A.
yadhav.krishnan@bell-labs-consulting.com

Miguel Hernandez

Senior Consultant
Bell Labs Computing/Nokia
740 Waterside Drive, Aztec West, Bristol, BS32 4UF, U.K.
miguel.hernandez@bell-labs-consulting.com

Furquan Ansari

Partner
Bell Labs Computing/Nokia
601 Data Drive, Plano TX 75075, U.S.A.
furquan.ansari@bell-labs-consulting.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Study 3	
1. Challenges faced by MSOs	3
2. How SD-WAN addresses these challenges.....	4
3. SD-WAN revenue potential	5
3.1. SD-WAN connectivity market assessment	5
3.2. MSO revenue projection	10
3.3. Revenue sensitivity	12
4. SD-WAN Business Case Evaluation.....	13
Conclusion.....	16
Abbreviations	17
Bibliography & References.....	17

List of Figures

Title	Page Number
Figure 1 - SD-WAN Connectivity Projection	6
Figure 2 - Enterprise Site Map by Existing WAN Service	6
Figure 3 - Growth of Existing WAN Service Connections.....	7
Figure 4 - Sensitivity of Connections Market to Speed of Diffusion.....	9
Figure 5 - SD-WAN Connectivity Map	11
Figure 6 - MSO Revenue Projection	12
Figure 7 - Sensitivity of MSO Revenue Projection.....	13
Figure 8 - Business Case Components	14
Figure 9 - SD-WAN Solution Components.....	14
Figure 10 - SD-WAN Business Case Financials.....	15
Figure 11 - SD-WAN Business Case NPV Sensitivity	16

Introduction

Software Defined Wide Area Networking (SD-WAN) services are gaining increasing traction in the enterprise communication market because of the confluence of two developing trends – growing cost and complexity of WAN connectivity services, and emergence of virtualization-enabled network programmability. Enterprise IT managers must contend with increasingly complex branch office communication needs being shaped by significant bandwidth growth, emergence of cloud services, proliferation of end points with diverse connectivity requirements and the need to ensure adequate WAN protection against emerging threats, while holding firmly to the expense budget.

Simultaneously, however, network programmability based on policy-based routing, network function virtualization, and analytics-driven zero-touch automation has greatly simplified traditional WAN management. SD-WAN captures these functionalities through a range of alternatives that tailor the WAN solution to specific enterprise needs cost effectively. Enterprises can move low-priority sites from Multi-Protocol Label Switching (MPLS) to internet-only connectivity, augment MPLS with internet to optimize local traffic distribution at selected sites or deploy internet breakout at some sites for direct cloud connectivity. SD-WAN also enables flexible branch networking with deployments varying from “thin” end points, which focus primarily on connectivity, to “thick” ones that additionally provide value-add services.

While SD-WAN benefits from the complementarity of these two trends, an MSO needs to evaluate the resulting economics for itself considering its target enterprise market and the required networking and platform expenses. To this end, after briefly describing current enterprise WAN challenges and discussing how SD-WAN addresses them, this paper focuses on how a typical MSO can develop SD-WAN’s revenue and profit potential projections and enhance them through sensitivity analyses of key demand and cost parameters. A distinctive aspect of the proposed model is its utility-based approach that correlates an enterprise’s SD-WAN adoption to its value – measured in terms of desired WAN attributes – relative to its current WAN service. We present an illustrative case study for a generic MSO; the results indicate that SD-WAN can significantly impact the operator’s enterprise revenue while delivering an attractive business case with a payback period of 2.25 years. The proposed modeling approach is extendible to other operators, regions and markets worldwide.

Study

1. Challenges faced by MSOs

An MSO faces increasing challenges along multiple dimensions for its enterprise business.

Market threats: MSOs have a growing presence in the enterprise segment that is largely limited to small and medium businesses (SMBs). They face declining revenues per bit, while public and 3rd party cloud providers present new challenges due to their high levels of service ubiquity and agility, in addition to strong pricing economics.

MSOs are also constrained by their limited connectivity to enterprise customers; the high cost of cable rollout severely limits new network build. While Data Over Cable System Interface Specification® (DOCSIS) 3.1 has significantly expanded the range of their services, they remain threatened by the expanding fiber footprint of telecom operators.

WAN complexity: Current static and inflexible MSO networks are unable to cope with growing WAN complexity faced by enterprises because of increasing site bandwidth needs, diverse branch connectivity requirements and seamless networking with public and 3rd party clouds, while ensuring appropriate scalability. An enterprise's ability to rapidly add new branches, terminate temporary locations, and move functionalities across sites requires a level of network programmability and location-independence of routing and other network functions that current MSO networks are unable to offer.

Traditional MSO WAN services also do not provide adequate network visibility to enterprises, nor do they lend them the required network control and monitoring capability. This is a key reason why MSOs are unable to attract many enterprises who prefer to remain unmanaged instead of ceding control to traditional managed services.

WAN economics: As enterprise IT managers seek to reduce their networking costs in the face of growing bandwidth usage, MSOs are ideally positioned to offer solutions to help them. To avail of this market opportunity, MSOs are challenged to come up with innovative approaches to improving branch networking efficiency through leaner, more automated operations, and by realizing the economic benefits provided by new technologies, such as Software-Defined Networking (SDN) and Network Function Virtualization (NFV).

Increasing security risks: Network virtualization and increasing use of cloud services also open new areas of vulnerability. Flexible, service-specific protection is needed against threats from untrusted sources, new and larger attack surfaces, more frequent software and life cycle management updates, and rapid growth in east-west traffic. Greater customer usage of cloud services, and consequent regulatory considerations such as General Data Protection Regulation (GDPR), also implies a greater need for MSOs to securely segment an enterprise's operations while keeping overall WAN complexity and cost in check.

Low service agility: Many traditional MSOs and other network operators carry the legacy of slow sales cycles, high overhead costs and complex product catalogs that constrains their ability to rapidly offer new services. Expensive and time-consuming truck rolls and manual integration steps for deploying new sites, and the strong dependence of major provisioning and assurance processes on manual tasks further diminish MSOs' service agility.

2. How SD-WAN addresses these challenges

SD-WAN helps mitigate many issues described in Section 2.

Market opportunity: SD-WAN opens new revenue opportunities by expanding service portfolios and markets. It enables MSOs to couple value-add services (VAS), such as security and unified communications, to provide "multiple service, single access" connectivity to replace the prevailing "multiple service, multiple access" deployments resulting in significant cost reduction for both enterprises and themselves. It also simplifies MSO connectivity offers to off-net enterprises and, for business sites passed by its cable footprint, it positions the MSO as a secondary carrier for providing underlay services. By placing many WAN routing functionalities in the MSO network, it provides a range of flexible deployment options.

WAN simplification: Relative to traditional WAN services, SD-WAN provides enterprise IT managers with significantly greater network visibility and control to help them translate their application-level policies into effective routing decisions. In addition to enabling MSOs to offer more granular services with performance guarantees tailored to the specific needs of individual sites, it also provides enterprises and MSOs with a better understanding of the drivers behind their traffic growth to help contain

networking cost increases. For many enterprises with unmanaged WAN, this can be the tipping point for adopting MSO-managed services.

Separating the data, control, and management planes allows SD-WAN to distribute complex WAN functionalities between the enterprise site, its data center and the MSO. The complicated WAN path computation and routing functions in traditional on-premise solutions are substituted by relatively simple data forwarding engines that are located on site under the control of policy-driven SDN controllers placed in the MSO network to facilitate easy integration of diverse site connectivity and performance requirements. Greater automation also leads to simplified deployment, flexible service provisioning and assurance, and enhanced solution scalability in terms of enterprise sites and service SLAs.

Improved WAN economics: Distribution of WAN functionality results in thin on-premise equipment that leads to lower MSO capital investment per site. SD-WAN deployment costs are further reduced through automated provisioning that minimizes costly truck rolls and labor-intensive integration steps. Automation also reduces operating expenses incurred for routine care, fulfillment and assurance operations – such as software upgrades, fault management, inventory management, and performance monitoring – as well as longer-term capacity management operations. SD-WAN also enables significant transport cost savings by rationalizing site connectivity in several ways – such as augmenting MPLS connectivity with internet for channeling low-value traffic, providing internet breakout for cloud connectivity, and by moving low-priority sites to internet-only connectivity (and potentially dropping MPLS altogether).

Improved security: SD-WAN draws upon an MSO’s virtualized platform functionalities to provide programmable, policy-driven services that can be scaled to meet customer-, site- and application-specific security needs. Segmentation of enterprise DCs, branches and cloud sites into zone-based security policy groups helps reduce attack surfaces and work across heterogeneous resources such as virtual machines (VMs), containers, and physical appliances to protect against untrusted sources. Vulnerability mitigation measures, such as microsegmentation, help address threats arising from east-west traffic. Analytics-driven automation enables SD-WAN to detect many security threats ahead of their actual realization and take effective remediation measures.

Improved agility: Greater automation helps SD-WAN reduce the time to market new services and to upgrade existing services. It also eliminates expensive and time-consuming manual steps in the service assurance process. Use of self-service portals eliminates the complexity of product catalogs and improves customer experience while simultaneously reducing MSOs’ new order provisioning time and cost.

3. SD-WAN revenue potential

We adopt a two-step approach to estimating SD-WAN revenues for a typical MSO. First, we develop a 5-year SD-WAN connectivity revenue projection for the MSO’s target market. Next, we build the MSO’s share of overall market connectivity revenue and combine it with estimated SD-WAN enabled value-add services revenue to determine overall 5-year SD-WAN revenue.

3.1. SD-WAN connectivity market assessment

The 5-year SD-WAN connectivity market is built bottom up by considering the existing number of enterprise sites and their growth, current WAN services at these sites, projected SD-WAN adoption rates and estimated SD-WAN connectivity prices. Figure 1 illustrates the individual steps of this assessment framework that are now described.

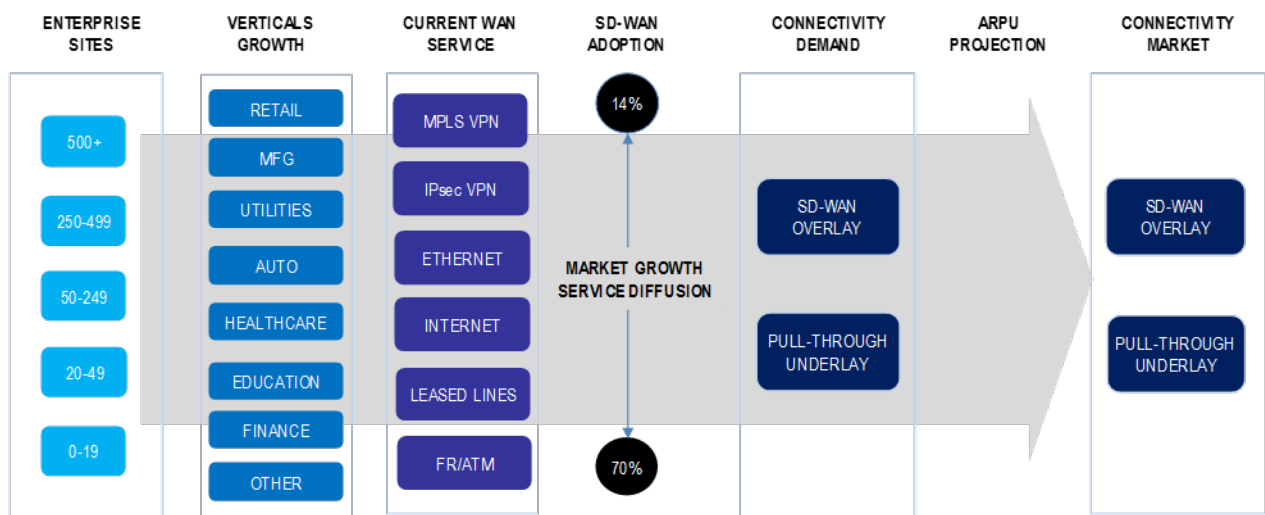


Figure 1 - SD-WAN Connectivity Projection

Number of enterprise sites

First, we estimate the current population of enterprise sites in the MSO's target market, broken down by their sizes – micro (0-19 employees), small (20-49), medium (50-249), large (250-499) and very large (500+) – and their verticals.

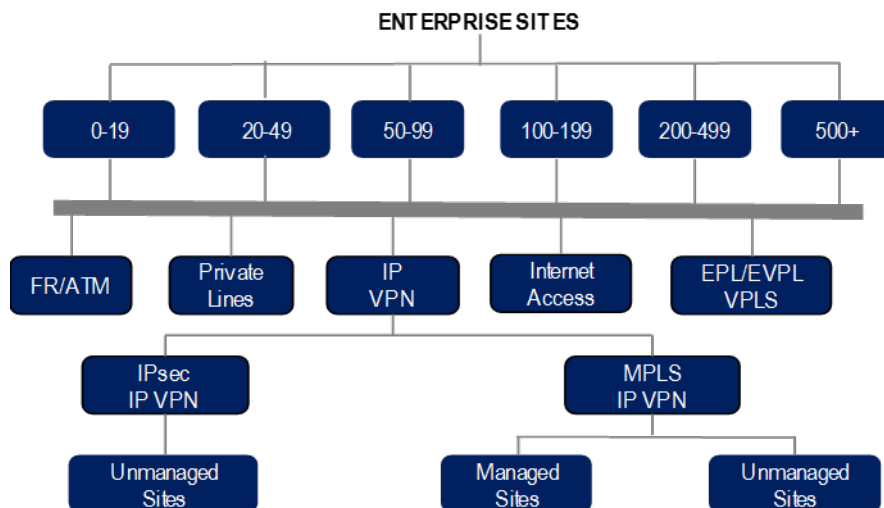


Figure 2 - Enterprise Site Map by Existing WAN Service

Growth of verticals

Next, we forecast future enterprise site count and bandwidth needs by projecting the current site population along the expected growth (or decline) trajectory of these verticals over the next 4 years. This step yields a 5-year map of enterprise sites by their sizes and the verticals they belong to.

Existing WAN services

At this step, we project the deployment of current WAN services across the various enterprise sites. These services include Frame Relay, ATM, private lines, internet access, ethernet service – including VPLS, EPL and EVPL – and IP VPN (an MSO may not offer some of these services). As shown in Figure 2, we break down IP VPN sites into MPLS- and IPsec-based sites, and furthermore, MPLS IP VPN sites into managed and unmanaged sites. We also develop the future deployment of these services at various enterprise sites by incorporating their growth trends shown in Figure 3.




Figure 3 - Growth of Existing WAN Service Connections

SD-WAN adoption and connectivity demand

The key premise of the market assessment model is that an enterprise's SD-WAN adoption is predicated on the incremental utility, or the *value*, that it provides relative to the enterprise's current WAN service. We measure the value of a WAN service along the following 8 dimensions:

1. **Service cost:** Total cost of WAN operations to the enterprise with a given WAN service. It comprises the cost of the WAN service, such as the monthly managed service fee charged by a service provider, and internal expenses incurred for providing network connectivity.
2. **Service performance:** Key performance capabilities of the WAN service. It includes the range and quality of supported SLAs, service reliability, and its ability to support service chaining and secure connectivity to public and 3rd party clouds.
3. **Implementation complexity:** Ease and speed of deploying the WAN service. It is a measure of the enterprise's ability to rapidly set up new site connectivity, terminate temporary connections, and add new services.
4. **Control:** Ease and speed with which the WAN service can translate enterprise policies into WAN routing decisions. This includes, for example, the ease with which the enterprise can reconfigure service priorities across individual applications, change traffic routing decisions at individual

sites, etc. This dimension reflects a key reason why many enterprises remain unmanaged as traditional managed services do not allow them to exercise the desired level of control.

5. Implementing value-add services: Ease and economics of deploying value-add services. It measures the ability of the WAN service to efficiently support enterprise applications relating to network security, unified communications, and other enterprise-specific needs.
6. Scalability: Ability to scale both size and scope of services supported at each site. This includes the ease with which the WAN service can scale bandwidth allocated to each site, connect a site to public and 3rd party clouds, and move applications across private, public and 3rd party clouds.
7. Network management support: The WAN service's ability to provide required levels of performance monitoring and reporting, and network assurance and fulfillment support.
8. Security: Overall protection provided by the WAN service, determined by its ability to prevent, detect and mitigate threats, especially in the context of increasing interaction with external clouds, more dynamic fulfillment and provisioning, and more frequent life cycle updates.

Each dimension is broken down into *attributes* to enable more granular characterization. Each attribute is assigned a *weight* that reflects its importance for a given site deployment. For example, a site currently supported by managed MPLS service has higher weights for service performance and network management support, and relatively lower weights for service cost. We assign each WAN service a utility *score* for each attribute that captures the extent to which that attribute is reflected in that service. For example, internet over broadband will score high on cost economics but relatively low on performance and security. The assignment of attribute weights and WAN service-specific attribute scores reflects WAN managers' preferences gleaned through various consulting engagements and analyst reports.

With this construct, the model determines the value of each WAN service for a given site as the weighted sum of its attribute scores. The probability with which a given site will migrate to SD-WAN is proportional to its relative value over the current WAN service. SD-WAN's overall adoption rate for the market is the sum of these probabilities over all enterprise sites in the target market. Note that this adoption rate indicates SD-WAN's share of the WAN services market in the eventual, steady state. We model its market share in the interim years through a diffusion process using an approach proposed by Bass (see [1] for example) to determine the number of SD-WAN connections created during each year of the 5-year planning horizon. Figure 4 shows how the diffusion rate can impact the number of connections during the interim years.

In order to determine the net new connectivity market created through SD-WAN adoption, we also compute the loss of connections due to the cannibalization of existing WAN services, and incremental connectivity gains resulting from pull-through underlay services. [Note that SD-WAN adoption does not always result in the displacement of current WAN service; it is deployed in many cases, such as Hybrid SD-WAN, also to augment existing services.] The first stream captures connections lost to service providers as their existing customers churn from their current WAN services to SD-WAN. The second stream addresses the new underlay internet and MPLS connections created to support the adopted overlay SD-WAN services. For example, an enterprise site switching from current Frame Relay service to SD-WAN over MPLS triggers new revenues on account of both overlay SD-WAN and underlay MPLS services; there is also a simultaneous loss of cannibalized Frame Relay revenue.




Figure 4 - Sensitivity of Connections Market to Speed of Diffusion

At the end of this step, we have a map of new overlay SD-WAN connections and new underlay internet and MPLS connections, as well as a count of cannibalized existing WAN service connections.

Average Revenue per Service (ARPS) Projection

In order to monetize the projected number of SD-WAN connections, we first estimate the average revenue per service (ARPS) for individual SD-WAN connectivity services, and next determine the revenue generated annually at each enterprise site based upon the type of service provided to that site.

We determine the basic connectivity ARPS for speeds ranging from 30 Mb/s to 1 Gb/s. It is estimated through simultaneous consideration of demand- and supply-side economics, and market prices. From customers' perspective, ARPS must reflect their willingness to pay for the incremental utility received over their current service. From MSOs' perspective, it should provide the targeted operating margins. From the market's perspective, ARPS should reflect the prevailing prices of services currently available.

The pricing model captures these 3 considerations; it reflects the market prices for currently available services (primarily at lower connectivity speeds), while ARPS at higher speeds is computed using a utility-based approach similar to one described earlier for determining SD-WAN adoption rates; this approach ensures that there is an increase in net utility for customers switching to SD-WAN. The model also verifies that MSO operating margin requirements are met; it is further validated through SD-WAN business case evaluation discussed in Section 5.

Overall Connectivity Market Revenue

The final step for arriving at the SD-WAN enabled connectivity revenue requires combining ARPS values with connection projections and netting out cannibalization losses. We include additional revenues generated through options, such as redundant CPE devices and high availability packages, to round up the overlay connectivity market projection.

To determine the overall SD-WAN enabled connectivity market, we also include revenues generated through additional pull-through underlay services – primarily internet and MPLS – required to support SD-WAN for the migrating sites that do not currently have these services.

3.2. MSO revenue projection

An MSO's 5-year SD-WAN revenue comprises its share of the SD-WAN connectivity market for both overlay and underlay services and pull-through revenues generated by value-add services enabled by SD-WAN. We focus on determining the *net new* revenues generated by SD-WAN after considering all revenue streams impacted by SD-WAN adoption.

Connectivity Revenue

We first consider the MSO's current enterprise customers. As described in Section 4.1, we estimate the number of current sites, project their 5-year growth in terms of their count, bandwidth requirements, and WAN services. Combining these connections with estimated services ARPS yields a baseline MSO revenue map prior to SD-WAN adoption.

Using the incremental utility-based evaluation and Bass diffusion approach discussed earlier, we next develop SD-WAN adoption rates to project the number of SD-WAN connections added each year across these sites. We estimate ARPS for various SD-WAN services using the 3-pronged approach discussed in Section 4.1 to develop 5-year SD-WAN connectivity gross revenue projection.

In order to determine SD-WAN's net new revenue impact, we additionally consider the following adjacent revenue streams that are triggered by SD-WAN adoption:

1. Revenue cannibalization: For many customers, SD-WAN will substitute the existing WAN service, resulting in the cannibalization (and loss) of current revenues. [Note that SD-WAN can be viewed as protecting the revenue streams with these customers, as otherwise they would be lost to other providers.]
2. Hybrid SD-WAN: However, in some cases – especially with MPLS services – SD-WAN will be adopted to augment current services leading to hybrid SD-WAN deployments with no cannibalization. For example, Verizon notes that SD-WAN has improved its MPLS sales [2].
3. Customer churn: Some migrating customers may switch to other SD-WAN providers, leading to loss of current underlay service revenue.
4. Pull-through underlay revenue: Some SD-WAN customers will require new internet or MPLS underlay connections, as well as LTE or other backup lines.
5. Cloud connect: Some customers will also opt for dynamic and secure connectivity to 3rd party and public clouds.
6. Options and additional services: Many SD-WAN customers will buy additional options such as dual CPEs, and professional services involving configuration support, etc.

In addition to generating new revenues from existing customers, SD-WAN also helps the MSO expand its current customer base. There are two modes through which new customers can adopt SD-WAN with the MSO: they can either churn into SD-WAN from their existing services with other providers, or they can adopt SD-WAN overlay from the MSO while retaining their current providers for underlay services. Enterprises adopting the latter path provide the opportunity for the MSO to rapidly build its market share by reaching out to off-net customers outside its cable footprint without having to make expensive investments.

To compute the revenue impact of new customers, we classify them into two groups – those covered by the MSO’s cable footprint and others who are not. The revenue potential of the first group is computed much the same way as the existing customers with the obvious recognition that the enterprise sites under consideration cover all locations in the MSO’s target market that are not its current customers. Other than cannibalization and customer churn, the other adjacent revenue streams apply here as well. For the second group, we primarily consider the revenue impact of overlay SD-WAN connection and a subset of options and additional services as other adjacent flows do not apply.

Figure 5 shows a typical enterprise customer map for an MSO that captures several SD-WAN connectivity variants discussed above.

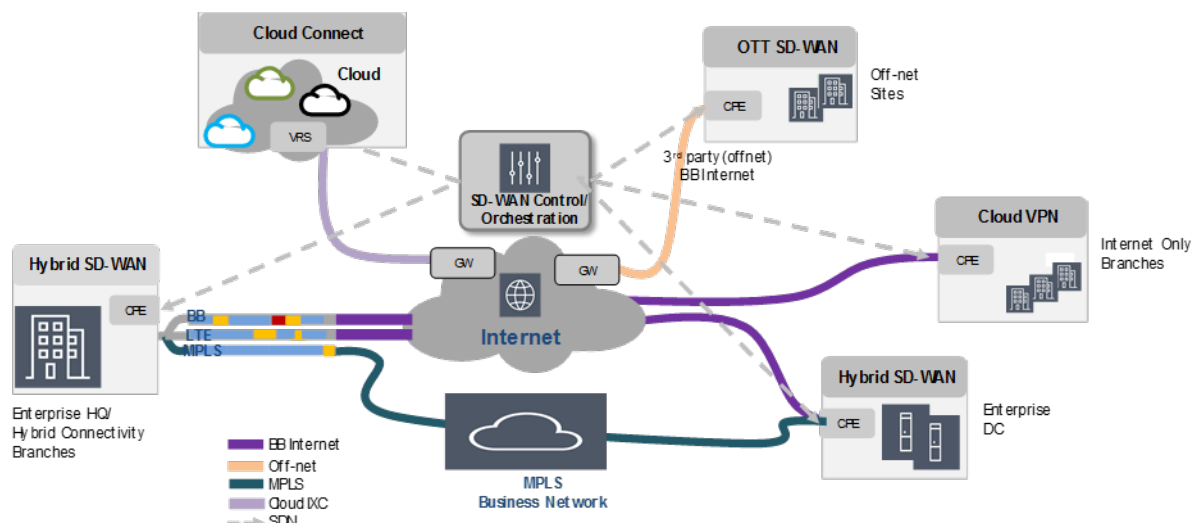


Figure 5 - SD-WAN Connectivity Map

Value-add Service Revenue

In addition to connectivity, SD-WAN also enables new value-add services. We include two groups of these service in the revenue model - security services and unified communications – that are closely related to network connectivity. However, depending upon the MSO’s enterprise strategy, the scope of these value-add services can extend far beyond to span other enterprise applications and business processes, and result in revenues exceeding those generated through basic connectivity.

As noted earlier, network security services are a natural adjunct to SD-WAN connectivity. We include basic services – such as anti-virus, IPS, DDoS mitigation, web filtering, antispam, and mobile security – in the model. Similarly, we include voice, email, messaging, collaboration sites, web and video conferencing services as part of the unified communications package. Adoption rates of value-add services are determined through the utility-based approach discussed earlier, and they are evaluated at prevailing market prices.

Figure 6 presents the 5-year MSO revenue broken down into SD-WAN overlay, pull-through underlay and value-add services. Overlay services are clearly dominant, accounting for 68% of the overall \$135 million revenue over 5 years. However, value-add services also contribute significantly by providing 20% of the total revenue; as noted earlier, the share of value-add services revenue can be more substantive if the MSO is willing to offer a broader range of these services as part of its enterprise strategy. Also worth

noting is the 12% revenue uplift generated by pull-through underlay services that can play a significant role in determining the viability of the overall business case.




Figure 6 - MSO Revenue Projection

3.3. Revenue sensitivity

It is important to understand the sensitivity of SD-WAN revenue projection to market and business parameters. Figure 7 shows the impact of two such parameters – the pace of SD-WAN market growth and the MSO’s realized share of this market. Gradual market growth relates to the steady state penetration being reached in 7 years; balanced and aggressive growths reduce this time interval to 5 and 3 years, respectively. Likewise, we bound MSO market shares at Reach (30%) and Pessimistic (20%) levels around the target share of 25%. The nominal 5-year revenue of \$135 million, following the growth shown in Figure 6, is projected for a balanced market growth and target MSO market share of 25%.

As shown in the table, MSO revenue can vary considerably across the scenarios. Specifically, if the market grows aggressively, it presents significant upside potential for the MSO. An MSO planning for the target market share must nonetheless be ready to avail of any sales opportunity as even a 5% increase in market share translates into \$36 million to \$61 million additional revenues over 5 years depending upon the pace of market growth. On the other hand, a gradual market growth does not significantly diminish overall revenues; this is a strong plus for SD-WAN. For example, at target market share, the MSO’s 5-year revenues drop only by 11% (from \$135 million to \$120 million) if the market grows gradually. On the other hand, with aggressive market growth, these revenues increase by 36% (from \$135 million to \$184 million).




Figure 7 - Sensitivity of MSO Revenue Projection

4. SD-WAN Business Case Evaluation

The 5-year SD-WAN business case for the MSO is built upon the revenue projection described above through a bottom-up computation of incremental capital expenses (capex), and network and non-network operating expenses (opex) required to support this revenue.

Figure 8 gives the key revenue, capex and opex components considered in the business case evaluation. Capex is driven by the investment in SD-WAN solution elements shown in Figure 9 that comprises both hardware, such as CPE devices and self-service portals, and software required for the orchestrator, SDN controller, virtual network functions (VNFs) and other components of the solution stack. Capex also includes incremental capital expenses incurred for common infrastructure, such as the DC network, that SD-WAN shares with other network services. Configuring and sizing equipment hardware, as well as establishing software license requirements, is driven by parameters such as the number of subscribers, required features and options, and processing capacity demand. Market prices are used to estimate the cost of procuring SD-WAN elements and the common infrastructure. We also build in the cost of overall system integration (including interworking with existing OSS/BSS systems) into capex computation.

SD-WAN opex includes network expenses to support hardware and software maintenance, network operations such as fulfillment, assurance and capacity planning, and for meeting the network power and real estate needs. It includes right-to-use fees for 3rd party software; a large share of overall opex also comes from non-network expenses on account of sales & marketing, general & administrative items, and customer care and billing services.

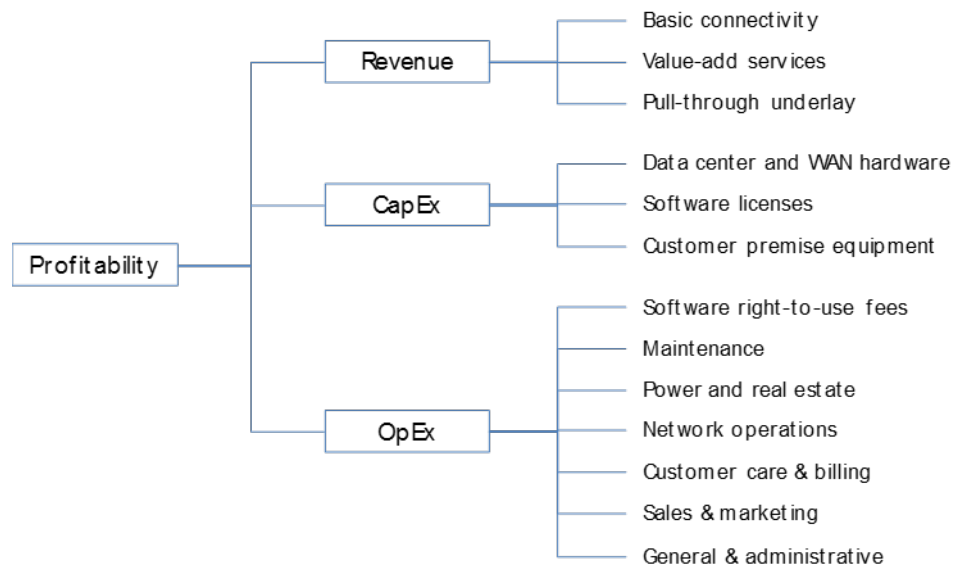


Figure 8 - Business Case Components

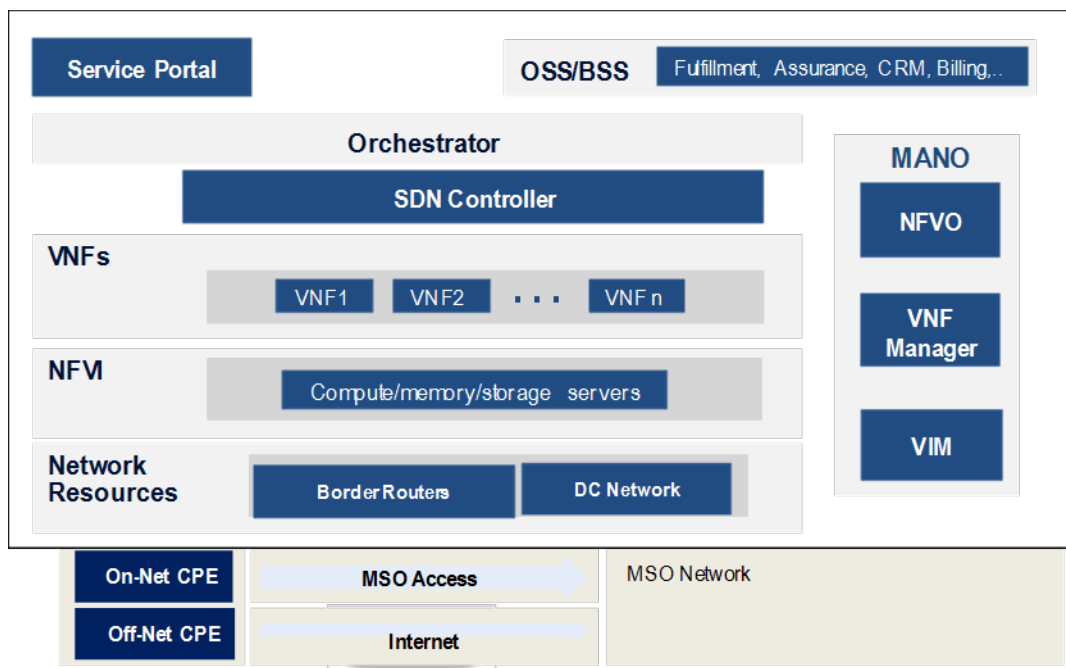


Figure 9 - SD-WAN Solution Components

Figure 10 captures the individual elements of the business case financials, and also shows the resulting profitability in terms of the cumulative discounted cash flow (CDCF) through each of the 5 years. These values indicate a strong SD-WAN business case for the MSO with a 5-year net present value (NPV) of \$36 million and a payback period of 2.25 years.

SD-WAN capex is driven primarily by the cost of the CPE device and the self-service portal, both of which scale with actual deployment. A key implication of this cost structure is that it significantly reduces the inherent financial risk of SD-WAN deployment in that major capital investment is backed by immediate customer order and revenue inflow. Other capex elements, such as the virtualization platform, account for about 32% of overall 5-year capex with the bulk of these investments coming in the later years.

Major drivers of SD-WAN opex are right-to-use fees for 3rd party value-add services software, sales & marketing expenses, general & administrative costs, SD-WAN software maintenance fees, and customer care expenses.

In addition to establishing baseline SD-WAN financials, it is also necessary to understand their sensitivity to key market and business parameters. It is difficult to precisely predict how the market will behave in future, what the actual SD-WAN take rate will be, and how competition will respond. It is also important for the MSO to understand the implications of its own actions, such as market entry timing. In order to capture some of these sensitivities, we simulate the results of 3 scenarios and contrast them with the baseline financials discussed above. The scenarios investigated are:




Figure 10 - SD-WAN Business Case Financials

Scenario 1 – MSO late to a rapidly growing market by 1 year. MSO’s entry is delayed by 1 year into a market that is growing more rapidly than one considered in the baseline.

Scenario 2 – Slow market adoption of SD-WAN. While SD-WAN eventually achieves the baseline market penetration, there is a 2-year delay due to slower diffusion during the interim years.

Scenario 3 – Phased VAS introduction. MSO strategy to deploy only connectivity services at a site initially, and delay VAS deployment at that site by 1 year.




Figure 11 - SD-WAN Business Case NPV Sensitivity

Figure 11 captures the NPV implications of these 3 scenarios. Late entry in a growth market under Scenario 1 has the most severe impact on the MSO's business case. Not only is its NPV reduced significantly throughout the 5-year planning horizon, the NPV growth trajectory in the later years (and beyond) is seriously compromised as well because the MSO's late market entry reduces its market share and ARPS, and increases customer churn, in the face of early competitive presence in a growth market.

While a slowly developing SD-WAN market under Scenario 2 reduces 5-year NPV by 50% and increases the payback period, the MSO may still find the business case to be viable, especially as SD-WAN deployment risk, and the risk of stranded investments, is inherently low for the reasons discussed earlier.

A phased VAS deployment approach under Scenario 1 has minimal impact during early years, and the baseline payback period is generally preserved. It also staggers initial investment costs with a relatively small impact on SD-WAN connectivity market share. However, the delayed realization of higher-margin VAS services results in lower NPV over the 5-year horizon. A viable catch-up strategy for the MSO under this scenario would be to pursue VAS deployment more aggressively in the later years.

Conclusion

In this paper, we present a utility-based approach to developing SD-WAN revenue and profit projections for a generic MSO. Market and revenue projections are based on SD-WAN adoption rates driven by its incremental value measured in terms of WAN service attributes desired by individual enterprise sites. We show that the SD-WAN market is significant and growing, and its adoption provides the MSO a viable business case with a payback period of 2.25 years. We also note that SD-WAN has inherently low financial risk as the bulk of the required capital investment is based on actual deployment, and therefore, it scales naturally with revenue.

While SD-WAN financials based on expected market growth remain attractive in of themselves, the MSO needs to also be ready for demands exceeding these projections in order to maximize its revenue opportunity. Increasing virtualization, growing enterprise adoption of cloud services, and need to simplify WAN operations and reduce costs in the face of expanding branch office connectivity and application needs are trends that portend significant potential for faster SD-WAN growth.

A word on SD-WAN technology maturity. Backed by a strong ecosystem, SD-WAN has reached a level of performance stability and maturity that makes it ready for mass deployment. It provides a range of solution alternatives that makes it easily deployable in a phased manner. As it expands its capabilities through increasing automation and analytics support, SD-WAN will enable an MSO to further distance its WAN solutions from traditional alternatives.

Abbreviations

ARPS	average revenue per service
CDCF	cumulative discounted cash flow
CPE	customer premise equipment
DC	data center
DDoS	distributed denial of service
EPL	ethernet private line
EVPL	ethernet virtual private line
MPLS	multi-protocol label switching
MSO	multiple system operator
NFV	network function virtualization
NPV	net present value
OSS/BSS	operational support system/business support system
SMBs	small and medium businesses
SDN	software defined networking
SD-WAN	software defined-wide area networking
SLA	service level agreement
VAS	value-add service
VM	virtual machine
VNF	virtual network function
VPLS	virtual private line service
VPN	virtual private network
WAN	wide area networking

Bibliography & References

1. V. Mahajan, E. Muller, F. M. Bass: New Product Diffusion Models in Marketing: A Review and Directions for Research, *Diffusion of Technologies and Social Behavior*, IIASA, Austria, 1991
2. Verizon: SD-WAN Boosting MPLS Sales, Light Reading, June 14, 2017, <http://www.lightreading.com/carrier-sdn/sd-wan/verizon-sd-wan-boosting-mpls-sales/v/d-id/733682>

How to Finally Conquer Wi-Fi in the Home: Service Provider Style

A Technical Paper prepared for SCTE•ISBE by

Nav Kannan

VP Product Management
ARRIS International plc
3871 Lakefield Drive, Suwanee GA 30030
+1 215 435 2231
nav.kannan@arris.com

Charles Cheevers

CTO CPE Products
ARRIS International plc
3871 Lakefield Drive, Suwanee GA 30030
+1 678 473 8507
charles.cheevers@arris.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	5
The Evolving Wi-Fi Home	10
1. Evolution of the Home Wi-Fi Network	10
2. Growing diversity of Wi-Fi devices in a residential setting.....	12
3. Customer Expectations and the Burden of Cable Service Providers	16
Unwrapping the Goodies.....	19
3.1. Choose Your primary AP	20
3.2. Choose Your Multi-AP Topology Strategy	20
3.3. Choose Your Multi-AP Device Architecture	20
3.4. Choose Your Meshing Solution.....	21
3.5. Choose Your Steering Solution.....	21
3.6. Choose Your Cloud RRM/SON solution	21
3.7. Choose Your Wi-Fi Telemetry Strategy	21
3.8. Add Your Machine Learning and AI Roadmap	22
4. Wi-Fi improvements Worth Noting	22
5. In-Home Device Implementations Catch-up to Standards.....	22
5.1. 802.11ax: Aimed at Improving Efficiency and Performance	22
5.2. Wi-Fi EasyMesh	24
5.3. 802.11v and 802.11k to help “Roam (in) Sweet Home”.....	27
5.4. Wi-Fi Easy Connect	29
6. AI / ML techniques.....	30
7. Cloud Based Analytics	37
8. Goodie-Bag Summary.....	37
Addressing Consumer Needs with the Goodies	37
9. Connectivity & Security	37
10. Coverage & Performance at Range	39
10.1. Multi-AP for Network Extender Communication.....	41
10.2. Zero Touch Provisioning of New Wi-Fi Extenders	41
10.3. Cloud Assisted and Policy-based Management	41
10.4. Roaming and AP Steering.....	42
10.5. Band Steering.....	42
10.6. BSS Steering.....	42
10.7. Dynamic Channel Selection	43
10.8. Mesh Link Optimization	43
10.9. Airtime Management	43
11. Subscriber Visibility & Device Management	43
Conclusion.....	45
Abbreviations	45
Bibliography & References.....	46

List of Figures

Title	Page Number
Figure 1 - Customer Expectations of Wi-Fi Speeds.....	5
Figure 2 – Mesh Network with ‘Hub and Spoke’ as well as ‘Daisy Chain’ Topologies	8
Figure 3 - Simplicity of QAM Video	9
Figure 4 - Complexity of IP Video Making Troubleshooting Complicated	9
Figure 5 - The Home Network is Set to Evolve.....	10
Figure 6 - High Capacity Device Increase	11
Figure 7 - Service Bandwidth Increase	11
Figure 8 - Smart Home / IoT Device Increase	12
Figure 9 - The Four Main Categories of Home Service	13
Figure 10 - Static and Mobile Device Groups	14
Figure 11 - Device Groups Based on Location of the Device	15
Figure 12 - Diversity of Wi-Fi Client Devices	16
Figure 13 - Expectation: Wi-Fi Speeds = Access Speeds	17
Figure 14 - Reality: Attenuation and Interference Affect Performance	18
Figure 15 - Typical Wi-Fi Device Distribution in a Home	20
Figure 16 - Steering and Meshing: Defined	21
Figure 17 - New Technologies Demanding drive for Wireless Performance.....	23
Figure 18 – Key Benefits of 802.11AX.....	23
Figure 19 - 802.11ax Performance Gains over 802.11ac.....	24
Figure 20 - Wi-Fi Certified EasyMesh™ (source: http://www.wi-fi.org).....	25
Figure 21 - Wi-Fi EasyMesh Architecture.....	26
Figure 22 - Role of WFA Standards Based Components	26
Figure 23 - Criticality of AP Steering.....	27
Figure 24 - Signal Strength Alone is not a Unique Determinant of Performance	28
Figure 26 - Role of AP Utilization in Steering Decisions.....	29
Figure 27 - Device Provisioning with Wi-Fi Easy Connect.....	30
Figure 28 - Consumer Behavior	31
Figure 29 - Policy to Adopt for High-Bandwidth Set-top Box.....	31
Figure 30 - Policy for Static High Definition TV	32
Figure 31 - Policy for High-Bandwidth Mobile Devices.....	33
Figure 32 - Visual Representation of Device Coverage Health	34
Figure 33 - Policy for Static High-Priority Elements	34
Figure 34 - Policy for IoT devices	35
Figure 35 - Policy for Devices at the Edge of the Home.....	35
Figure 36 - Policy for Devices Outside the Home.....	36
Figure 37 - Wi-Fi Client Onboarding Process	38
Figure 38 - Device Fingerprinting Essential with the Explosion of IoT devices	39
Figure 39 - Correlation of Poor Performance to Client Characteristics	39
Figure 40 - Architecture of a Typical Cloud-Assisted Wi-Fi Management Solution	40
Figure 41 - Gateway Controller for Wi-Fi Management.....	40
Figure 42 - Zero-Touch Configuration.....	41

Figure 43 - An Example Implementation of Wi-Fi Management Function	42
Figure 44 - Mobile Application Screens to Highlight Data Usage	44
Figure 45 - Mobile Application Screens for Parental Control and Device Access Restrictions	44

List of Tables

Title	Page Number
Table 1: Cost Delta with Tri-band Configurations	19
Table 2: Categories of Data Analysis problems that Ai/ML Techniques Excel at ?	36

Introduction

Service delivery over IP is now a reality, and in a home, most of these services are consumed with a wireless device. In addition, many of these services are overlays from Over-the-Top (OTT) content providers. Service providers have to adapt to this changing world by being able to manage the optimal delivery of these services to the devices by offering “Wi-Fi” itself as a service, or they risk being relegated to the role of just the access provider, getting the bits in and out of the home at the point of termination, but having no control or even role in the distribution of the data inside the home. Wireless technology is inherently complicated and is also evolving at a more rapid pace than the access technologies that current cadence of device replacement for service providers. Retail devices are also now entering the market that are designed to cater for ease of use, coverage and rely typically on Multi Access Point S/W management and Mobile App to give the user features to manage connected devices. These devices offer consumers the promise of ease of set-up, manageability, and visibility, usually with the assistance of a mobile application. This is doubly painful for the service providers, in that they don’t get to participate in the commercial transaction of these retail devices, while sharing an unfair burden of having to answer for the “poor Internet service” should any of these retail devices not stand up to the promise that they were offered. An additional downside to these retail devices is that they will also inhibit the service provider to roll out new IP services that rely on being able to manage and touch all end devices. Retail APs often provide their own DHCP scope addresses and NAT out the service provider.

The challenges presented above may seem to portray any solution process as a daunting task. However, there is a way that the service provider can compete and in particular adding tools to their arsenal to take advantage of and take control of the Wi-Fi in the home. To start with, let us look at the key performance requirements from the perspective of a subscriber:

Connectivity: ‘I need to be able to seamlessly onboard and connect a variety of Wi-Fi devices’

Coverage: ‘I need to be able to make use of my Wi-Fi devices throughout my home’

Performance: ‘These Wi-Fi devices must adequately perform to meet my needs’

Happiness NPS “If the applications I use on Wireless devices work well – I’m a happy customer”



Figure 1 - Customer Expectations of Wi-Fi Speeds

In addition, customers who are now used to the new generation of mobile applications and retail devices want the ease and comfort of simple and intuitive ways to manage and control the various devices, while also wanting the reassurance that there will not be any compromise of either privacy or security. These expectations are hygiene factors and are listed below.

Security: ‘I need the Wi-Fi devices to not compromise data security’

Manageability: ‘I need to be able to control and manage these devices’

Visibility: ‘I need to be able to visualize and monitor these devices to ensure they are working properly’

All of this while understanding that Wi-Fi is like a “Utility” and it really should just work. For a consumer to pick up a smartphone and use a mobile app to manage Wi-Fi it needs to be:

- Useful: Services like ‘Mealtime mode’ where all devices are paused on Wi-Fi
- Infrequent use but intuitive to use: Onboarding new Extenders or devices; Solving simple Wi-Fi issues
- Worthwhile: Adding security services to Wireless connectivity

Practically all the above needs can now be adequately addressed by the Service provider. The Wi-Fi standards are rapidly evolving not just from the perspective of functionality, but also from the perspective of manageability and interoperability. New Wi-Fi goodies have arrived and it’s time to unwrap them and take control of Wi-Fi. We now have what we need to make the Service Provider Access Points (APs) control the Wi-Fi clients in the home. Service provider class APs can now support 802.11k, v, r, u, and ‘ai’. These IEEE letters allow the AP to be able to command the Wi-Fi clients to be able to move to any commanded AP, provide information on hidden AP, and better report their telemetry. Add the introduction of 802.11ax, and we also have even better control of client battery life.

It is worth pausing here and reflecting on a couple of key points that are the fundamental decision points for service providers:

1. The right Multi Access Point strategy
 - a. There is a huge desire to go to an all Wireless Mesh solution. Simplifies things for install. Makes performance more complex.
 - b. There is also the question of how to architect the Wi-Fi mesh. True mesh technologies like full 802.11s implementations where all Wi-Fi nodes can see and reach all Wi-Fi nodes tend to be over complex for home solutions.
 - c. The service provider path is probably aligned more with the regular path of being the care keeper for the home
 - i. A single AP for lowest cost acquisition
 - ii. Minimum number of devices for range
 - iii. Minimum number of devices for performance

This tends to converge then to more of a hub (main Gateway) and spoke – direct attach extenders. The first one which typically solves issues to 95% of US homes. Additional spokes can also be added for 360 coverage.

For larger homes the addition of another multi hop extender is also simpler than true mesh solution.

DESIRE FOR MESH ; RELIABILITY WITH WIRE

Mesh



Desirable – Simple User Install

Physics – 2x2 have one channel for backhaul and fronthaul

Physics – AP needs to have enough backhaul to Fronthaul

Mesh that works – works with Tri-band extender but **increased Cost**

Hub and Spoke



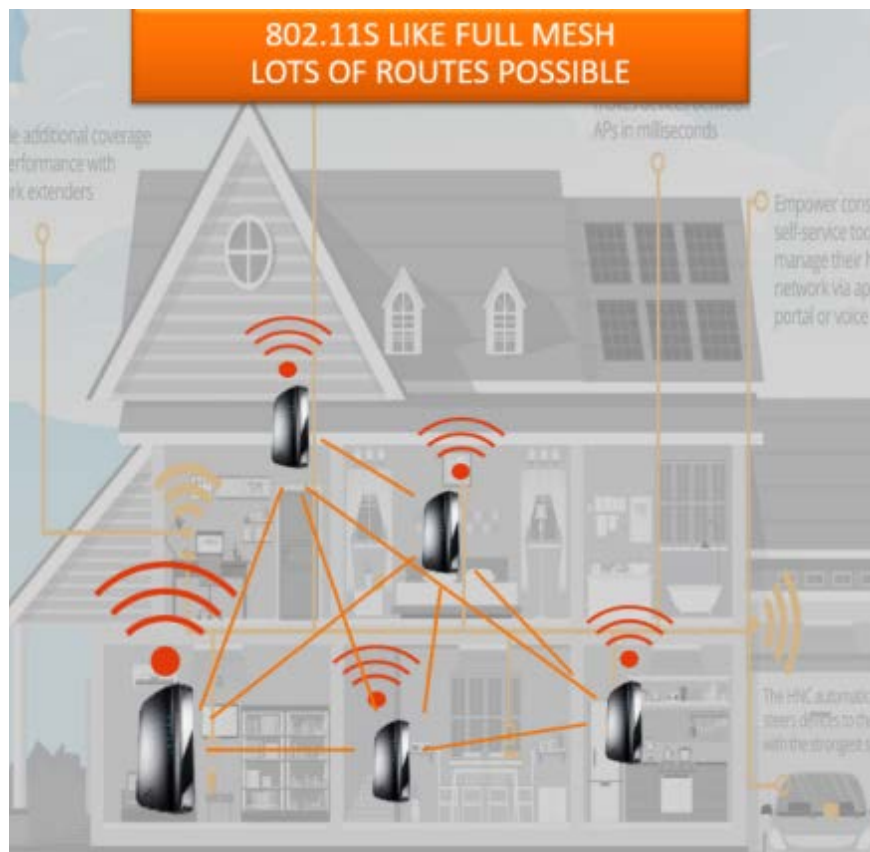
Daisy Chain



Flexible – Can be wire or Wi-Fi

Physics – Wired backhaul like Ethernet, MoCA or G.hn does not use Airtime

Physics – Satellite AP can be placed far from Root AP closer to where its needed without backhaul Wi-Fi restrictions






Figure 2 – Mesh Network with ‘Hub and Spoke’ as well as ‘Daisy Chain’ Topologies

2. The right set of tools to trouble shoot Multi-AP Wi-Fi. As you can see below transitioning from QAM video to IP video introduces complexity to troubleshoot. You can see below that QAM video has been simple to troubleshoot. TV is not working the problem is either into the STB or in the STB.



Figure 3 - Simplicity of QAM Video

If you add in a new Multi AP Wi-Fi architecture, now when the TV is not working – the broadband and in home Wi-Fi architecture are implicated with 6 points for failure to assess. It is therefore important for all Wi-Fi solutions to resolve packet loss to the multi hop or meshed architecture that is put in place.

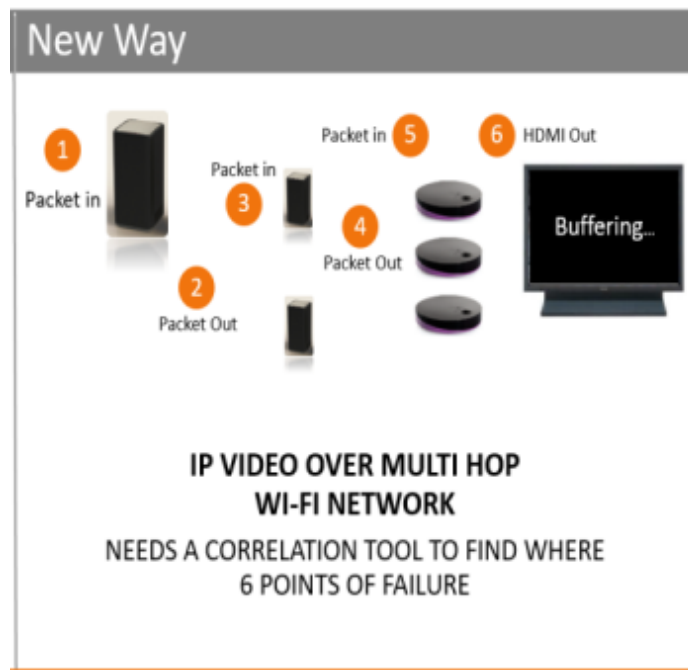


Figure 4 - Complexity of IP Video Making Troubleshooting Complicated

The Home Wi-Fi devices range from static high bandwidth devices like set-tops, to roaming devices, to critical Internet of Things (IoT) Wi-Fi devices. Using Data Mining and Machine Learning to create policies that maximize the performance of these devices through controlling Wi-Fi transmissions and client connections is now fully in the hands of the service provider back office solutions and gateway controllers. Add the recent standardization of the Wi-Fi Alliance (WFA) Multi Access Protocol to allow different vendor Wi-Fi AP devices to talk to each other in a common language and you have the perfect toolbox for service provider managed Wi-Fi services – just in time for the migration from Quadrature Amplitude Modulation (QAM) Video to Video over IP over Wi-Fi. This paper covers these topics in depth.

The Evolving Wi-Fi Home

1. Evolution of the Home Wi-Fi Network

The in-home Wi-Fi network and associated services are constantly evolving, and this presents both an opportunity and challenge for the service provider. The evolution is instigated by both the service providers and the consumers in parallel.

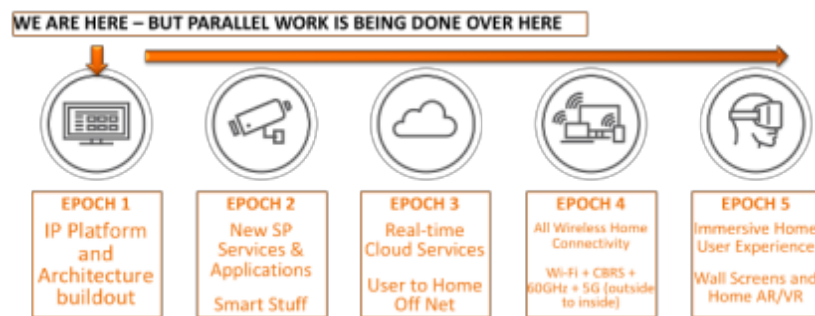


Figure 5 - The Home Network is Set to Evolve

The type of devices in the home are also diversifying and there is an increasing number of devices that have high capacity in terms of bandwidth consumption. A typical US household has 2.6 family members, and a distribution of the various device types is captured in the figure below:




Figure 6 - High Capacity Device Increase

Of particular interest is the consumption of video, a technology which continues to evolve independently, with 4K, UHD and even 8K over time. Virtual Reality / gaming consume high bandwidth video, and the following figure shows the bandwidth consumption patterns.

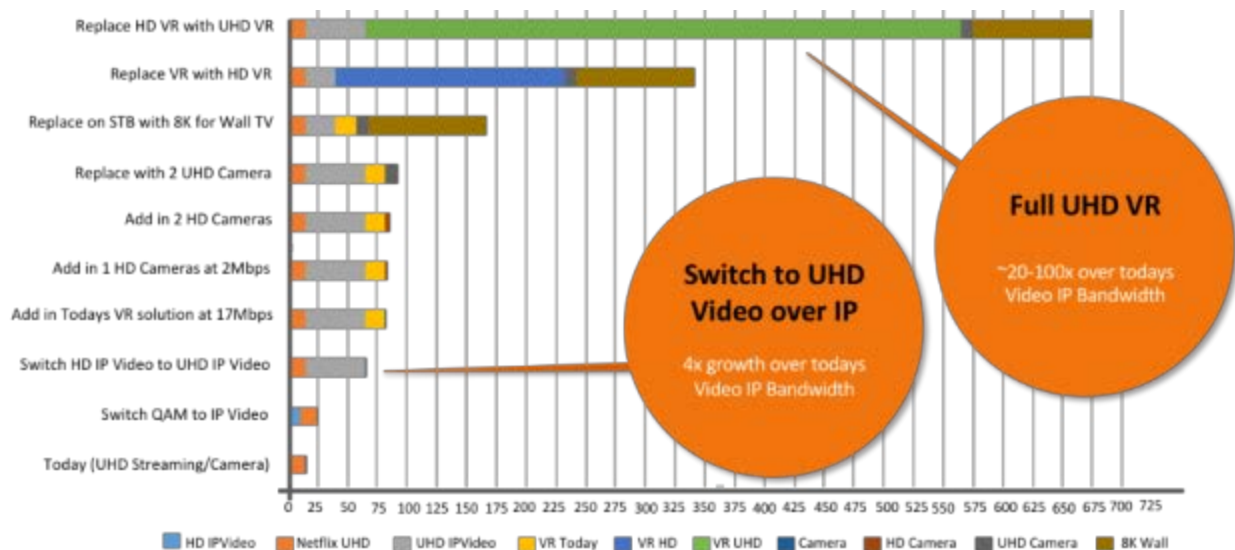


Figure 7 - Service Bandwidth Increase

While video and related services present the challenge of ever-increasing bandwidth requirements, the advent of Internet of Things (IoT) and associated services, there are now additional challenges. The evolving home with IoT is shown in the figure below.



Figure 8 - Smart Home / IoT Device Increase

2. Growing diversity of Wi-Fi devices in a residential setting

The ubiquitous adoption and the continuous improvement of the Wi-Fi standards has led to an increasing number of Wi-Fi enabled devices in the home. Consumers are now able to enjoy the practical benefits as well as the aesthetic appeal of not having to deal with the clutter of wires, irrespective of the nature of the use-case. Such is the appeal of the need to “get rid of the wires”, that we have a growing market for even wireless chargers for mobile phones.

The simplistic umbrella term ‘Wireless Device’ is also deceptive, since it hides the complexities that arise of the nature and use of the specific device. There is a growing diversity of these devices and consequently there is an ever-growing gap between the expectations of an end-user in terms of the functionality and performance of any ‘wireless device’, and the actual reality. If services that are provided by an entity such as a cable service provider, any issue due to a device will potentially taint the perception on the service itself. Since the trend of the subscriber is to shift as many devices to be wireless, it is inevitable that the onus of ensuring the wireless connectivity performance of all devices, whether provided by the service provider or bought directly by the subscriber will unfairly fall on the service provider.

The wide diversity of devices can be studied by grouping Wi-Fi devices into very simple categories.

Four main categories of Home Service for Wi-Fi are:

- **Video Players** – Highest bandwidth consumption and in the case of the SP’s own STB – the most likely to cause customer churn if it does not work properly
 - Additionally to add differentiated Video services like 4K HDR and 8K will require a robust Wi-Fi network above the level of OTT streaming solutions which rely on ABR
- **Broadband** – Emphasis changing to lower latency and Gbps burst to get the user a snappier experience as the lines blur between the benefits of 500 Mbps, 600 Mbps to 1 Gbps WAN SLA
- **Voice and Audio** – Small packet services that cannot have delay, crackles or drops – so high value and high customer churn if voice over Wi-Fi does not work properly
- **IoT** – Wi-Fi IoT security services like cameras become ‘must always work’ services. They can also drive hard demands on Wi-Fi if installed at range and outside home

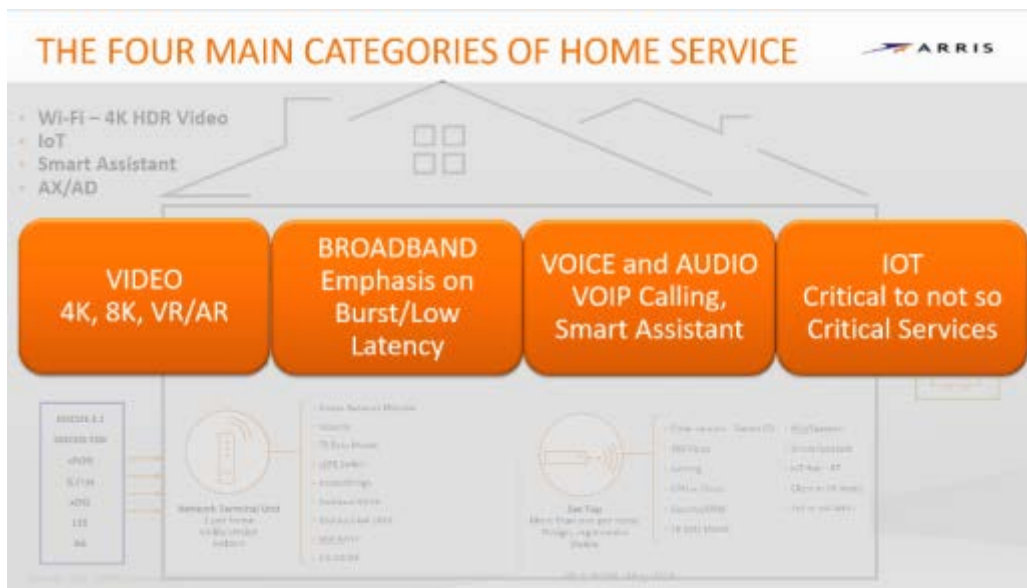


Figure 9 - The Four Main Categories of Home Service

There are also two basic groups of devices in the home as they relate to Wi-Fi:

- Static Devices – STB, Smart TV all the way to 4K Security cameras
- Mobile Devices – particularly smart phone tablets but also devices like Wi-Fi enabled vacuum cleaners

It is important for a service provider Wi-Fi solution to make this simple categorization of devices. In particular for example – STB need to be prioritized for quality video delivery. They are not expected to steer to other APs often and also may have specific policies to only work on 5 GHz bands. Additionally as AP utilization is a key metric of where clients attach in a Multi-AP home – the contribution to the AP utilization of a 4K STB makes it a client that forces other clients to steer off the AP when in high utilization levels.

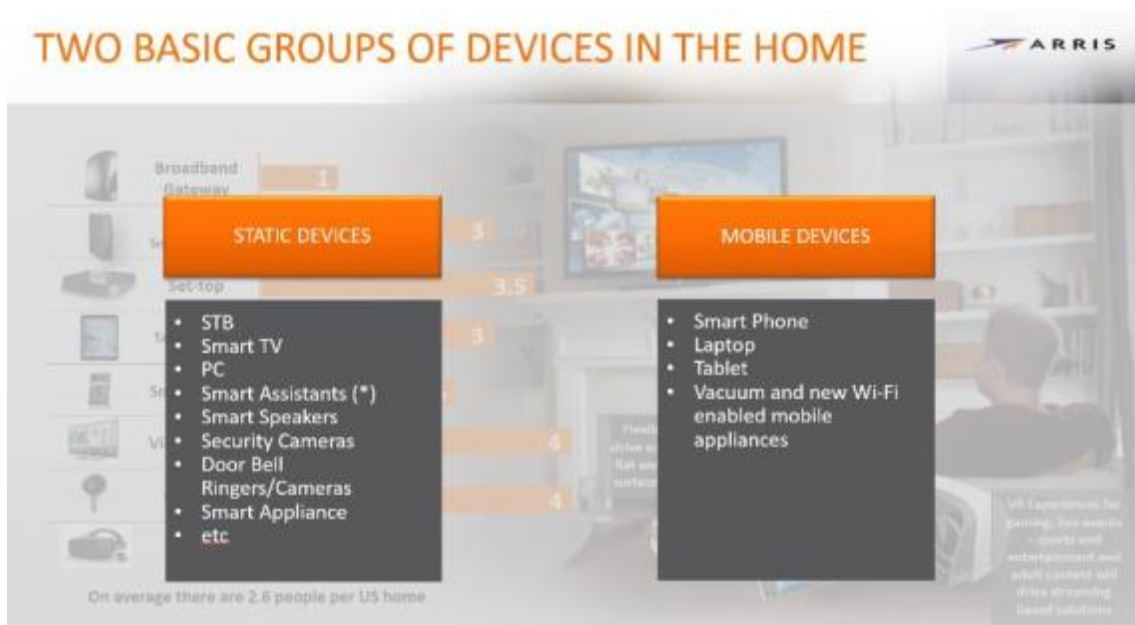


Figure 10 - Static and Mobile Device Groups

In these two basic groups it is important then for all Machine Learning and Wi-Fi algorithms to discern the location of these devices. One coarse subgrouping is

- Devices in house reasonably close to the AP(s)
- Devices exterior to the house or furthest away.

This is key because more and more the home has cliff edge Wi-Fi events that drive all the issues or the changes in Wi-Fi architecture. A simple example is the addition of a Wi-Fi based camera doorbell on the outside of the house. This then can generate a major change in Wi-Fi from

- Poor connectivity at the point of install driving an extender to be added
- Lower MCS device bringing down potential airtime for all the other devices
- A continually transmitting device that keeps adding to the AP utilization and always running – affecting even 802.15.4 and Bluetooth devices as they try and compete with Wi-Fi airtime transmission slots.

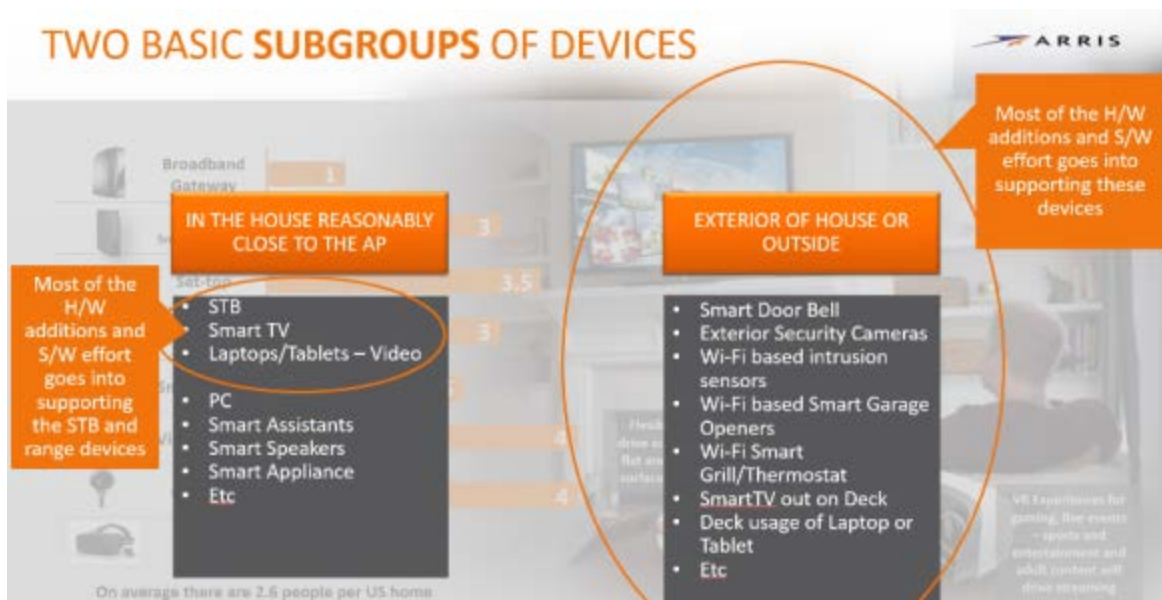


Figure 11 - Device Groups Based on Location of the Device

There are additional categories as well that any service provider Wi-Fi solution algorithm needs to also be cognizant of and use the inherent capability of the device, Service Overlay (if any, whether managed by the service provider directly or OTT applications that the subscriber uses) and finally the usage pattern. These three categories and the various choices or aspects within each result in a plethora of combinations, making the management challenge daunting. Let us examine the diversity in greater detail below.

The key parameters of this category include the:

- Generation of the Wi-Fi standard and associated advanced features that the device supports - determined to a large part by the Wi-Fi silicon capabilities
- Device-specific features based on hardware implementation including the number of spatial streams and the overall transmit power amplification and antenna design
- Support in software through device drivers for advanced features to aid the performance of the device

Since the start of the Wi-Fi standards evolution nearly two decades ago, there have been multiple generations of Wi-Fi standards, starting with 802.11, 802.11a, b, g, n, ac and now ax, that even today, 802.11n devices are extant.

The second category is the nature of the Service Overlay. These could be explicit services provided by the MSO such as Linear and On-demand video, Telephony, Security, and Home management along with the basic data services.

The last category is the usage pattern, in terms of locality & mobility, temporality and consumption.

- Locality & mobility distinguish static clients such as a Wi-Fi set-top box in a fixed location, versus a portable device such as a laptop and a totally mobile client such as a smartphone which could move around even as it is being used.
- Temporality refers to the time-based usage pattern typically exhibited by the device. Some devices stay connected, but are rarely used for data consumption, as in a SMART TV which is

connected to an external device for source. Some others may be used during specific times of the day, and finally a few that are always on and transacting data, such as an IoT device.

- Consumption refers to the amount of data usage typically exhibited by the client, and may be broadly grouped as Low data usage, Bursty data usage, and Heavy data usage.

The figure shown below captures the diversity of the Wi-Fi Clients, and one can imagine the number of combinations, especially when factoring the manufacturer of the device.




Key Category	Main Aspect	Parameters that define the variations to the category			
Device Capability	Standards				
	Implementation	Spatial streams from 2x2 to 8x8	Power Amplification	Antenna Design	
Service Overlay	MSO Services				
	OTT Applications				
Pattern of Use	Location	Fixed location	Portable but fixed	Mobile	
	Temporality	Always ON	Daily Pattern	Rarely used	
	Consumption	Low	Bursty	High	

Figure 12 - Diversity of Wi-Fi Client Devices

In addition to all of the above that relates to the nature and use of a given Wi-Fi client, the geographical location of the device, in relation to the Access Point (AP) is an exceedingly important factor in the performance. It is not just the distance, but the presence of walls and material that have an attenuating impact on the signal, directly affecting the maximum possible throughput that a device can hope to achieve. Another factor that has an impact on the Wi-Fi performance is the level of interference due to Wi-Fi signals in the house due to other devices within or from outside the specific residence.

The purpose of discussing the diversity of Wi-Fi client devices is to highlight that any solution must consider this reality. Knowledge of the exact nature of the device is essential for a more informed decision process, be that a manual one or an automated algorithm.

3. Customer Expectations and the Burden of Cable Service Providers

Cable service providers face the challenge of providing Internet service through the complicated medium of Wi-Fi. Should everything work as needed, most subscribers would assume this as the 'basic expectation being met'. However, at the first sign of any issue with connectivity, coverage, performance or such aspect, the service provider is drawn into the problem, notwithstanding that the source of the problem could be elsewhere in the home, unrelated to the internet service itself. As we have seen in the section before, connectivity, poor performance and coverage depend on a wide variety of factors.

Sub-optimal conditions can also affect the service overlay, and should the service be one that is offered by the service provider, then implications fall on the provider notwithstanding the device capability. It is for

this reason that the providers as of now insist on equipment supplied by them to provide a specific service (such as a Wi-Fi set-top box to provide video over IP over Wi-Fi service) as opposed to relying on a device that is customer owned and managed (COAM) device.

Yet another factor that works against the incumbent service provider unfairly is the asymmetric expectation across modalities of service delivery. The same subscriber who is willing to allow IP video to an OTT device that occasionally buffers or has artifacts will not tolerate a freeze of video or reduced quality in linear video content provided by the customer. This increased tolerance level is due to the fact the subscriber is compensated somehow otherwise, as in the case of OTT video, the convenience of content selection or cost. This is not different than how consumers were willing to tolerate the poor quality of cellphone voice calls because of the advantage of mobility that the cellphones were able to provide. However, it is believed that Services providers will want to offer highest profile IP video to managed IP Video STB devices and not rely on Adaptive Bitrate to be the core of their service offering. This should differentiate the SP offering for HDR and immersive video services.

Perhaps the single biggest gap between customer expectations and reality is in terms of the performance of Wi-Fi throughout the home. The terms that we often encounter here are “Coverage”, “Performance at Range” and so on. It is important to dwell on this subject a bit. The problem space as well as the solution space is sufficiently complex, and we must expand on the statement made earlier in the introduction when we mentioned that customers need their Wi-Fi Speeds to match the Access Speeds (that they pay for).

Performance of a Wi-Fi client is highly dependent on how close the device is to the Access Point (considering obstructions to the RF signals, more than just distance). The effective distance (considering the obstructions, resulting in a normalized measurement called attenuation) is also referred to as “Range”. By the laws of physics, the signal strength (and therefore, the resultant performance) degrades over range, the fall off being more rapid and dramatic in the 5 GHz channels than with the 2.4 GHz channels. Given that 5 GHz channels offer higher throughputs, the fall over a shorter range is even more significant.

The following diagrams highlight the gap between expectations and reality visually, from the perspective of a typical home.




Figure 13 - Expectation: Wi-Fi Speeds = Access Speeds

From a purely commercial perspective, this seems very reasonable. A centrally placed access point (4x4 DBC or 4+4+4 TBC) can offer whole home Wi-Fi services to 500 Mbps speeds in 2,500-5,000 square

feet homes. However, centrally placed APs and GWs are hard to engineer and often the AP is biased to the access network connection point at the outer wall of the home. This then typically necessitates a second AP to get the desired throughput to the extremes of where consumers' devices need Wi-Fi connection. After all, the Service Tiers that are advertised by the service provider are based on Access Speeds and given that the consumers pay steeper monthly access fees for increasing service tiers, there is a natural expectation of the availability of the same speeds on any device, at any location in the room.

Here is a picture of (what we would like) the current reality: The picture illustrates that additional APs are added – which increases the complexity of management and with Wi-Fi those devices also cause additional congestion and interference in home and to the neighbors' home unless properly managed:




Figure 14 - Reality: Attenuation and Interference Affect Performance

Closing the gap between expectations and reality, with the use of additional access points in the home whose operations are coordinated for optimal performance – has been a focus of the industry for the last couple of years. It has spawned the following initiatives

- Retail Multi AP strategies to align with the consumers frustration when Wi-Fi goes wrong
 - SP's own analysis for the right Multi-AP strategy – form and function of the extension AP
 - AC outlet mounted lower power 2x2 type managed repeaters
 - Tri-band 4+4+4 ultimate Extenders
 - Wired vs Wireless Wi-Fi extension of backhaul
 - Wired works – but causes OPEX costs and dissatisfaction with consumer. Rarely self-install as it requires Ethernet wire pulls or potentially simple MoCA filter.
 - Wireless – lowest OPEX cost and minimal effort from consumer (when done right).
- Challenge to get a Wireless Mesh that is not a repeater and inefficient on the Wi-Fi airtime.
- 2x2 Wireless Mesh devices – are still repeaters but can be made more efficient with mesh management software
 - Tri-band extenders are more ideal for the applications but have challenges of
 - Size and Cost

Table 1: Cost Delta with Tri-band Configurations

Extender Type	Cost	Comment
2+2+2	\$	Can fit on AC outlet
2+2+4	\$\$	Needs a Fan or lower Power Wi-Fi
4+4+4	\$\$\$	Maximizes 1W for range but table top only

- Software to manage the APs
 - Radio Resource Management (RRM) and Self-Optimization Networks (SON) initiatives
 - Cloud Wi-Fi analytics
- Standardization
 - Multi OEM and Multi Silicon solutions highly

Retail Wi-Fi devices that act as secondary or tertiary access points have entered the marketplace, promising easy setup, performance and management. This is a clear and present threat to the service providers, with the power to relegate their role to just an access provider, bypassing the complete data distribution. Without access to the devices beyond the network termination unit, the service provider will not be able to bundle services in the same manner as they are able to do now.

While the retail devices have acted as an interim solution for a consumer to Wi-Fi problems, they are not necessarily solving the full problem. Most solutions that are available today, while definitely slick on industrial design, ease of setup and deployment, and ease of self-service by the consumers, come nowhere near to addressing the full expectations of Wi-Fi Speeds matching access speeds.

The picture painted above, with the bewildering variation in the Wi-Fi clients as well as the tough road for incumbent service providers, is not necessarily one with a gloomy ending. There is help on many fronts: Standards are continuing to evolve, pushing performance boundaries, while also showing signs of maturity in terms of manageability and steps towards interoperability, signifying a ‘coming of age’ stage for the Wi-Fi technology. Concomitant availability of practical tools to take advantage of the advances in areas like Artificial Intelligence (AI), Machine Learning (ML) and Big Data will help drive complexity out of the solutions to solve the challenges. Finally, mature cloud-based solutions are available to help manage and visualize the connectivity, coverage, performance and other aspects.

In the following sections, we will examine the goodies that are available to us to ensure that the service provider can wrest the control back and provide meaningful managed Wi-Fi services to their subscribers.

Unwrapping the Goodies

The toolkit for enabling better management is now in our hands. We will examine developments along the following major categories, all of which are crucial to have a technically feasible and yet commercially viable set of solutions.

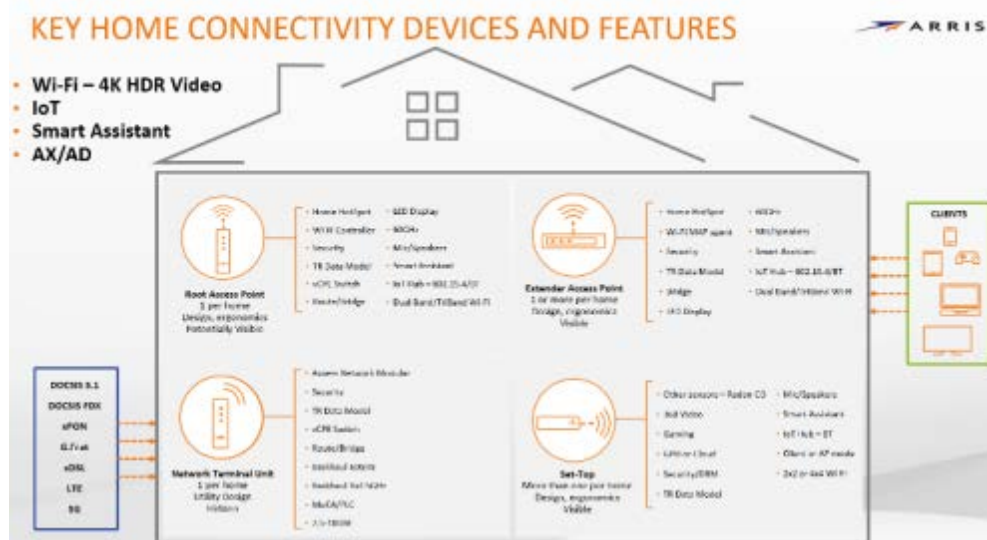


Figure 15 - Typical Wi-Fi Device Distribution in a Home

3.1. Choose Your primary AP

- All in one Access and Wi-Fi Gateway
 - Lowest CAPEX for single smaller homes
 - Tends to bias towards the wall for location
- 2 box solutions
 - Standalone e-MTA and ONU
 - Standalone AP device with Ethernet WAN
- Decisions then on what level of Wi-Fi to add to the Primary GW
 - 4x4 DBC is now the standard -> moving to AX
 - 4+4+4 Tri-band is a desire and moving towards with AX
 - 8x8 DBC with ability to manage radios
 - All have different performance and cost points and size constraints. The industry is trending towards Tri-band to
 - Support high > 1 Gbps on Wi-Fi
 - Allow one of the 5 GHz channels/radios to be allocated to Wi-Fi backhaul when additional extenders are added to a Multi-AP home

3.2. Choose Your Multi-AP Topology Strategy

- Wired AP solution – using Ethernet, MoCA, G.hn
- Wireless AP solution – using Wi-Fi meshing

3.3. Choose Your Multi-AP Device Architecture

- DBC – managed repeater
- TBC – optimized for airtime efficiency and performance
- Table top mounted AP
- AC outlet mounted AP
- Add additional services like IoT and Smart Assistant

3.4. Choose Your Meshing Solution

- In Gateway Wi-Fi controller only
- In Cloud Wi-Fi controller only
- Hybrid Wi-Fi Controller in GW and Cloud
- Full mesh 802.11s
- Partial mesh
- Hub and Spoke
- Hub and Spoke with Multi Hop (EasyMesh architecture)

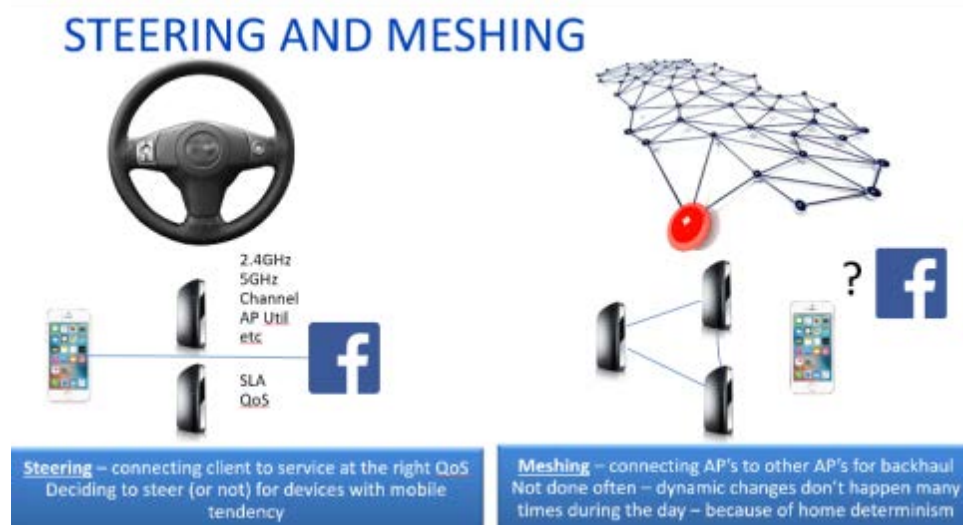


Figure 16 - Steering and Meshing: Defined

3.5. Choose Your Steering Solution

- Blacklisting
- BTM steering
- Fast Transition

3.6. Choose Your Cloud RRM/SON solution

- Cloud managed Wi-Fi RRM to manage across Home Wi-Fi domains or MDU
- Local GW based Dynamic Channel Selection solutions

3.7. Choose Your Wi-Fi Telemetry Strategy

- Choose your Wi-Fi data model in the AP and Extender devices
 - TR181+
 - WFA Data Elements
 - BBF USP
- Choose your pull and push strategy and Cloud Connection
 - RESTful

- WebPA
- USP
- COAP
- MQTT
- WFA Data Elements

Decide frequency of collection, what to collect, what to compress, and what to filter.

3.8. Add Your Machine Learning and AI Roadmap

- Analyzing Wi-Fi telemetry to update Gateway Controller policies for Wi-Fi steering, Device policies, service policies

4. Wi-Fi improvements Worth Noting

There are a number of areas that this paper will expand on – they fall into 3 categories below

- **Work of standards bodies** coming to implementation fruition. Sometimes seen as an undue constraint on the speed of innovation, this bridle is a necessary step in coordination across the entire industry, placing emphasis on interoperability and the importance of a viable ecosystem that will allow for network effects to take place and costs to be manageable.
- Group of three inter-related technologies of **Artificial Intelligence (AI), Machine Learning (ML) and Big-Data Analysis**. These technologies are now at a level of maturation to allow for their application to solving problems across multiple disciplines, and Wi-Fi management is no exception. We will examine specific problems in the area of Wi-Fi management as examples where the application of these technologies is apt.
- The third category is more a paradigm shift, than a specific technology. This involves moving the platform for solving several management challenges from the devices to the cloud. Headend and back-office infrastructure always existed, and cloud-based service or management is not new. What we identify here are specific examples where the visibility across homes, the availability of compute power and data storage, and the inevitable association with AI / ML and Big-Data Analysis.

5. In-Home Device Implementations Catch-up to Standards

One of the issues with Wi-Fi Controller/SON or RRM solutions was that they were proprietary in nature, agent based and while they flattered to allow potential porting to be done – in reality they did not allow this easily or at all. Standardization in this area was required and this spawned the work that ultimately ended up being certified by WFA as Wi-Fi CERTIFIED EasyMesh. Additionally SCTE 2018 marks the time when the most important thing to happen to Wi-Fi since it first emerged is now real and relevant – 802.11ax. An IEEE standard that forms the basis of 90% of all home connections. The following sections briefly outline the importance of these 2 standards to the service provider's strategy.

5.1. 802.11ax: Aimed at Improving Efficiency and Performance

Perhaps the most important standard that is being driven from an implementation perspective to fruition is the 802.11ax. There is enough material on this topic that in the context of this paper, we will just mention the importance of this, rather than explaining the technology itself.

Earlier in the paper, we noted that number of devices in the home, increasing bandwidth requirements of video, and the advent of IoT devices as key triggers to demand increased performance, and the 802.11ax standard strives to address all of these. There is a clear market demand and a causal connection that is illustrated in the diagram below:




Figure 17 - New Technologies Demanding drive for Wireless Performance

The standard improves spectral efficiency in dense client environments (that should help when the number of Wi-Fi client devices in the home increases), with a concomitant increase in effective throughput, taking us closer to realizing the goal of Wi-Fi speeds matching access speeds. See the simple chart below that illustrates all the benefits of the new AX standard.




Figure 18 – Key Benefits of 802.11AX

However, the key feature of ax is the ~27% of improved efficiency over ax that can be immediately realized particularly in Multi-AP architectures. See the chart below – where we can achieve 1 Gbps thru 2 sheetrock walls vs 846 Mbps in 802.11ac. This is an enormous benefit for minimizing the number of extender APs and also increasing the flexibility of location of the extender device.

802.11ac Performance (5 GHz UNII-1 Band @ 30 dBm) Balanced Client Device



Backhaul UNII-3	Link Rate to Root AP			
Configuration	0 Wall	1 Wall	2 Walls	3 Walls
80 MHz 4x4 ac	1290.27	1290.27	846.34	360.75
80 MHz 4x4 ax	1958.04	1538.47	1009.72	430.42
160 MHz 4x4 ac	2709.39	1763.04	736.27	355.55
160 MHz 4x4 ax	3711.18	2128.47	890.81	430.42

WARNING!!
160 MHz Channels Require use
of DFS Frequency Band and
Power

Backhaul DFS	Link Rate to Root AP			
Configuration	0 Wall	1 Wall	2 Walls	3 Walls
80 MHz 4x4 ac	1290.27	947.12	550.99	176.76
80 MHz 4x4 ax	1958.04	1129.83	657.41	210.80
160 MHz 4x4 ac	2709.39	1763.04	736.27	355.55
160 MHz 4x4 ax	3711.18	2128.47	890.81	430.42

Figure 19 - 802.11ax Performance Gains over 802.11ac

Another important design goal for 802.11ax is the consideration for devices (especially IoT devices) that are battery operated, by way of improving the efficiency of operations to allow for low power-consumption. 802.11ax allows for efficient allocation of low data-rate connections, and for improved battery life of sensors. Power savings modes have been enhanced, to include longer sleep intervals and scheduled wake times. Many IoT devices implement a 20 MHz channel only, and the standard now takes that into account to have a “20 MHz channel-only mode”, to take such devices into account.

5.2. Wi-Fi EasyMesh

Even with 802.11ax and the increased efficiency, it is impossible to achieve full coverage across a large home with a single access point, however powerful it is, since it is usually hard to argue with the laws of physics. The introduction of multiple access points in the home immediately gives rise to several technical challenges.

The primary and foremost challenge relates to the problem of onboarding the additional access points to ensure that these units have the same SSID and passphrase as the main access point. After all, while the multiple access points can be utilized as separate entities, it is obviously not a meaningful proposition, since it would require the consumer to segregate the Wi-Fi clients and statically associate them to different access points, each of which have a different SSID. The obvious solution is for all the access points to share the SSID and Passphrase information seamlessly. This is in general referred to as “Auto Configuration” or “Zero Touch Configuration”. The configuration should not only be done at initial set-up time, but also whenever there is a change in the configuration at the main access point (like the consumer changing the password), that these changes are propagated.

The process of auto configuration (or zero touch configuration) has to be solved independent of whether the secondary access points are connected to the main access point via a wired connection (such as Ethernet or MoCA), or a wireless connection either by sharing the 5 GHz band or a dedicated radio in Tri-band configurations.

The challenge cited above is not purely technical in nature but reflective of the need for interoperability. To ensure the setup and ongoing synchronization, there needs to be a common protocol for communication across the access points. While there have been proprietary solutions, true multi access point solution has relied on the development and adoption of a common protocol.

The Multi Access Point Protocol (MAP) started as a Special Interest Group activity, and since then has been adopted by Wi-Fi Alliance (WFA), the standards body. It has since been renamed as Wi-Fi EasyMesh™ and includes a certification process.




Figure 20 - Wi-Fi Certified EasyMesh™ (source: <http://www.wi-fi.org>)

The EasyMesh architecture supports both wired and wireless backhaul links from the secondary (additional) access points to a central gateway. The Gateway is that Access point which is connected to the WAN network. The architecture accounts for two software components, the ‘Agent’ and the ‘Controller’ and in general supports topology variations. The following diagram depicts a typical scenario.

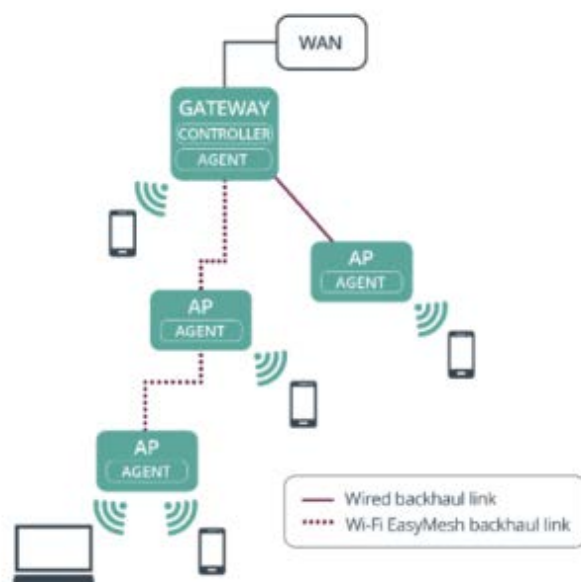


Figure 21 - Wi-Fi EasyMesh Architecture

The key communication protocol is the contract between the “Agents” that run on different platforms, and are responsible for discovery, onboarding and subsequent information exchange. The “Controller” typically runs on a gateway (as has been depicted in the diagram above and below) and manages the various network clients in the home. This standardization allows any EasyMesh compliant Extender to be managed by a nominated controller AP. This then allows a SP to mix and match Extender OEM and silicon providers easily.

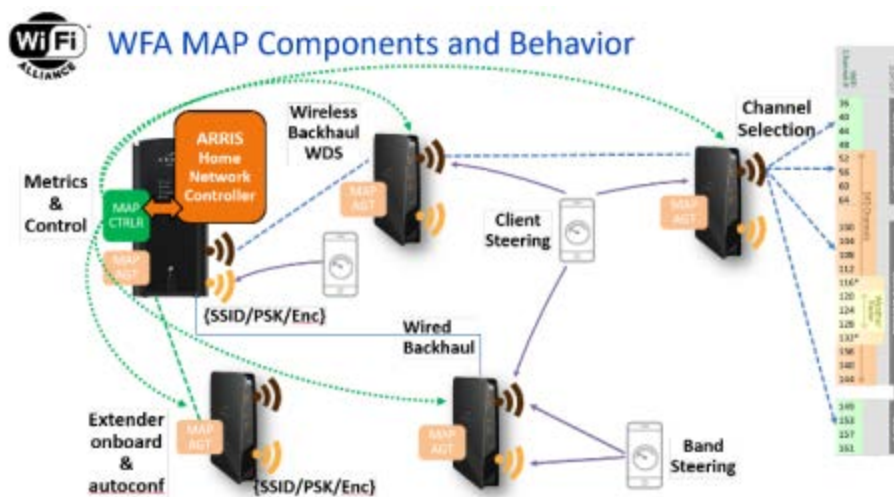


Figure 22 - Role of WFA Standards Based Components

The EasyMesh MAP (Multi-access Point) controller – creates the topology of all the agents APs in the discovered network. The MAP controller will coordinate sending information to and from the MAP agents. These commands are Wi-Fi telemetry and Wi-Fi commands such as change channel, band steer, change power – the suite of Wi-Fi control commands you would expect to have in a Wi-Fi AP network.

However, the MAP controller does not make Wi-Fi algorithmic decisions. It needs a Wi-Fi controller (algorithm) to tell it what to tell the MAP agent-based APs. This makes the algorithms agnostic of Wi-Fi command/control and allows the SP to select a Wi-Fi algorithm solution – as implementation in GW/AP above the MAP controller level.



Figure 23 - Criticality of AP Steering

The architecture, while securing interoperability, is flexible enough to allow intelligent Radio Resource Management (RRM) and Self-Optimization Networks (SON) schemes to be implemented by solution providers to differentiate as well as address the diverse needs of the service provider community.

5.3. 802.11v and 802.11k to help “Roam (in) Sweet Home”

While Wi-Fi EasyMesh addresses the basic connectivity and information-exchange challenge across the multiple access points, the second level challenge is what is referred to as the “Sticky Client” problem. A mobile client that is associated with a given access point may move (who has not walked around the home while talking on the cellphone?) to another location in the house where the signal strength relative to the currently associated access point is so weak that it affects the performance of the said device. The presence of a nearby access point (which might offer a better signal), if not taken advantage of, will be absolutely useless.

Having the client be associated with the most appropriate client dynamically, without any action on the part of the user (a problem that is already solved in the enterprise space), becomes important. Moreover, when there are many clients, it may be prudent to balance the load on the various access points so that the load is equitably distributed. In addition, there may be preferential treatment to specific clients based on the services that are consumed by them.

In order to achieve these solutions, especially in a multi-vendor home ecosystem, it is imperative that the various access points (and the clients) share information that help and assist the seamless movement of the client association from one access point to another.

802.11k standard defines creation and transmission of neighbor report lists. Neighbor reports contain information about the neighboring access points (from the perspective of the AP providing the list), and are transmitted to a client that supports this protocol. These neighbor reports allow the client device to have a clearer picture of the Wi-Fi surroundings and allows the client device to prune the list of channels that it has to scan before finding a suitable candidate neighboring AP to associate with.

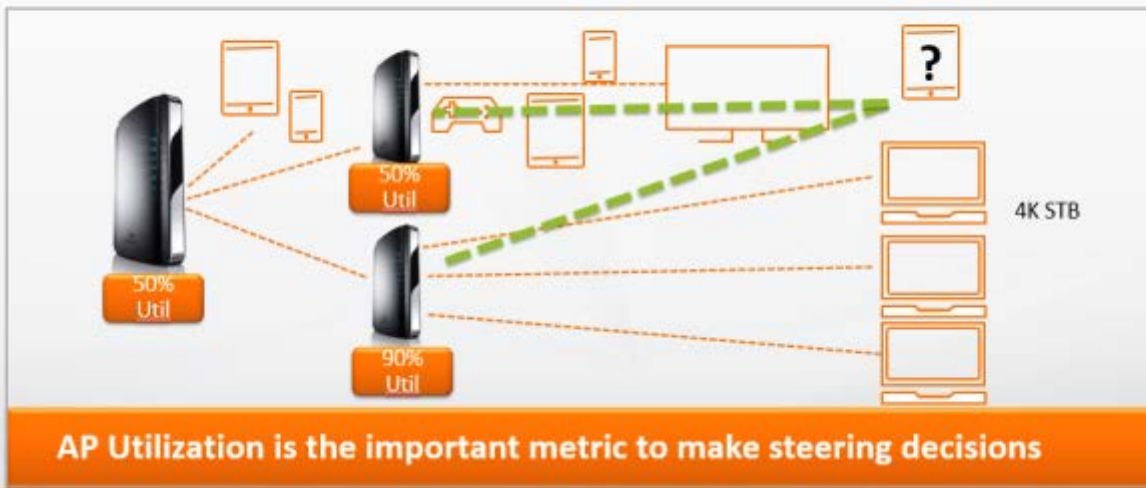
802.11v standard is aimed at smooth and fast transition of a client to another access point, directed by a controller. This transition is initiated by the Controller and allows for the use of intelligent algorithms to make such a decision. Access points that support 802.11v can direct clients (that support the 802.11v protocol) to roam to another AP which presumably is intended to provide a better Wi-Fi experience for the client device. The client devices will have to accept and respond to the Basic Service Set (BSS) Transition Management (BTM) frames.

The example highlighted below shows that BTM steering in particular allows the SP for the first time to control which AP the device connects to vs the device making all the decisions.



Figure 24 - Signal Strength Alone is not a Unique Determinant of Performance

Best Connection may = $\text{Airtime}/\text{Avail}/\text{Util}$



Best Connection may = $\text{Airtime}/\text{Avail}/\text{Util}$

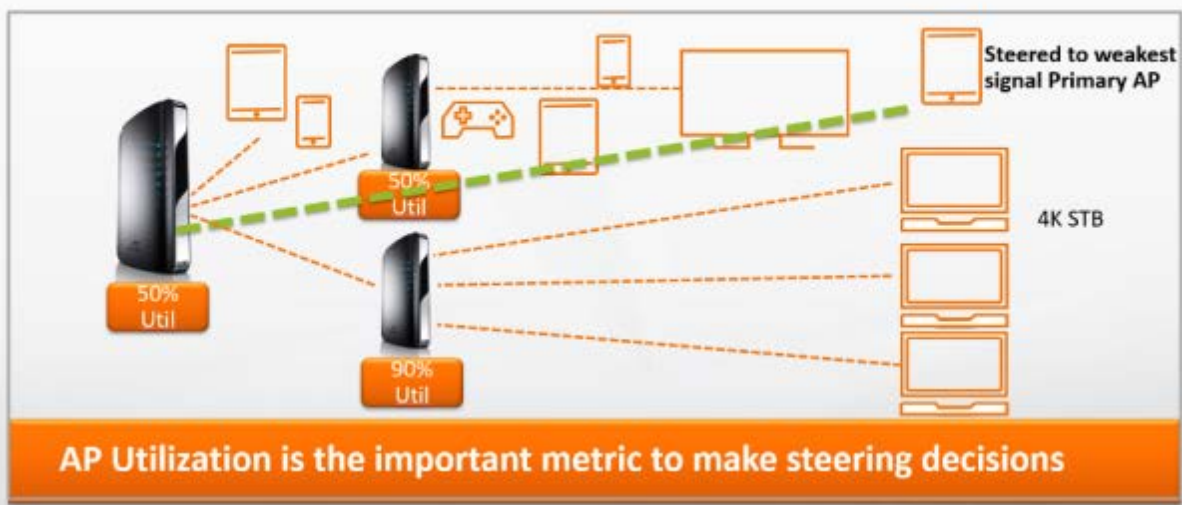


Figure 26 - Role of AP Utilization in Steering Decisions

5.4. Wi-Fi Easy Connect

While the standards listed above deal with the Access Point infrastructure inside the home, we need to turn our attention to adding the ever-growing list of Wi-Fi clients to the home network. Given the wide diversity of the clients, the onboarding of retail devices to the home network is anything but trivial. As engineers, we tend often to overlook the complexity involved for an average consumer. Onboarding a typical retail device to be a client of the home Wi-Fi network involves multiple steps and is often left to the retail device manufacturer to solve it in whatever manner that they consider the easiest.

Wi-Fi Easy Connect is an emerging standard that is part of the Wi-Fi Alliance set of emerging standards. The architecture proposed allows for a simple process that involves ‘one-touch provisioning’ assuming the various actors in the architecture follow the Device Provisioning Protocol Specifications that are available to WFA members.



Figure 27 - Device Provisioning with Wi-Fi Easy Connect

There are requirements imposed on access points, clients, and a requirement on a configurator, which is typically a mobile application. Client devices are required to have QR code as part of the device, in order to be enrolled as clients to the AP. (Source: <http://www.wi-fi.org>)

Both EasyMesh and Easy Connect are examples of technologies that require a critical mass of supported products in the entire ecosystem in order for network effects to kick in and make them ubiquitously used.

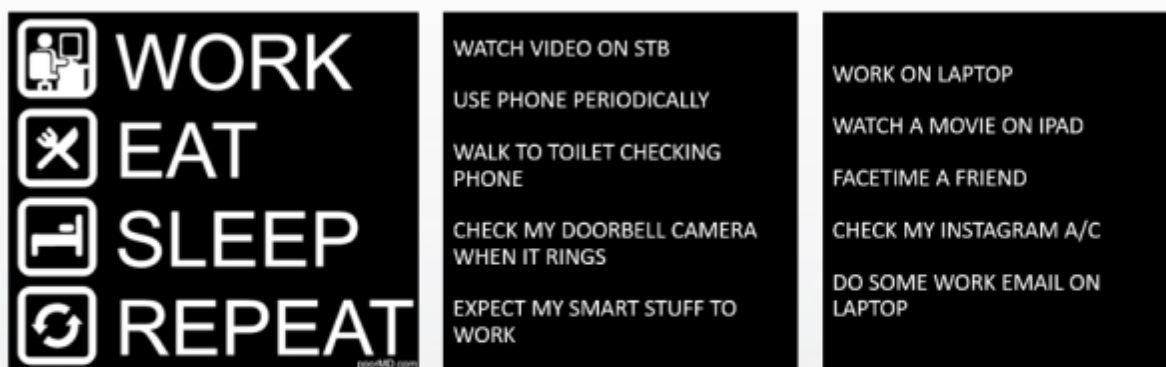
6. AI / ML techniques

Artificial Intelligence, Machine Learning and Big-Data Analytic techniques are inter-related technologies that have moved from being just hype-phrases to providing practical toolsets to solve problems that are amenable to the use of such techniques. We don't go into details here, given that these are huge subjects and any treatment will be not do any level of justice. Our main intention is to point out that some of the problems that we encounter in Wi-Fi management lend to solutions that utilize these techniques. And in particular to point out that Wi-Fi homes are extremely deterministic in nature and can be graded into a small subset of categories – that help with the management services to the home, again keeping with the simple theme – and the simple groupings of device types and service types above.

The human behaviors chart below shows that each person, household tends to have a specific signature:

- Number of people
- Number of devices
- Type of devices – high bandwidth, security, etc.
- Mobile usage
- Outdoor in garden usage
- Security cameras
- Times they are at home
- How they use their devices at times in the day
- How they set their home up on vacations
- And more

THE DETERMINISTIC HOME WE CLUSTER AND ARE CREATURES OF HABIT



PEOPLES HABITS ARE VERY DETERMINISTIC ; THEIR HOMES REMAIN DETERMINISTIC TOO

Figure 28 - Consumer Behavior

The STBs in the home in particular – need to use the Wi-Fi toolbox as a high need device. We know that if the consumer has any problems with primary TV viewing this can be a churn event. So these devices need to be signature found and policies applied to make them work better than OTT devices.

SIMPLE POLICIES FOR THE WI-FI HOME

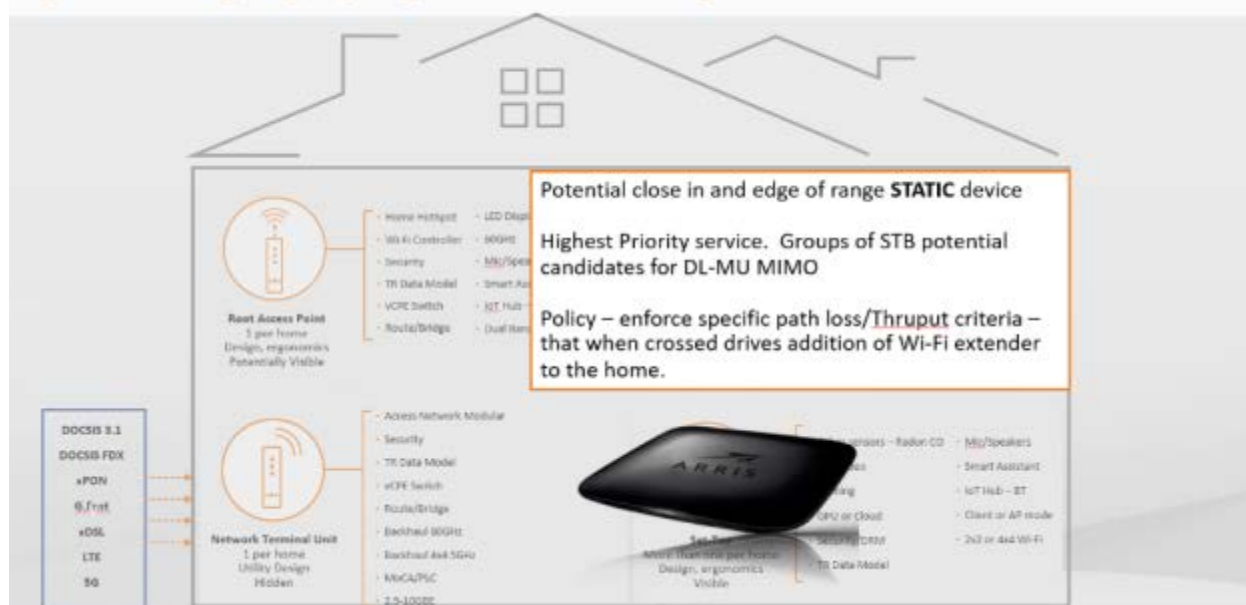


Figure 29 - Policy to Adopt for High-Bandwidth Set-top Box

Other static devices are OTT video consumption devices that are also high value and need the ML solutions to signature their use and setup management profile to work with them.



Figure 30 - Policy for Static High Definition TV

Mobility in the home and the assessment of the Mobility Index (how often people are mobile with traffic or not) is something that is also key for the machine learning elements of Wi-Fi management.

SIMPLE POLICIES FOR THE WI-FI HOME



Figure 31 - Policy for High-Bandwidth Mobile Devices

Tools like the one below tracking the mobility of devices can also help to assess how to enforce steering and Wi-Fi policies as well as device architectures. The images below

- Color shows health – Green Good, Yellow Ok, Red below SLA
- Size of bubble – amount of data consumed on device
- Right to left bubble track – the RSSI location of the device

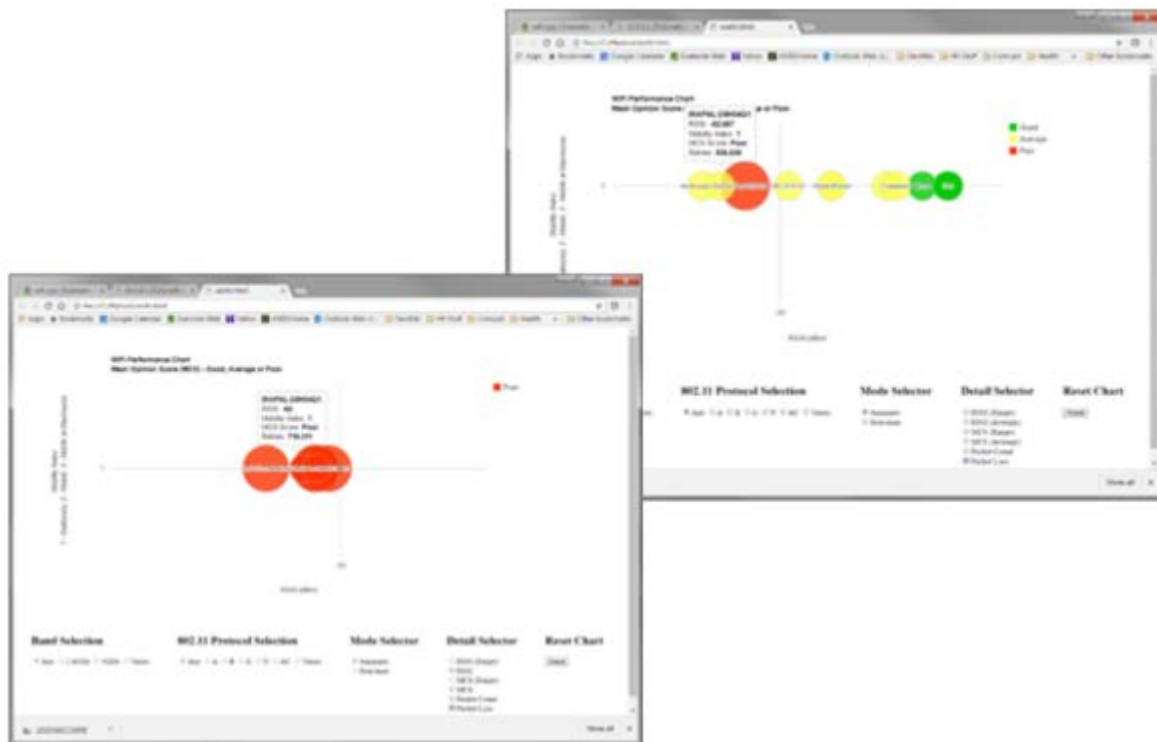


Figure 32 - Visual Representation of Device Coverage Health

The Machine Learning and AI policies also need to signature capture medical or security devices which have high worth metric and potential service SLAs on them. They need to get best Wi-Fi too.



Figure 33 - Policy for Static High-Priority Elements

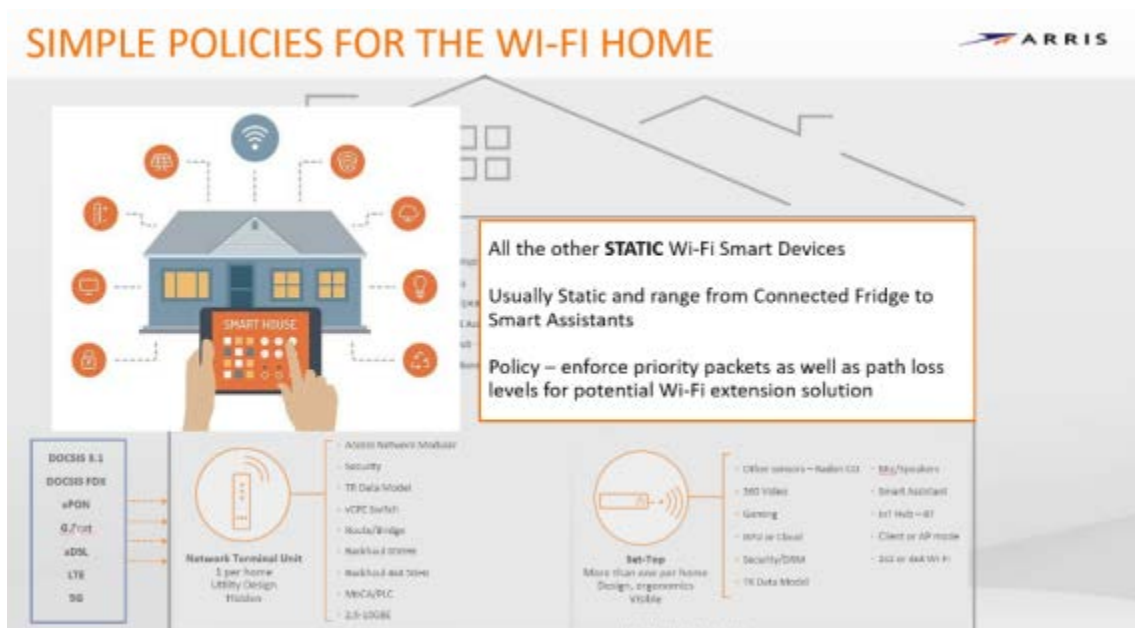


Figure 34 - Policy for IoT devices

The popularity of Wi-Fi devices like Ring doorbells and Wi-Fi controlled garage door openers has had a big effect on both Wi-Fi performance (long range low modulation bad apples when added to home) as well as driving the need for Wi-Fi extension additions. Many of the doorbells and cameras also prompt consumer to buy a same brand Wi-Fi extender to ensure that their service is optimized. The ML algorithms used need to grade/categorize the home – for this particular static device type and create special policies to manage.



Figure 35 - Policy for Devices at the Edge of the Home

One of the other key machine learning and categorization of the home – is whether the consumer is using their Wi-Fi devices outside the range of the home. The garden and deck usage of mobile devices like tablets which are not permanently attached static outside security or IoT devices but through mobility to the back garden the consumer is getting a poor Wi-Fi experience – and will blame the SP for poor Wi-Fi. Using Wi-Fi location and trending of mobility performance – the ML algorithms can infer that the user is trying (maybe unsuccessful) to use mobile device with high bandwidth at extreme distance from the home.

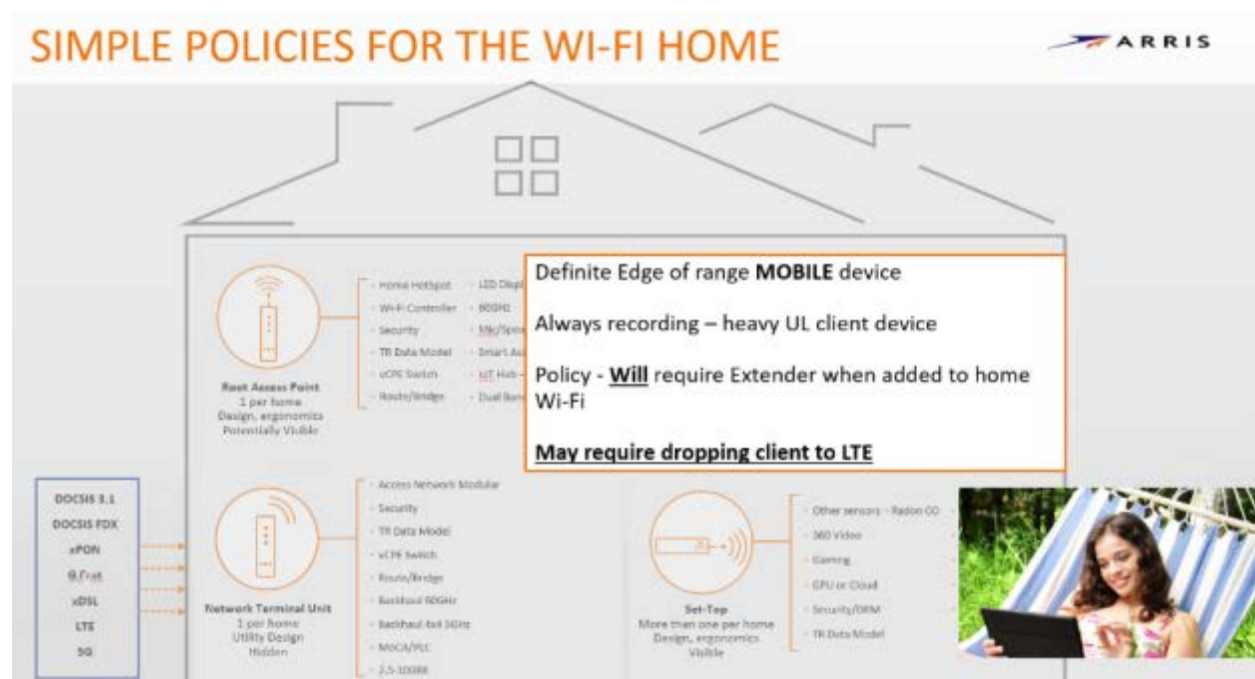


Figure 36 - Policy for Devices Outside the Home

The following table lists the data analysis categories that AI/ML tools excel at, and identifies potential applications relevant to Wi-Fi management:

Table 2: Categories of Data Analysis problems that Ai/ML Techniques Excel at ?

Analysis Category	Description	Potential Application
Classification	Identifying the category to which an object belongs	Device Fingerprinting
Regression	Predicting a continuous valued attribute associated with an object	User Behavior and prediction
Clustering	Automatic grouping of similar objects into sets	Device categorization; Resource Optimization
Dimensionality Reduction	Visualization and increased efficiency	Actionable Reports and Analytics
Model Selection	Comparing, validating. And choosing parameters and models	
Feature Extraction	Preprocessing and normalization	

There is a plethora of open source and commercial tools that offer a starting point for the development and fine-tuning of algorithms and solutions as applicable to the challenges in Wi-Fi management.

7. Cloud Based Analytics

Cloud-based SW architecture is not new. Back-office systems that support control Wi-Fi is not new either. In the Enterprise segment, Access Point controllers are traditionally based in the cloud and perform the management functions as a matter of routine. What we are highlighting here is the use of Cloud-based platform to help the analysis of data to support Wi-Fi management not only in a home setting, but also across multiple homes (like an MDU or neighborhood).

While Gateway based software-controllers offer low-latency and provide fast turnaround decisions for steering and roaming, they lack two crucial advantages that a cloud-based system can offer. The gateways don't have the compute power or the storage capacity to handle large amounts of data over time, to do time-series analysis for trend predictions and other such statistical analysis. Secondly, the gateways don't have visibility to other Wi-Fi devices such as neighboring access points and other wireless devices that will have an impact on the performance because of interference or such reasons. Lastly, without a cloud infrastructure, there is no way for one Gateway to learn from the knowledge, data or experience of other gateways.

For these reasons, a hybrid architecture where the local software-controller takes care of low-level execution, being informed of a management policy that gets articulated by a cloud-based system makes immense sense.

8. Goodie-Bag Summary

In the foregoing sections, we touched upon the various tools and standards that are available to us to be able to address the needs of the consumer in a meaningful way. In the next section, we will take each of the consumer needs and exemplify how the techniques and tools help us to tackle the challenges cited.

Addressing Consumer Needs with the Goodies

In the previous sections, we have seen in some level of detail, the challenges of Wi-Fi management, and also the tools and standards that are available for us to create the solutions. It is time for us to revisit each of the six critical needs that we stated as consumer needs for the Wi-Fi Home network. In this section, we go through each of these needs, and provide examples of how some of these needs are being addressed.

9. Connectivity & Security

The wide diversity of Wi-Fi clients imposed on us the challenge of onboarding as well ongoing management and maintenance of these clients.

In the section on Standards, we alluded to the upcoming standard for Wi-Fi Easy Connect via the WFA organization. However, there are many retail devices that are already in the field and many more that will be manufactured before the standards are implemented and the devices have the requirements (such as QR code meeting the specifications) met.

The problem itself can still be adequately addressed as long as the solution component includes a suitable tool for the role of the configurator such as a mobile application. The following diagram shows how a

retail device that has either a QR code or a WPS button can be utilized along with an application to initiate simple onboarding process.

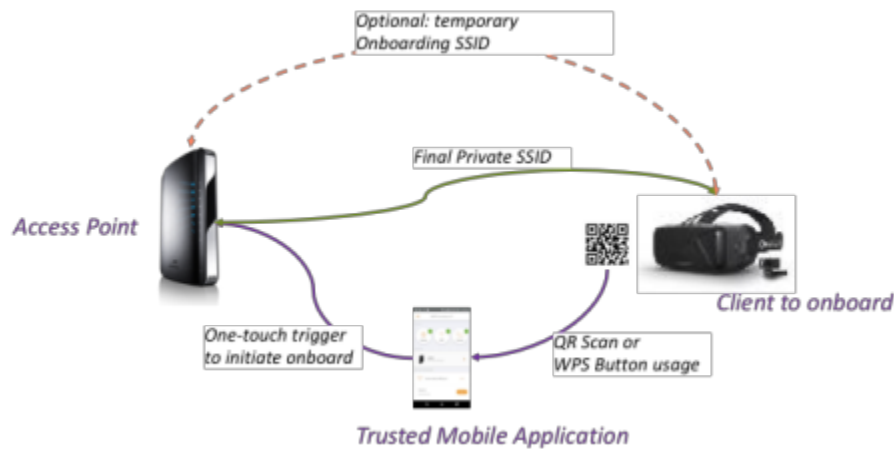


Figure 37 - Wi-Fi Client Onboarding Process

The use of the temporary onboarding SSID and the QR scan are intended to ensure that an arbitrary device will not be able to get onboard the home network. The arbiter here is the mobile application, which typically needs the administrative access to the access point for its operation, and hence assumed to be in possession of the custodian of the access point.

In case the device has a WPS button, there is still no need to have the device close to the access point to be able to press the WPS simultaneously, since the mobile application can be used to “soft initiate” the WPS action on the access point.

Ongoing maintenance of the client, in terms of its client association, data usage and other connectivity statistics are things that can be handled by the gateway and information provided to the mobile application for display.

While we are on the topic of clients, it is important to note that any of the solutions relative to coverage, security, performance and such need a clear understanding of the nature of the client. This cannot be achieved through simple device query, and in many cases may need to be deciphered based on the data traffic and other parameters. The process of identifying a device to a very fine degree of granularity is the process of “Device Fingerprinting”.



Figure 38 - Device Fingerprinting Essential with the Explosion of IoT devices

Device fingerprinting uses machine learning techniques and the learning process gets better as more and more clients across multiple homes are onboarded and recognized. This fundamental process comes in handy in correlating performance and security related functions.

As an example, consider a customer complaint about poor Wi-Fi performance from a specific Wi-Fi client device. If the system is able to figure out more detailed, granular information about the device (like its hardware capability or the lack thereof, software versions and such), a potential explanation could be provided explaining why the lack of performance is not due to any service provider-supplied device (such as the access point itself) but more to do with the client itself.

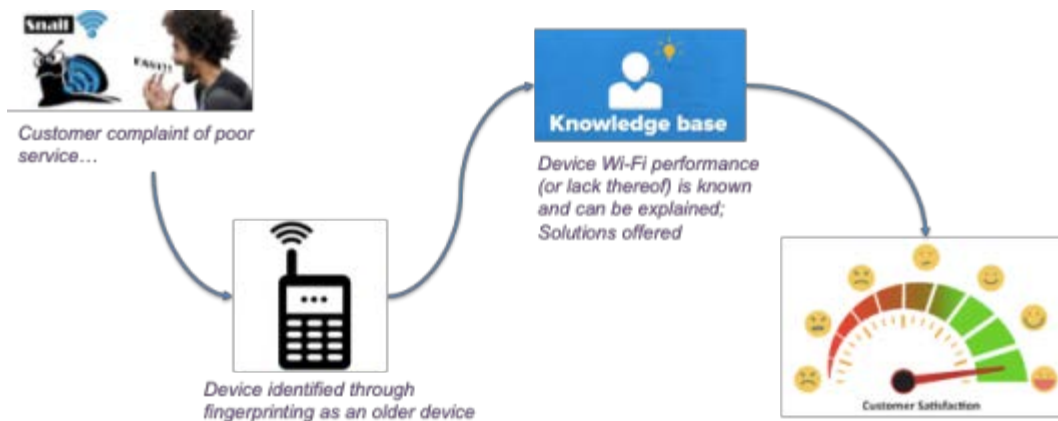


Figure 39 - Correlation of Poor Performance to Client Characteristics

Device fingerprinting is becoming an essential need in Wi-Fi management systems, especially when the management of IoT devices or security solutions are part of the service offering.

10. Coverage & Performance at Range

As we already noted, Coverage that goes hand in hand with “Performance at Range” is the touchiest subject given the customer expectations of ensuring Wi-Fi speeds are equal to the access speeds. We also noted that even with the upcoming 802.11ax specifications, a typical home needs to have multiple access points, and doing so, should solve the problems of auto-configuration and also allowing the clients to

associate with the right AP through steering. In this section, we will examine one typical implementation that combines the use of standards and other techniques to solve the problem.

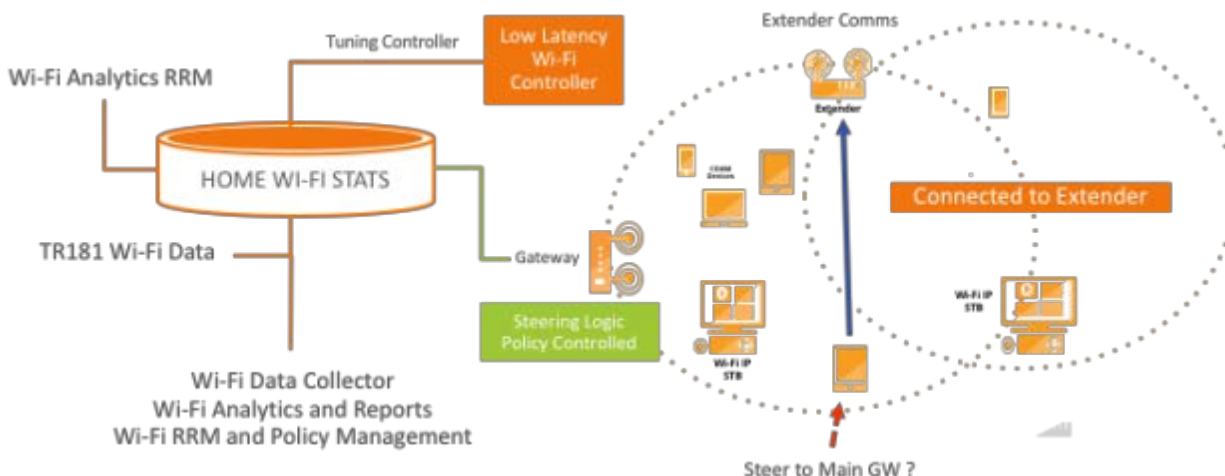


Figure 40 - Architecture of a Typical Cloud-Assisted Wi-Fi Management Solution

The above figure depicts the architecture of a typical cloud-assisted Wi-Fi management solution. There are three major software components: The controller & communication agent in the gateway, the communication agent in the secondary access points (extenders), and the cloud software modules. Let us examine the details of each of these entities and the role that they would play in enhancing the coverage. A typical architecture for the controller (along with the communication agent) in the gateway (or the main access point) is shown in the diagram below:

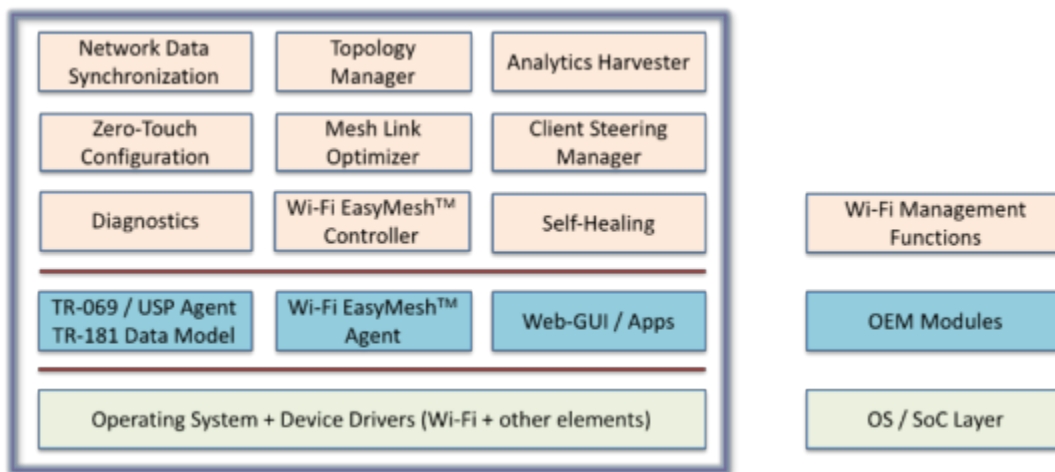


Figure 41 - Gateway Controller for Wi-Fi Management

The architecture depicted accounts for the controller software to be portable across multiple vendors. The OS/SoC layer typically brings the operating system and the associated device drivers, critical of which is the Wi-Fi device driver.

Typically, the OEM will have additional modules that are expected of any implementation given that these modules have a strong tie to the underlying operating system and device drivers. However, many of

these should have standard north-bound APIs (and equivalent data elements) that allow for Value-added functions (like Wi-Fi management functions) to be developed on top of them. In this case, the TR-069 agent (or the Broadband-forum defined User Service Platform (USP) agent) acts as the module that is standards-based implementations that would talk to a TR-069 Auto Configuration Server (ACS) and interoperate with implementations of the ACS from multiple vendors. In addition, inter-AP communications should migrate to the use of Wi-Fi EasyMesh agent implantation that will allow for APs from different vendors to interoperate. The north-bound API for such an Agent is also well-defined, allowing for the implementation of the Wi-Fi EasyMesh controller, which can (and usually must) be enhanced to allow for intelligent control of clients inside the home.

The upper layer of software is where the real intelligence is, and these modules serve to implement the sophisticated algorithms that are used to make informed decisions to handle the Wi-Fi clients in the home. Typical functionality handled by these modules are listed below:

10.1. Multi-AP for Network Extender Communication

the gateway-based manager should be standards compliant. It manages the communication between all APs in the home and utilizes the Multi-AP 1905 protocol to discover and configure new access points on the home network. The onboarding process can be accomplished through Wi-Fi or wired Ethernet. The (Multi-AP) MAP Controller or a MAP Agent sends a 1905 Topology query message to the network and start with the MAP controller discovery. The auto configuration process starts according to MAP specifications.

10.2. Zero Touch Provisioning of New Wi-Fi Extenders

Once an access point is discovered, the zero-touch configuration manager provides zero-touch provisioning of the new Wi-Fi extender by synchronizing the extender's Wi-Fi SSID and password with the gateway's configuration to create a single Wi-Fi network in the home.



Figure 42 - Zero-Touch Configuration

10.3. Cloud Assisted and Policy-based Management

The gateway Wi-Fi manager resides on the residential gateway and communicates with extenders using MAP. The controller is configured remotely via TR-069 but makes autonomous policy decisions locally – because the policies that it enforces are local to the home and do not require remote support. Therefore,

the policy decisions and actions are low latency and available even if there is a disruption of the broadband service.



Figure 43 - An Example Implementation of Wi-Fi Management Function

Policy events, such as steering actions or channel changes, are logged such that they may be retrieved and used as telemetry to troubleshoot problems and as feedback to optimize policy configuration in the cloud server.

10.4. Roaming and AP Steering

Many mobile devices exhibit the ‘sticky client problem’, where they maintain an association with an AP until the last gasp of connectivity is available – regardless of whether there is a better candidate AP to associate with.

The client-steering manager solves the sticky client problem by evaluating link quality to detect this condition and then forcing a client device with a low link quality to move to an AP with a stronger signal.

Clients are also steered to alternate APs to reduce contention. For example, when an AP is highly utilized by an individual client, other clients associated to that AP may be moved to another to balance the network and optimize performance.

10.5. Band Steering

The client-steering manager provides band steering to solve performance, throughput, and quality problems caused by Wi-Fi congestion by moving impacted clients to a different band. This feature is relevant even in homes with only one access point since most modern residential gateways are dual band.

10.6. BSS Steering

BSS steering enables service providers to establish separate BSS to separate the Wi-Fi policies and authentication credentials e.g. for community hotspots, guest Wi-Fi and STB streams. This is where the

standards such as 802.11k and 802.11v come into play, as they make the fast transition of active clients possible without any additional incurred delays.

10.7. Dynamic Channel Selection

Some access points select the best channel on boot, some perform regular scans to determine if conditions have changed and there is a clearer channel that should be used. However, even those devices often don't scan if they are busy – when a channel change may be needed the most.

The Self-Healing module scans regularly, regardless of how busy it is. A typical scan requires using the antennae for 10 ms, minimizing disruption with the benefit that a clearer channel may be available and switched.

10.8. Mesh Link Optimization

The Mesh-Link Optimization module performs policy-based management of mesh topology. Rather than leaving backhaul selection to individual devices making ad-hoc decisions, the EasyMesh Manager uses its knowledge of the topology, device capabilities, and loading of the network to select backhaul links and thereby optimize traffic flow across the mesh.

10.9. Airtime Management

Airtime fairness is well known to prevent clients that are slow or have poor connectivity from monopolizing airtime and starving other clients. However, airtime configuration can also be used to reserve airtime for certain devices, which establishes a minimum quality of service. For example, airtime can be reserved for Wi-Fi set-top boxes to ensure that managed video has sufficient airtime to ensure a quality user experience.

11. Subscriber Visibility & Device Management

From the perspective of the subscriber, there is an increasing need for having a clear view of the devices in the home and having some level of control over the devices to the extent that they are entitled to handle. Increasingly, functions like the ability to perform speed tests, to be able to control Internet access as part of parental control or just being able to draw all members of the family to dinnertime and being able to understand bandwidth usage patterns.

An ideal application will provide sufficient detail to the subscriber and even if not used frequently, can serve to be the first place that a subscriber would go before picking up the phone to call the customer. We depict here a few sample screens from a typical application for Wi-Fi management that includes many of the features described above.



Figure 44 - Mobile Application Screens to Highlight Data Usage

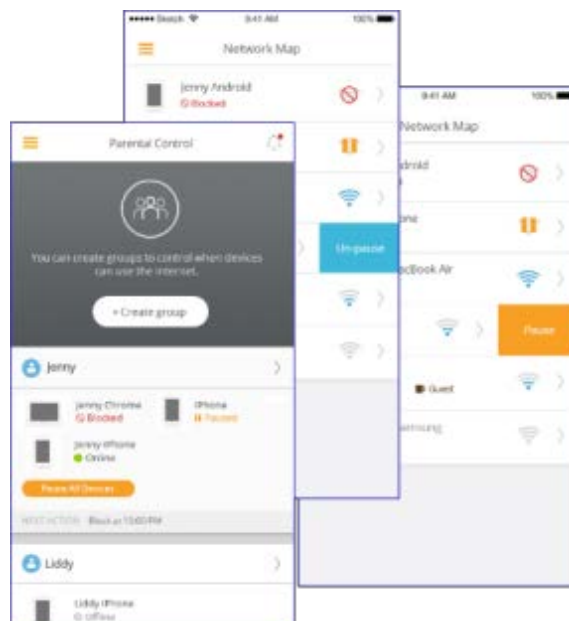


Figure 45 - Mobile Application Screens for Parental Control and Device Access Restrictions

Conclusion

Wireless Technology, given the inherent nature of physics and being intangible to start with, creates an enormous challenge for the service providers. The accelerated adoption of the technology by the consumer, and the concomitant pace with which innovations have unfolded in this space, while exhilarating for the enthusiast, is also one that is fraught with challenges from the perspective of a service provider. With consumer expectations of ease of use, performance, and visibility evolving rapidly, there is a gap between such expectations and what can reliably be addressed by the service provider in a meaningful way. Disruptive players are ready to step in, especially given that control of the home for service and other content delivery is highly valued. Additionally, retail players are out to make most of the situation with well-designed products that at least superficially address coverage and connectivity.

Challenging as the situation may be, we also have tools at our disposal, and there are more and more devices that are implementing standards that allow interoperability across multiple vendors. With standards, availability of tools that leverage AI and ML techniques, and the general paradigm shift to cloud-based technologies, the service provider can deploy intelligent software solutions and wrest control of the Wi-Fi management challenge.

Abbreviations

AC	Alternating current
ACS	Auto Configuration Server (In the context of TR-069 protocol)
AI	Artificial Intelligence
AP	Access Point
API	Application Programming Interface
AX	802.11ax (AX for short)
bps	bits per second
BSS	Basic Service Set
BTM	BSS Transition Management
CAPEX	Capital Expense
COAM	customer owned and managed
DBC	dynamic bonding change
Gbps	Gigabits Per Second
GHz	Gigahertz
GW	Gateway
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
IP	Internet Protocol
ISBE	International Society of Broadband Experts
MAP	Multi Access Point Protocol
Mbps	Megabits per second
MCS	Media & Communications Systems
MDU	Multiple Dwelling Unit
ML	Machine Learning
MoCA	Multimedia Over Coax Alliance
ms	Millisecond(s)

MTA	Multimedia Terminal Adapter
OEM	Original Equipment Manufacturer
ONU	Optical Network Unit
OPEX	Operating expense(s)
OS	Operating System
OTT	Over-the-Top; refers to service overlay on top of data services
QAM	Quadrature Amplitude Modulation
QR	Quick Response
RRM	Radio Resource Management
SCTE	Society of Cable Telecommunications Engineers
SLA	Service Level Agreement
SoC	Silicon on Chip
SON	Self-Optimization Networks
SP	Service Provider
SSID	Service set identifier
STB	Set-top Box
UHD	Ultra High Definition
USP	User Service Platform
WAN	Wide Area Network
WFA	Wi-Fi Alliance
WPS	Wi-Fi Protected Setup

Bibliography & References

Wi-Fi CERTIFIED EasyMesh™ Technology Overview; <https://www.wi-fi.org/discover-wi-fi/wi-fi-easymesh>

Multi-AP Specification v1.0; <https://www.wi-fi.org/discover-wi-fi/wi-fi-easymesh>

Wi-Fi CERTIFIED Easy Connect™ Technology Overview; <https://www.wi-fi.org/discover-wi-fi/wi-fi-easy-connect>

Device Provisioning Protocol Specification v1.0; <https://www.wi-fi.org/discover-wi-fi/wi-fi-easy-connect>

802.11ax and other WLAN related documents; <http://www.ieee802.org/11/>

Impact of Access Environment in Cable's Digital Coherent System – Coexistence and Full Duplex Coherent Optics

A Technical Paper prepared for SCTE•ISBE by

Zhensheng (Steve) Jia, Ph.D.

Distinguished Technologist

CableLabs

858 Coal Creek Circle, Louisville, Colorado 80027

303 661 3364

s.jia@cablelabs.com

L. Alberto Campos, Ph.D.

Fellow

CableLabs

858 Coal Creek Circle, Louisville, Colorado 80027

303 661 3377

a.campos@cablelabs.com

Mu Xu, Ph.D., CableLabs

Haipeng Zhang, Ph.D., CableLabs

Jing Wang, Ph.D., CableLabs

Curtis Knittle, Ph.D., CableLabs

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	5
1. Coherent Optics for Access Applications	5
2. Deployment Scenarios of Coherent Optics in Cable.....	7
3. Coexistence Testing with Legacy Optical Channels	8
3.1. DWDM Components	8
3.2. Coexistence Using 8-port Mux (Analog + Coherent).....	9
3.3. Coexistence Using 16-port Mux (Analog + OOK + Coherent)	11
4. Full Duplex Coherent Optics	14
4.1. The Need for Single Fiber Connections	14
4.2. The Existing Approach	14
4.3. Full Duplex Coherent Optics Approach.....	15
4.4. How Does It Work in a Cable?	16
4.5. Testing Setup and Results	16
Conclusion.....	19
Acknowledgement.....	19
Abbreviations	19
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 – Optical Technology Evolution	6
Figure 2 – Coherent Optics for Cable Access Applications	7
Figure 3 – Optical Spectra of Mux and Demux	9
Figure 4 - Experimental Setup for Coexistence Evaluation Using 8-port Mux/DeMux	9
Figure 5 – Optical Spectra for both Analog and Coherent Signals	10
Figure 6 – Experimental Results of Coexistence with 8-port Mux/DeMux.....	11
Figure 7 – Optical Spectra for Analog, NRZ, and Coherent Signals.....	12
Figure 8 – Experimental Results	13
Figure 9 – Today's Single Fiber Use Percentage	14
Figure 10 – Dual-Fiber Approach.....	15
Figure 11 – Single-Fiber Approach with Two Lasers	15
Figure 12 – Full-Duplex Single-Fiber Approach.....	16
Figure 13 – Testing Case I: 50km, 80km, Attunator at Rx Sides	17
Figure 14 – Testing Case I: Results.....	17
Figure 15 – Testing Case II: Variable Backreflector	18
Figure 16 – Testing Case II: Results (RP: Reflected Power).....	18

List of Tables

Title	Page Number
Table 1 – Launched & Received Power, and Gain/Attenuation of Optical Devices	10
Table 2 – Optical Transmitted Power for Analog, OOK, and Coherent Signals	13

Introduction

Cable access networks have been undergoing significant technology and architecture changes driven by the ever-increasing residential data service growth rate and an increasing number of services types being supported, such as business services and cellular connectivity. Digital fiber technologies and distributed access architecture 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. The combination of the natural evolution of coherent optics technology, along with this increasing demand for capacity and the unique features of a cable-specific fiber access environment with only a few fibers available for a 500-household passed serving area, prompted the evaluation of coherent optics as an alternative for a long-term fiber access connectivity solution in next-generation cable access networks.

During its 2017 Winter Conference, CableLabs® announced the launch of the point-to-point (P2P) Coherent Optics specification project. The project looks into the evolution of cable's optical access network, addresses its fiber shortage challenge, and re-designs digital coherent system from long-haul and metro solutions to the access network applications. This specification allows operators to best leverage the existing fiber infrastructure to withstand the exponential growth in capacity and services for residential and business subscribers while keeping cost down as much as possible.

When cable operators look to deploy coherent optics into their access networks, they are typically faced with two options: deploy coherent optics on the existing 10G system or build a new coherent-only connection. The ideal network for deploying such coherent systems would be a green field deployment on fibers without any compensation devices such as dispersion compensation modules (DCM) and other wavelength channels. However, in practice, to make the upgrade cost-effective, only one or a few channels may be upgraded in many brown field installations, depending on capacity demand. That means many of these networks that are deployed already with WDM analog DOCSIS technology and/or 10G on-off keying (OOK) services will coexist with a coherent system to support a hybrid scenario over the same fiber transmission. Such a hybrid configuration needs to be studied, especially the cross-phase modulation (XPM) impairment in the fiber nonlinear regime, to provide this option for operators to effectively support 100G on their existing networks. In this work, we fill the gap by presenting extensive experimental verifications under various coexistence scenarios and provide operational and deployment guidance for such use cases.

Additionally, according to a recent operators' survey, 20 percent of existing cable access networks use a single-fiber topology. This means that downstream and upstream transmission to nodes takes place on a single strand of fiber. This number is expected to grow further in the near future. Therefore, bidirectional transmission is needed for coherent signals to support single-fiber topologies and to facilitate the business use and redundancy of optical links. CableLabs' Full Duplex Coherent Optics (FDCO) proposal and the experiments that demonstrate simultaneous bi-directional transmission over single fiber and single wavelength are described. This paper shows how FDCO effectively doubles fiber capacity in a coherent optics-based fiber distribution network. The major impairment in the FDCO system is optical return loss (ORL) or optical reflections including all discrete reflections (Fresnel) and continuous reflections (backscatter). In this paper, the impact of ORL for FDCO is also analyzed and quantified for various configurations.

Content

1. Coherent Optics for Access Applications

Coherent optics initially received significant research interest in the 1980s because of high receiver sensitivity through coherent amplification by a local oscillator, but its use in commercial systems has been hindered by the additional complexity of active phase and polarization tracking. In the meantime, the emergence of a cost-effective erbium-doped fiber amplifier (EDFA) as an optical pre-amplifier reduced the urgency to commercialize coherent detection, because EDFAs and wavelength-division multiplexing (WDM) extended the reach and capacity as shown in Figure 1. Traffic demand, combined with the requirement to reduce cost per bit per Hz, or spectral efficiency increases, as well as advancements in CMOS processing nodes and powerful digital signal processing (DSP), led to the renaissance of coherent optics technology. 2018 is the 10th anniversary that digital coherent optical technology was officially reintroduced to the world.

Commercial coherent optical technology was first introduced in long haul applications to overcome fiber impairments that required complex compensation techniques when using direct detection receivers. The first-generation coherent optical systems are based on single-carrier polarization division multiplexed quadrature phase shift keying (PDM-QPSK) modulation format and the achieved spectral efficiency (SE) is 2 bit/s/Hz over conventional 50-GHz optical grid, thus the system capacity has been increased to around 10 Tb/s in the fiber C-band transmission window. Leveraging further development of CMOS processing, reduction in design complexity, and price decreases on opto-electro components, coherent solutions have moved from long haul to metro and access networks. This migration model has been demonstrated in the optical industry before: the DWDM system technology started in the long haul and then migrated to metro and edge access; forward error correction (FEC) encoding and decoding follows the same pattern. Benefiting from initial long-haul technology development, coherent optics for access networks will be the next natural progression. Current development of application-specific integrated circuits (ASICs) for DSP chips, and corresponding optical modules head in the two directions shown in Figure 1. One path is to have a programmable and comprehensive coherent DSP which is capable of processing data rates from 100G to 600G per single wavelength, with the support of higher modulation formats like 32/64-QAM and high net coding gain (NCE) FEC. The second path is the development of reducing the power consumption and thereby meeting the size and cost requirements for access applications, which is the focus of this work.

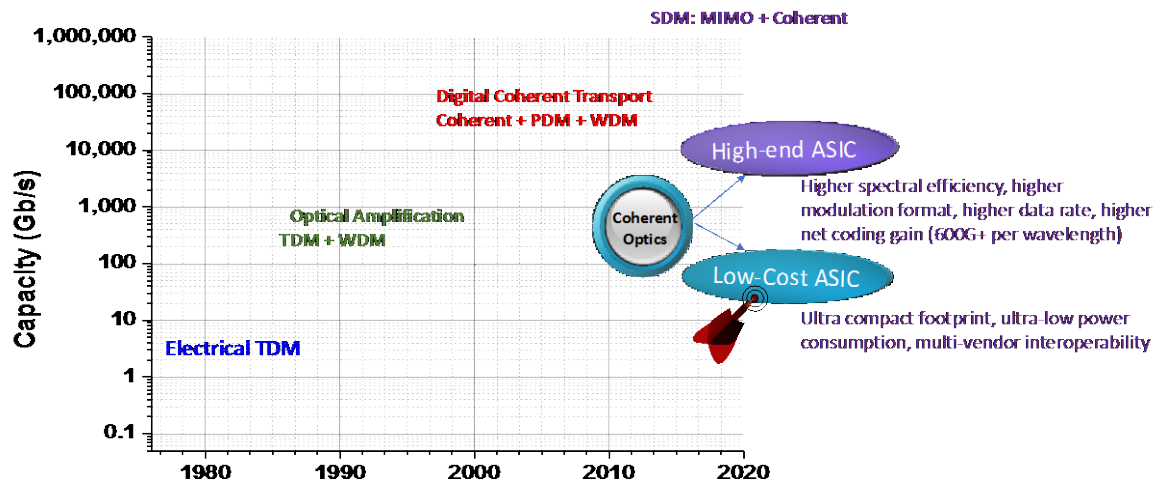


Figure 1 – Optical Technology Evolution

Coherent detection for access networks enables the superior receiver sensitivity that allows for extended power budget, and the high spectral efficiency enables dense WDM (DWDM). Moreover, the use of high-order modulation formats enables efficiently utilizing the spectral resource and benefiting future-proof network upgrades. In the cable access environment, coherent optics allows operators to best leverage the existing fiber infrastructure to deliver vastly increased capacity with even longer distances. However, the coherent technology in a long-haul optical system utilizes best-in-class discrete photonic and electronic components, the latest DAC/ADC and DSP ASIC based on the most recent CMOS processing node. The coherent pluggable modules for metro solution have evolved from CFP to CFP2 form factor for smaller footprint, lower cost, and lower power dissipation. However, it is still over-engineered, too expensive, too power hungry, and not interoperable. The access network is a totally different environment as compared to long haul and metro. It may need hardened solution for remote site locations, where temperature is not controlled. Another important factor to consider is standardization and interoperability. Standardization in the optical community is driven mainly by short-reach metro/aggregation applications, where optical performance is not a differentiator. Interoperability and a robust vendor ecosystem are therefore keys to providing a low-cost solution using coherent optics.

In 2017, CableLabs recognized the benefit of coherent optics and announced the launch of the point-to-point (P2P) Coherent Optics that allows the cable industry to support the growing requirements of broadband access as the industry evolves toward Node+0 architectures, and the volume of optical connections to intelligent nodes increases substantially. 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. 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 bring it to market quickly [3] [4].

Industry organization bodies such as the Optical Internetworking Forum (OIF) and IEEE are working on short-reach coherent optical standardizations. The OIF is defining a coherent standard for DWDM interfaces in DCI applications with reaches up to 120 km with multi-vendor interoperability, and IEEE is considering coherent optics for unamplified applications beyond 10 km distances. All of this standardization activity reinforces the view of coherent optics moving to shorter reach and high-volume applications.

2. Deployment Scenarios of Coherent Optics in Cable

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 edge-to-edge connectivity to the desired end-point as shown in Figure 2. Following capacity growth trends, it is obvious that initially the aggregation use cases are going to outnumber the direct edge-to-edge connectivity use cases. The aggregation use case supports any Distributed Access Architecture (DAA), including Remote PHY, Remote MAC-PHY, and Remote optical line terminal (OLT) architectures.

In the 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 a secondary Hub 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 Aggregation Node ranges from 20 to 80 km, and the distance from the Aggregation Node to each end point is less than 3 km. Each primary Hub can support multiple (~60) Aggregation Nodes for different services.

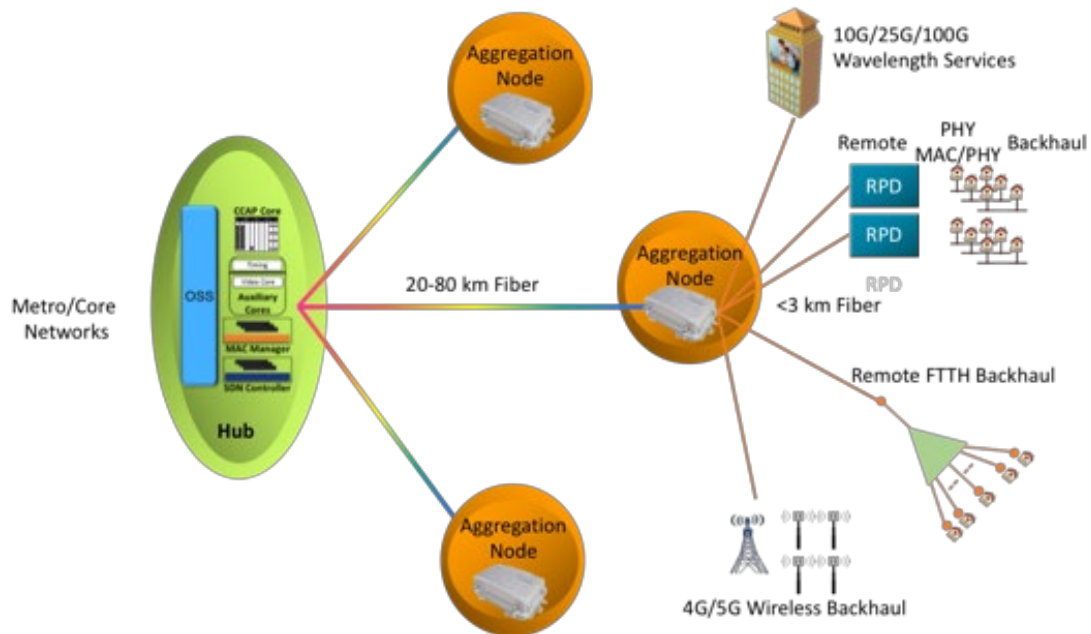


Figure 2 – Coherent Optics for Cable Access Applications

Commercial services have been a rapidly growing and high revenue segment in cable. Business connectivity, cellular backhaul and wireless access point connectivity, including 5G connectivity, are expected to play a bigger role in cable's future service portfolio [1]. These services demand very high bandwidth as well as robustness and flexibility for supporting a diversity of service levels. Coherent optics is a technology that can easily address the service requirements of this market segment, which is also shown in Figure 2. Direct wavelength services can overlay aggregation connections with 10G/25G intensity modulated signals or 100G coherent signals. This is lambda/wavelength deep for edge to edge services. In this use case, the coherent optic links are terminated at the edge customer and WDM

multiplexer/demultiplexer at the HE/Hub is used for aggregating multiple P2P optical links onto a fiber. The WDM systems can be a hybrid system with a mix of data rate and modulation formats [2].

3. Coexistence Testing with Legacy Optical Channels

The commercial coherent 100G transmission systems are showing excellent receiver sensitivity, robustness, and tolerance for channel impairments such as CD and PMD. Therefore, the ideal network for deploying such coherent systems would be a green field deployment on fibers without any compensation devices such as dispersion compensation modules (DCM) and other wavelength channels, which is called a coherent-only implementation. However, in practice there are many brown field installations, meaning many of these networks are deployed and have several WDM analog DOCSIS and/or 10G OOK (Ethernet over fiber or PON) services running over the existing fiber already. The expectation from cable operators has been that adding additional 100G coherent services by using free channels in the WDM grid is preferred without impacting the existing services. This will essentially create a hybrid 10G/100G network with multiple services coexistence. But the fact is that 10G signals based on analog amplitude modulation (AM) or OOK have a much higher power density than coherent 100G, causing them to have a much greater impact on the refractive index for nonlinear effects such as cross phase modulation (XPM) and four-wave mixing (FWM). Additionally, crosstalk penalties in ITU-T grid networks with mixed rates lead to system degradation due to optical Mux/DeMux in-band residual power or non-uniform channel grid allocation in DWDM systems.

To provide an option that enables network operators to effectively support 100G on their existing networks infrastructure, such as optical amplifier and Mux/DeMux, CableLabs took the initiative and has done experimental verification to quantitatively explore the performance challenges in such coexistence applications. In the previous effort [5], because of the limited availability of analog optical channels, three copropagating analog DOCSIS channels were tested along with single coherent channels. The experimental results show that coherent optics transmissions are robust, even in close proximity to much stronger analog optical carriers, and coherent optical carriers impose negligible impact on analog optical carriers. To further test the transmission performance of full-loading coexisting systems, the following experiments have been conducted with longer transmission distances.

3.1. DWDM Components

Three different kinds of optical multiplexors/demultiplexors have been evaluated in the testing. Figure 3 shows their optical spectra for two wavelength channels; they are 8-port thin-film filters (TFFs), and 40-48-port array waveguide grating (AWG) based optical multiplexors/de-multiplexors. TFFs use concatenated interference filters, each of which is fabricated with a different set of dielectric coatings designed to pass a single wavelength. As shown in Figure 3, TFFs have a better optical performance in terms of flatter passband ripple and higher isolation in neighboring channels. They work well for low channel counts, especially for analog WDM systems, but have challenges at higher channel counts and narrower spacing because they need several hundred layers of coating, which requires stricter error control. In contrast to TFFs, AWG devices use a parallel multiplexing approach that is based on planar waveguide technology. The key advantage of AWGs over TFFs is that their cost is not dependent on wavelength count making them extremely cost-effective for high channel count applications. The existing long-haul coherent DWDM systems are typically using AWG for Mux and Demux. In our experimental setup, the insertion loss is ~1.5 dB for TFFs and ~3.5 dB for 40-port AWG.




Figure 3 – Optical Spectra of Mux and Demux

3.2. Coexistence Using 8-port Mux (Analog + Coherent)




Figure 4 - Experimental Setup for Coexistence Evaluation Using 8-port Mux/DeMux

Figure 4 shows the experimental setup of the first case. All eight analog DOCSIS channels (up to 1.2 GHz) are multiplexed through a TFF based 8-port Mux, while four coherent channels are multiplexed via the 40-port AWG based Mux with 100-GHz optical grid spacing. The analog channels are selected in order to minimize nonlinear interference with each other, and the selection of coherent signals is expected to exhibit the worst coexistence condition. These two kinds of signals are combined via an optical combiner. The channel labels in the diagram correspond to the standard ITU-T wavelength grid. The purpose of selecting a nonuniform analog wavelength plan is to mitigate fiber nonlinear impairments, especially four-wave mixing (FWM). In the meantime, creating the worst nonlinear crosstalk impairments is the criteria for selecting coherent wavelength plans. The combined signals are then amplified by an EDFA that is designed for long-distance analog signal amplification. The maximum output power is

about 18 dBm. These amplified signals then transmitted over 80 km single mode fiber (SMF) and are split to reach the corresponding optical DeMux for analog and coherent channels respectively. The launched power of two kinds of channels and the gain/attenuation of optical devices along the optical links are shown in Table 1. Around 10 dB power difference is set between coherent and analog channels.

Table 1 – Launched & Received Power, and Gain/Attenuation of Optical Devices

Signal Type	Tx Output Power (dBm)	Mux Loss (dB)	Coupler Loss (dB)	EDFA Gain (dB)	Fiber Attenuation (dB)	Splitter Loss (dB)	DeMux Loss (dB)	Received Power (dBm)
Analog	9.5	-1.5	-3	+2.5	-5.5 for 26.1 km;	-3	-1.5	-2.5 for 26.1 km; -6.5 for 51.8 km
Coherent	-1	-3.5	-3	+5	-9.5 for 51.8 km	-3	-3.5	-14.4 for 26.1 km; -18.4 for 51.8 km

Figure 5 shows the optical spectra of all signals before and after optical amplification (a) and before and after optical fiber transmission (b). CH 23, 25, 26, and 27 are coherent channels with wider spectra and much lower power compared to eight analog channels.




Figure 5 – Optical Spectra for both Analog and Coherent Signals

In this system setup, the transmission performance of both analog and coherent channels is shown in Figure 6. Negligible penalty is observed after 26.3 km or 51.8 km fiber transmission with the impact of coherent signals on analog channel CH 52 as shown in Figure 6 (a) with 26.3 km transmission. Other analog channels show similar performance when we compare the transmission condition (with or without a coherent channel over the same fiber). In the case of the impact of analog channels on coherent channels, minor BER difference is observed for 8-QAM and 16-QAM based 200 Gbps channels with 0 dBm transmitter output power. When compared with back to back coherent signal sensitivity, less than 0.5dB power penalty is found for the transmission and analog overlay using analog EDFA amplification and the same fiber.

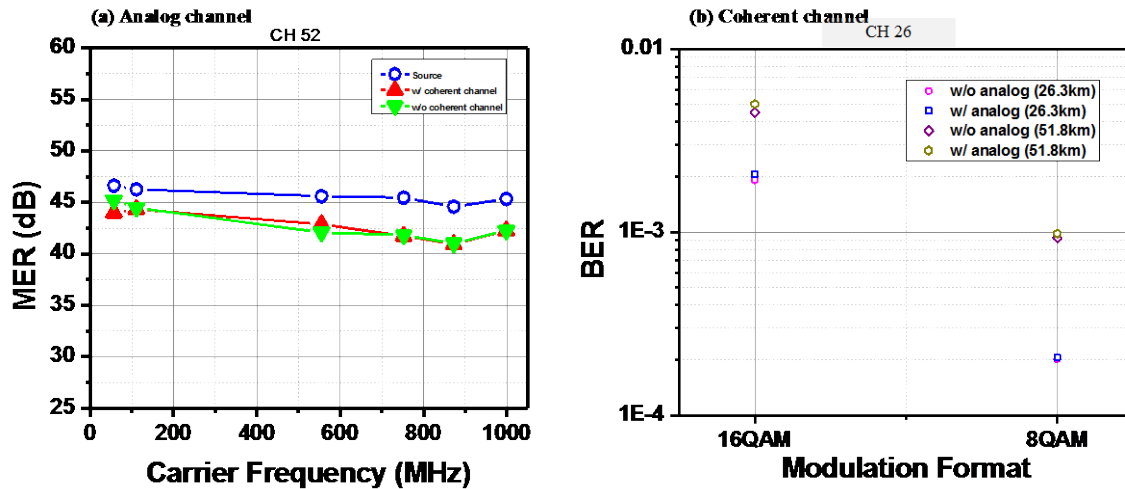


Figure 6 – Experimental Results of Coexistence with 8-port Mux/DeMux

3.3. Coexistence Using 16-port Mux (Analog + OOK + Coherent)

Next, the more complex coexisting setup was established with eight analog channels, two coherent 100G PM-QPSK channels (CFP2-DCO form factor), two coherent 400G channels, and two 10G NRZ channels. This coexistence hybrid scheme includes all the major modulation formats, and services under different data rates/ baud rates. A pair of 16-channel TFT based wavelength division multiplexers are used for channel multiplexing and demultiplexing. The optical spectra of these multi-channel coexistences are shown in Figure 7, with analog, coherent 100G, coherent 400G, and 10G NRZ marked with red, blue, green, and purple respectively.




Figure 7 – Optical Spectra for Analog, NRZ, and Coherent Signals

The corresponding input power level (right after the transmitter) is shown in Table 2, based on typical operational conditions for different detection schemes. The power levels are measured at the output port of each transmitter before entering the WDM Mux. Among them, the powers of the analog channels are set to around 9.5 dBm while 56GBaud 400G coherent channels have the power set to ~3 dBm. To improve the spectral efficiency and confine the optical power within each WDM channel, the coherent signals are shaped by square root raised cosine filters.

Table 2 – Optical Transmitted Power for Analog, OOK, and Coherent Signals

Application Scenarios	Channel Index	Input Power (dBm)
DOCSIS Analog	21	9.64
	24	9.48
	28	9.43
	33	9.64
	39	9.48
	52	9.11
	60	9.01
	62	8.92
400G Coherent	22	2.68
	61	3.15
CFP2 Coherent	36	0.08
	44	-0.14
10G NRZ	26	-0.89
	57	-0.75




Figure 8 – Experimental Results

The performance difference is insignificant for analog channels compared to coexistence using the 8-port Mux. In the case of the coherent channels, Figure 8 shows the normalized power requirements for 200G, 300G, and 400G data rates, at back-to-back and 50 km transmission, with and without analog plus NRZ channels. Less than 0.6 dB power penalty is observed in the coexistence scenarios compared to the non-coexistence case.

In summary, three main observations were found in the coexistence measurement experiments:

- Both coherent and analog/NRZ signals work well in the coexistence application with ~0.6 dB maximum power penalty in the case of 100 GHz channel spacing and 50/80 km fiber transmission distances for different nonlinearity tolerance scenarios.

- The legacy components/devices for analog systems are working well for coherent signals multiplexing and amplification, including analog EDFA, optical Mux and DeMux. Coherent signals show strong robustness when they are deployed in traditional analog DWDM systems.
- However, the conventional AWG-based optical Mux and DeMux configuration, which is typically used for coherent channels, is not good for conventional analog channels.

4. Full Duplex Coherent Optics

4.1. The Need for Single Fiber Connections

According to a recent operators' survey, 20 percent of existing cable access networks use a single-fiber topology as shown in Figure 9. This means that downstream and upstream transmission to nodes takes place over a single strand of fiber. It is estimated that over the next several years, this number will grow further. Therefore, to control the cost and fully utilize the existing infrastructure, bidirectional transmission over a single fiber is needed for coherent signals to support single-fiber topologies and to facilitate the redundancy of optical links.




Figure 9 – Today's Single Fiber Use Percentage

4.2. The Existing Approach

Today, achieving bidirectional transmission in an optical domain with a single laser requires two fibers. This is the standard practice using today's coherent optical technology. One laser in a transceiver performs two functions:

- as the optical signal source in the transmitter
- as the reference local oscillator signal in the receiver

Because of the use of the same wavelength from the same laser, a second fiber must be available for the other direction—one fiber for downstream and a second fiber for upstream as shown in Figure 10.




Figure 10 – Dual-Fiber Approach

The second typical approach is to use a single fiber but transmit at different frequencies or wavelengths, similar to the upstream and downstream spectrum split that we implement in our HFC networks. To accomplish this frequency/wavelength multiplexing approach, two lasers operating at different wavelengths are needed, as shown in Figure 11. Wavelength multiplexers and demultiplexers following a wavelength management and allocation strategy are needed to combine these different wavelengths over the same fiber. The second laser ends up costing a lot more than money—increasing power consumption, operational complexity, and transceiver footprint.

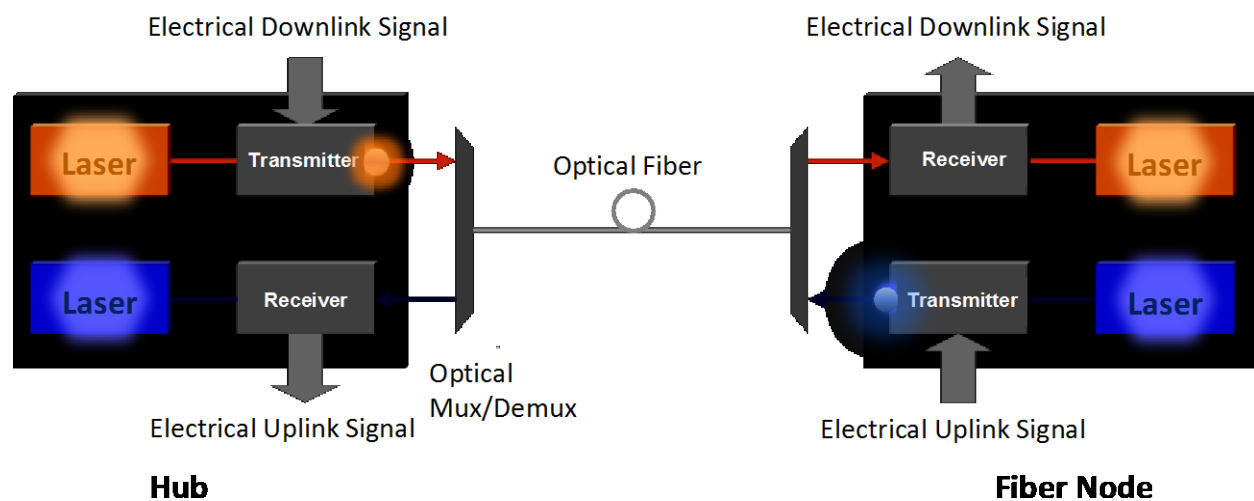


Figure 11 – Single-Fiber Approach with Two Lasers

4.3. Full Duplex Coherent Optics Approach

CableLabs proposes an alternative method achieving full duplex coherent optics. We leverage two optical circulators on each end in a special configuration. The circulator is a low-cost, passive, but directional

device—much like a traffic roundabout for cars, however this device is used for rerouting the optical path in different directions. Instead of using two fibers, a single fiber is connected for bidirectional transmission; most importantly, instead of using two lasers, a single laser is employed for single-fiber coherent systems. The scheme is shown in Figure 12.

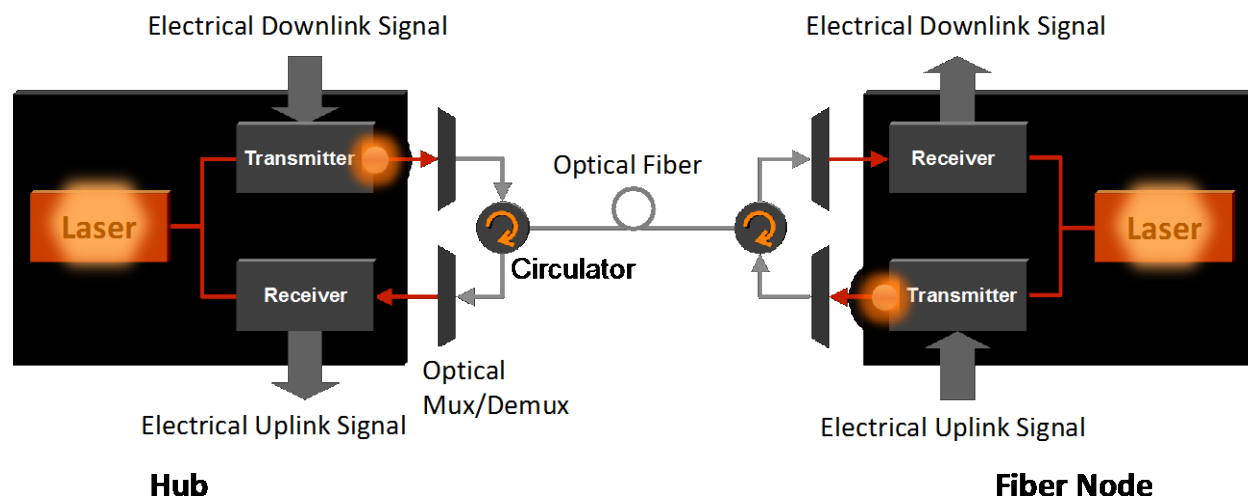


Figure 12 – Full-Duplex Single-Fiber Approach

4.4. How Does It Work in a Cable?

Many scenarios in cable focus on the access environment with limited transmission distances. Unlike backbone and metropolitan coherent optical networks, access networks don't require multiple directional optical amplifiers in cascade. When dealing with coherent signals, we have much higher Optical Signal to Noise Ratio (OSNR) sensitivity and higher tolerance to the impairments from the spontaneous Rayleigh backscattering (continuous reflection) and Fresnel reflection (discrete reflections), than intensity-modulated systems. The majority of existing analog optics employs angle-polished connector (APC), which provides excellent mitigation for return loss from Multiple-Path Interference (MPI) or jumper cable/optical distribution panels/fusion or mechanical splices. In addition, the threshold of the Stimulated Brillouin scattering (SBS) nonlinear effect is suppressed because of the nature of phase-modulated signals on reducing optical carrier power and increasing the effective linewidth. With this new dimension of direction-division multiplexing (DDM) in the optical domain, any coherent wavelength can be used twice, once in each direction, thus doubling the whole fiber system capacity. This full-duplex implementation is not wavelength-selective. It works for both short and long wavelengths, and it would cover not only the entire C-Band but, with different optical sources, the entire fiber spectrum.

4.5. Testing Setup and Results

Figure 13 shows the first test setup with the variable attenuator on the receiver side in each direction. This is the typical operational case to measure the power penalty with and without full duplex operation, where the power of both received signal and returned impairment is attenuated. The penalty comes from the Rayleigh backscattering and Fresnel reflection along the whole link.




Figure 13 – Testing Case I: 50km, 80km, Attunator at Rx Sides

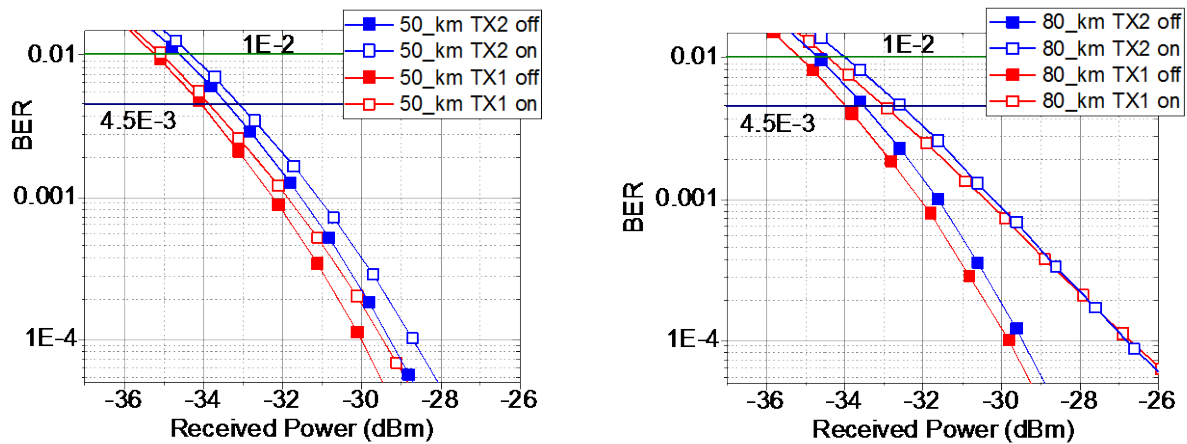


Figure 14 – Testing Case I: Results

The reflected power is measured as -34.7 dBm, as the output power of the transmitter (TX1 or TX2) is set to 0 dBm. Figure 14 shows the results for 50 and 80 km transmission distances. Around 0.5 dB and 1 dB power penalties are observed for 50 km and 80 km transmission, respectively, when compared with full duplex operation with single direction operation.

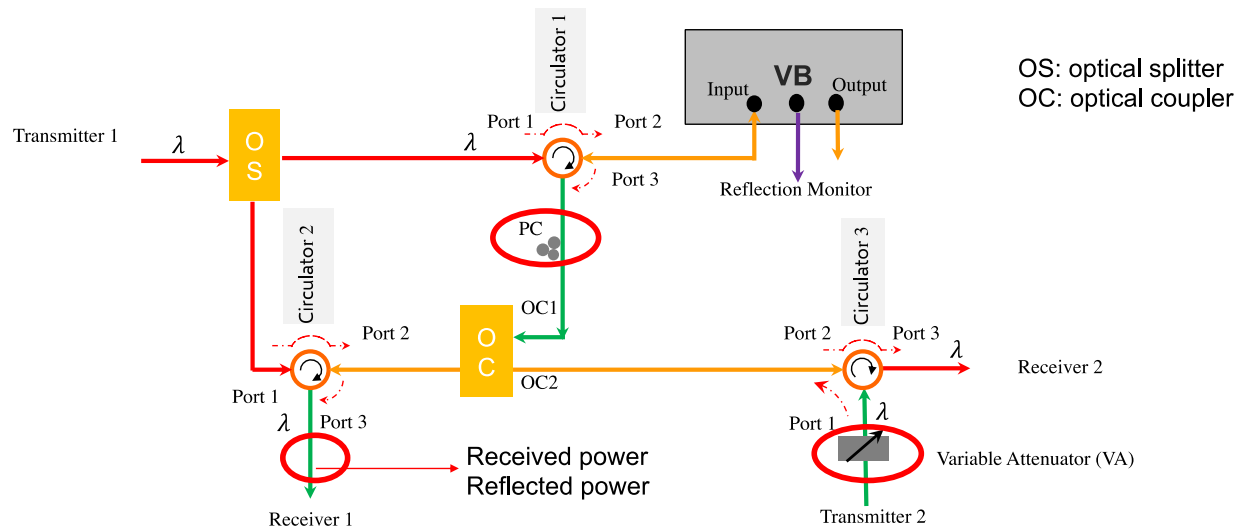


Figure 15 – Testing Case II: Variable Backreflector

Figure 15 shows the second test setup with the variable backreflector to measure the robustness of coherent signals at different return loss levels. Instead of a fixed reflection impairment used in the previous setup, we use a backreflector to purposely control the reflected power to the desired signal detection level. To achieve full duplex operation, there are two conditions that need to be satisfied at the receiver:

- The received power (the transmitted power – link loss) has to be larger than the power sensitivity requirement;
- The optical signal to noise ratio (from reflection power) has to be better than the OSNR sensitivity requirement.

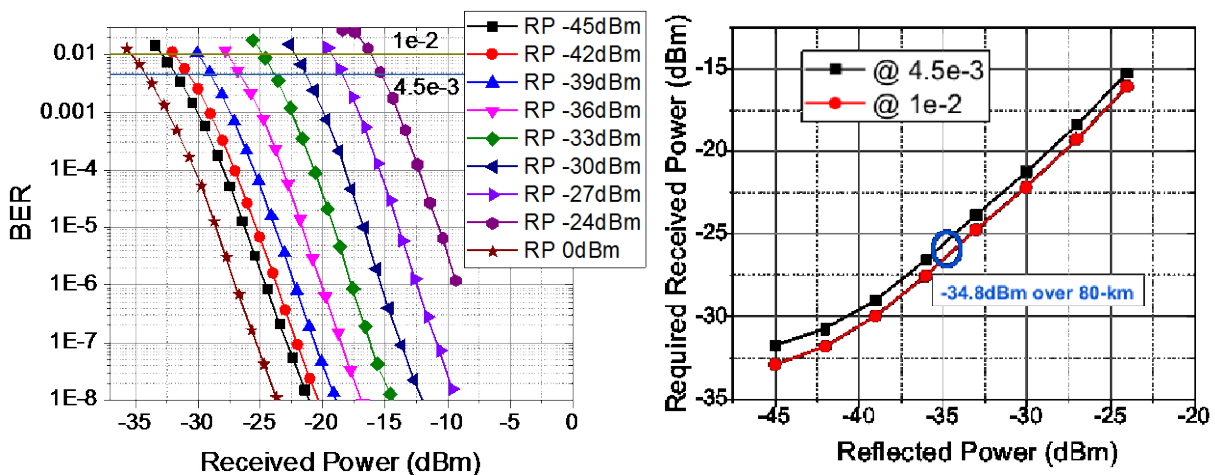


Figure 16 – Testing Case II: Results (RP: Reflected Power)

Figure 16 shows the BER vs. power curves under different reflected power levels for 100G PM-QPSK signals. The reflected power is measured before RX1 and the transmitter output power is set to 0 dBm. It is also noted that the polarization controller is inserted in the setup to emulate the worst case of polarization alignment. The required receive power almost linearly increases as the reflected power

becomes larger. For example, for the 80 km transmission case with -34.8 dBm reflected power in Test Case I, the required received power would be ~-26 dBm to maintain the required OSNR level. It is also noted that there is no error floor observed even if the reflected power is measured at -24 dBm.

Conclusion

As the industry evolves toward Node+0 architectures, the volume of optical connections to intelligent nodes will increase substantially compared to traditional architectures. Coherent optics technology offers a future-proofing solution for cable operators to meet bandwidth demand without the need for retrenching new fibers.

In this paper, we presented extensive experimental verification under different coexistence scenarios and provided operational and deployment guidance for such use cases. Less than 0.6 dB power penalty is observed with complexed hybrid scenarios. The results show coherent optics transmissions are robust, even in close proximity to much stronger analog and intensity modulated optical carriers. This means that the cable operators can effectively support 100G or higher coherent channels on their existing networks without the concerns of significant performance degradation.

Additionally, CableLabs' full duplex coherent optics proposal and the experiments that demonstrate simultaneous bi-directional transmission over single fiber and single wavelength are also discussed in this paper. The major impairment in the full duplex coherent optics system is ORL including all discrete reflections (Fresnel) and continuous reflections (backscatter). The impact of ORL for FDCO is also analyzed and quantified for various configurations. The quantitative results provide the cable operators an elegant solution to their single-fiber use cases with coherent optical systems.

Acknowledgement

The authors of this paper would like to extend their gratitude to Acacia, ARRIS and Ciena for providing analog optics and coherent optical transceiver components, in addition to their technical support for setup and operation.

Abbreviations

ADC	analog to digital converter
ASIC	application-specific integrated circuit
BER	bit error rate
bps	bits per second
CD	chromatic dispersion
CMA	constant modulus algorithm
CMOS	complementary metal-oxide-semiconductor
CMTS	cable modem termination system
DAC	digital to analog converter
dB	decibel

dBm	dB milliwatt
DCF	dispersion compensation fiber
DFB	distributed feedback (laser)
DMF	dispersion managed fiber
DSP	digital signal processing
DWDM	dense wavelength division multiplexing
ECL	external cavity laser
EDFA	erbium-doped fiber amplifier
EPON	ethernet passive optical network
ETDM	electrical time division multiplexing
EVM	error vector magnitude
FDCO	Full Duplex Coherent Optics
FEC	forward error correction
FWM	four-wave-mixing
Gbps	gigabit per second
GHz	gigahertz
HD	high definition
HFC	hybrid fiber-coax
HHP	household pass
Hz	hertz
I	in-phase
ISBE	International Society of Broadband Experts
km	kilometer
LD	laser diode
LO	local oscillator
LPF	low-pass filter
MHz	megahertz
MIMO	multi-input multi-output
MMI	multi-mode interference
MSA	multi-source agreement
MZM	Mach-Zehnder modulator
NRZ	non-return zero
NZDSF	non-zero dispersion shifted fiber
OIF	Optical Internetworking Forum
OLT	optical line terminal
OOK	on-off keying
OPLL	optical phase locked loop
PAM	pulse amplitude modulation
PBS	polarization beam splitter
PHY	physical layer
PM	polarization multiplexing
PMD	polarization mode dispersion
PON	passive optical network
Q	in-quadrature
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying
R-PHY	remote PHY
RF	radio frequency

RFoG	RF over glass
RIN	relative intensity noise
RPD	remote PHY device
SBS	Stimulated Brillouin scattering
SCTE	Society of Cable Telecommunications Engineers
SD-FEC	soft decision forward error correction
SMF	single mode fiber
SNR	signal to noise ratio
SOP	state of polarization
SPM	self-phase modulation
ULA	ultra-large area
ULLF	ultra-low loss
XPM	cross-phase modulation

Bibliography & References

- [1] Book Chapter, “Introduction to broadband access technologies and evolution of fiber-wireless systems”, in “Fiber-Wireless Convergence in Next Generation Communication Networks”. 2017, ISBN 978-3-319-42820-8.
- [2] L. A. Campos, Z. Jia, T. Liu, “Leveraging deployed fiber resources for the implementation of efficient scalable optical access networks,” Sept. SCTE/ISBE Cable-Tec Expo’16, 2016.
- [3] Cable Television Laboratories, Inc. “P2P Coherent Optics Architecture Specification”, June 29, 2018.
- [4] Cable Television Laboratories, Inc. “P2P Coherent Optics Physical Layer 1.0 Specification”, June 29, 2018.
- [5] 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.

Implications of 5G low-latency requirements on Hybrid Fiber-Coaxial Networks

A Technical Paper prepared for SCTE•ISBE by

Sanjay Dhawan

VP Strategy – Network Products and Solutions

Ericsson Inc.

6300 Legacy Dr., Plano, TX-75034

214-310-8779

sanjay.dhawan@ericsson.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
5G Requirements, Architecture and Use Cases	3
Conclusion.....	11
Abbreviations	11

List of Figures

Title	Page Number
Figure 1 - Use Case Evolution with 5G	4
Figure 2 - 5G Subscription Uptake and Traffic Growth.....	4
Figure 3 - 5G Drives Performance Boost in 8 Dimensions	5
Figure 4 - Optimizing 5G Deployments for Coverage, Throughput and Latency.....	6
Figure 5 - Maximizing HFC/DOCSIS Assets for 5G Deployment	6
Figure 6 - Achieving Low Latency with 5G.....	7
Figure 7 - Network Slicing a Key Enabling Technology.....	8
Figure 8 - 5G Network Architecture Functional View	8
Figure 9 - 5G Distributed Architecture	9
Figure 10 - HFC Throughput and Latency Requirements.....	10
Figure 11 - Coordination Considerations between HFC and 5G	10

Introduction

With the finalization of initial 5G standards in 3GPP during Q1 2018, 5G deployments are starting to gain momentum globally, with a few initial commercial launches during 2018. For Cable MSO's, 5G adds new revenue opportunities, in terms of extending Mobile Broadband service to own users, and Fixed Wireless Access where there are challenges with DOCSIS/Fiber deployment, as well as IoT use cases.

One of the key attributes of 5G is the significant reduction in latency. Unlike traditional 4G LTE systems, latency requirements in 5G vary with use cases. As an example, for a traditional smartphone web browsing service, 15-20 ms round trip times may be acceptable. However, for a use case such as autonomous driving, round trip latency requirements need to be under 10 ms. Other use cases that require sub-10 ms round trip latency include industrial robotics control, drone control, web gaming, and connected, collaborative multi-site live concerts (band members playing a song across multiple locations). Typically, low latency demands tend to be localized with communication over a short distance. In order to fulfil the low latency requirements for such use cases, new architectures need to be implemented in the network. These include implementing control functionality and local switching at the edge, which in a DOCSIS network can even be a hub site. Micro servers that can support virtual applications will need to be deployed at hub sites or cell sites and in close proximity to users. In addition, while network slicing can support multiple use cases from each radio site and can fulfil use case specific routing, bandwidth and latency requirements will need to be deployed across the networks. Present DOCSIS networks are typically designed for a median latency requirement of ~10-15 ms, which can continue to work well for traditional Mobile Broadband use cases. Also, where network slicing with Edge Servers are deployed, the current cable infrastructure may be able address the 5G requirements.

5G also introduces a Virtual RAN architecture, where Layer 3 (higher layer) RAN functionality is centralized in the cloud. The one-way latency objective between the 5G radio site and the VRAN node is typically 5 ms. To fulfill such an objective, it becomes important to maximize fiber and optical switching in the access transport network. Layer 3 ethernet switching, which can add significant delays, can be deployed between the VRAN and the Core.

The 5G scheduler is hungry, which means that it will try to get the data it receives as soon as possible to the target user. For mmWave, the scheduler has a transmit time interval (window) of ~250 micro seconds which is extremely time sensitive. The faster the data can be transferred from Core to the radio, the faster it can be forwarded to the users.

Eliminating latency bottlenecks in the transport network will be key towards maximizing the overall throughput experience of 5G networks.

5G Requirements, Architecture and Use Cases

5G is gaining momentum with extensive interest from MSO's and MNO's. In fact, the race to be 1st to the market has already begun. 5G provides an evolution from current 4G smartphone services, while at the same time adding new revenue streams for service providers. Attributes such as Gigabit throughput experience enables Fixed Wireless in Urban areas, while ultra-low latency enables autonomous automotive control and remote robotic control for manufacturing automation.

5G will enable an enhanced user experience for industrial use cases. For the above manufacturing example, 5G would enable remote control of robots with round trip latency of 10 ms. For automotive and drones, autonomous control would be achieved via ultra-low latency complemented by distributed computing. Similarly, for energy and utilities, 5G would enable real time control and automation of grids.

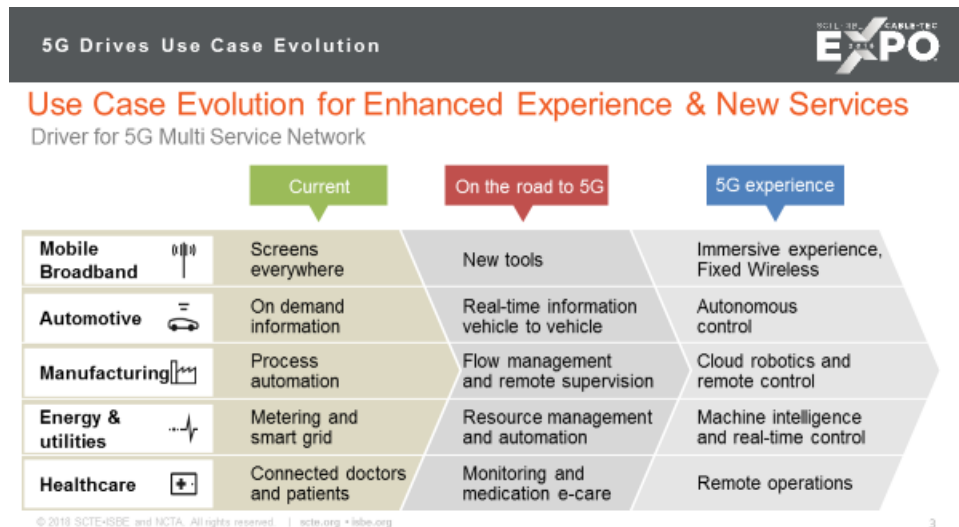


Figure 1 - Use Case Evolution with 5G

Given the initial momentum for 5G, Ericsson expects 48% of smartphones in North America to be 5G capable by 2023. Furthermore, as indicated in the Ericsson Mobility Report (2018), this is complemented by growth in IoT subscriptions from 100M today to 260M in the same time. From a user behavior standpoint, driven primarily by video, smartphone traffic is expected to grow 7 times from 2017 to 2023, to 49 GB/month.

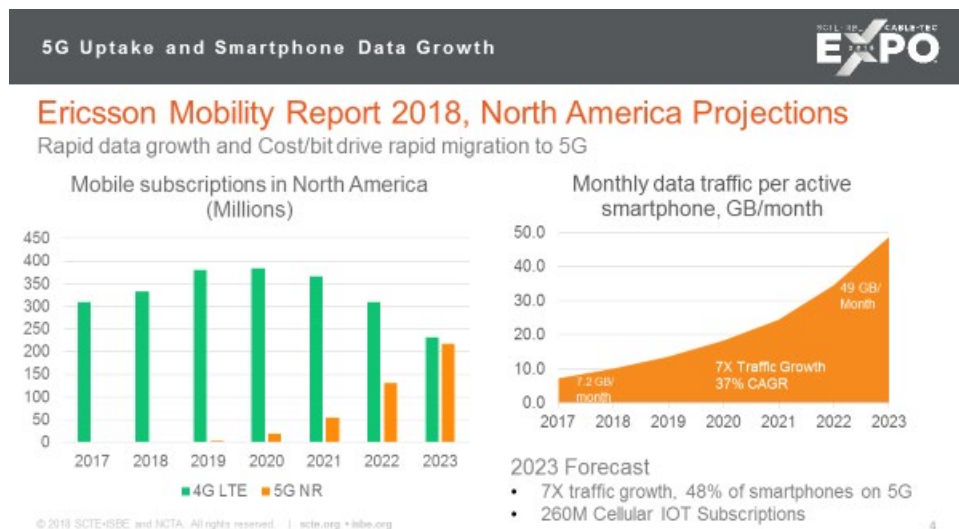


Figure 2 - 5G Subscription Uptake and Traffic Growth

5G enables performance boosts in multiple dimensions, which is key to supporting a diverse set of use cases. While the focus on initial use cases was peak Gigabit throughput, latency, reliability and positioning accuracy are emerging to be equally important along with new use cases. Several critical IoT use cases require sub 10 ms latency, 99.9999% reliability and even greater positioning accuracy than traditional GPS. On the other hand, the primary requirement for Fixed Wireless Access is a throughput experience that enables multiple 4K TVs in a home.

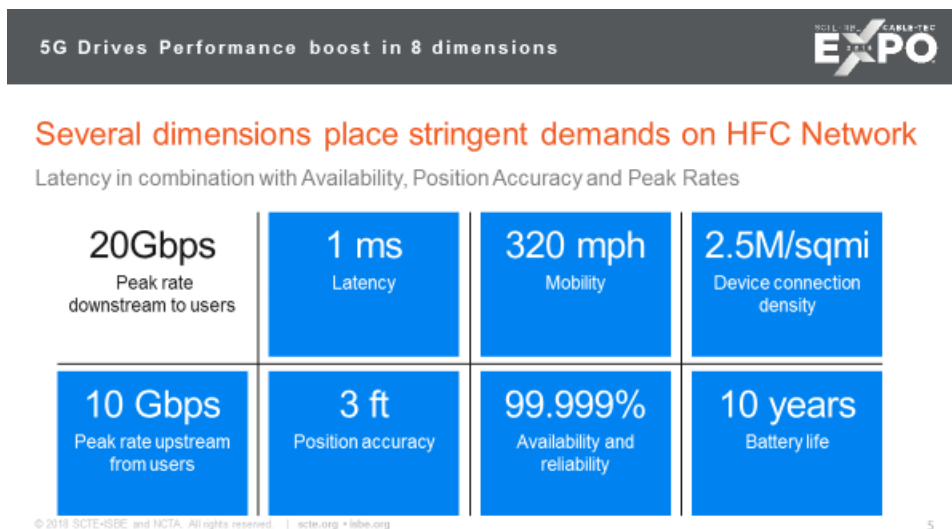



Figure 3 - 5G Drives Performance Boost in 8 Dimensions

Support for new IoT use cases complemented by exponential increase in traffic is expected to drive new ways of building networks. These networks are expected to be service slice aware, with ultra-dense small cell grids enabling Gigabit throughput along with ms latency, while macro networks provide ubiquitous coverage. Microcells leveraging mmWave (24, 28, 37 and 39 GHz) spectrum are ideal for such deployments, as mmWave provides 100 MHz – 1 GHz of spectrum per operator, and their low propagation characteristics enable an ultra-dense grid, where required.

Each class of use cases has different and distinct requirements on 5G from the perspective of coverage, bandwidth and latency. As an example, to support IoT use cases such as utility meters, ubiquitous coverage is required that includes even rural areas. Alternately, to support autonomous driving, a network that covers a wide area is required, while optimizing latency. To enable robotic manufacturing, localized indoor 5G optimized for ultra-low latency is required. Optimizing the combination of coverage with bandwidth, latency and reliability, is important for a well designed 5G network.

As can be seen from the figure below, a 5G network would require a combination of low, mid and high bands. The low bands would be ideal for providing wide coverage, while the high bands would be ideal for ultra-high capacity and ultra-low latency and require a dense grid. Similarly, mid bands are well suited for high capacity with moderate coverage. Tightly coupling 5G network elements serving low, mid and high bands would maximize the 5G experience for varying use cases.

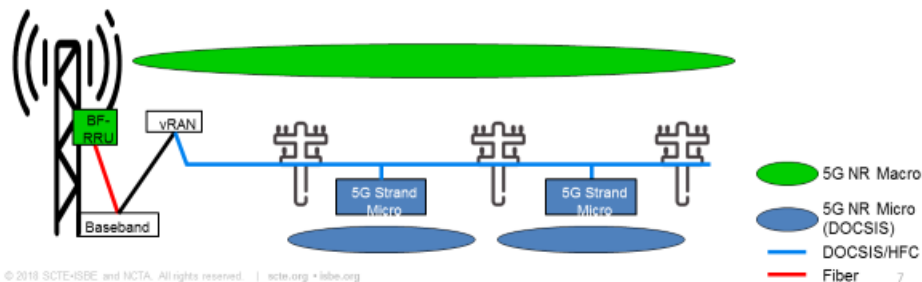
There are major fundamental trade-offs between capacity, coverage, latency, reliability and spectral efficiency in a wireless network. Due to these fundamental limits, if one metric is optimized for improvement, this may result in degradation of another metric."



© 2018 SCTE-ISBE and NCTA. All rights reserved. | scte.org • isbe.org

1

- ❖ Traditional Tower Top Deployments provide Wide Coverage Layer
- ❖ DOCSIS / HFC Compliant Deep Micro deployments for Throughput, Latency and Reliability
- ❖ Strand and Pole optimized solutions key to securing deep penetration



© 2018 SCTE-ISBE and NCTA. All rights reserved. | scte.org • isbe.org

As indicated above, for several 5G use cases, the most important attribute is sub 5-10 ms round trip latency, which in turn places <1ms latency requirements on the 5G Radio Access Network.

To enable low latency, several techniques are being implemented in 5G. To begin with, in mmWave, the scheduler has a time slot of 62.5-250 micro second. Within each slot the scheduler can serve 1 or more users, with Multi-User MIMO. Furthermore, techniques such as instant uplink access give ultra-low latency users instant access to the network for short data bursts, thereby keeping one-way latency to ~0.5-1.0 ms. Additional techniques such as mini-slot further reduce the transmit requirement to a subset of 1 timeslot. The scheduler is implemented as hungry, such that it will try to get the data out to users as quickly as possible, by maximizing the most important resource, i.e., spectrum, while managing users across excellent and poor radio conditions. As an example, if 100 MHz of spectrum is available, and there is only 1 user, it will be fully used for 1 or more slots, to get the buffered data out to the user as quickly as possible. If there are multiple users, the data push to users is optimized based on several factors, including Service QoS requirements, RF conditions, amount of data, etc. As radio conditions change rapidly, the scheduler needs to adapt. The adaptation of scheduling takes place on a timeslot basis.

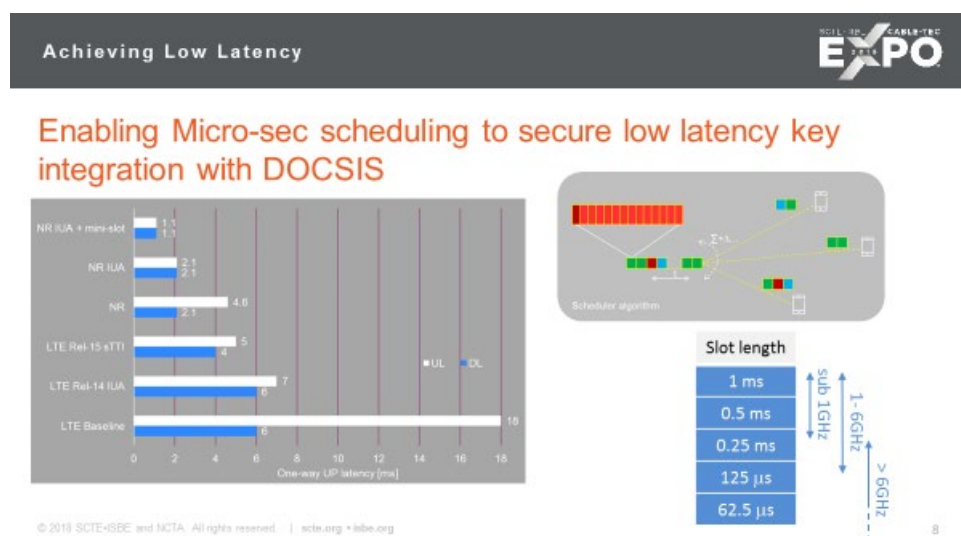


Figure 6 - Achieving Low Latency with 5G

An important architectural consideration in meeting diverse performance requirements is Network Slicing. Network slicing is akin to a VPN in IP networks, such that each VPN has its own bandwidth and QoS criteria. In 5G, where we can have hundreds of factories, each requiring dedicated bandwidth along with different QoS for different classes of devices (e.g., Robots, employees, etc.) in each factory, Network slicing is an optimal way to achieve such a requirement, without building dedicated networks for each factory. Essentially, each Network slice is a logical network serving a defined business purpose or customer, consisting of all required network resources, including Radio Access, Transport, Core and Cloud, configured together. It is created, changed and removed by management functions.

5G Networks are Slice Aware, serving diverse performance requirements for each slice

Attributes for Slicing Network: Orchestration, Virtualization, Software Defined, Distributed, Automated, Application Aware,...



A further zoom into an example of a Distributed Network Architecture with Virtual Network Functions and Application Servers in a Cable / HFC environment supporting three families of use cases (Critical IoT, enhanced Mobile Broadband and Massive IoT) is presented below.

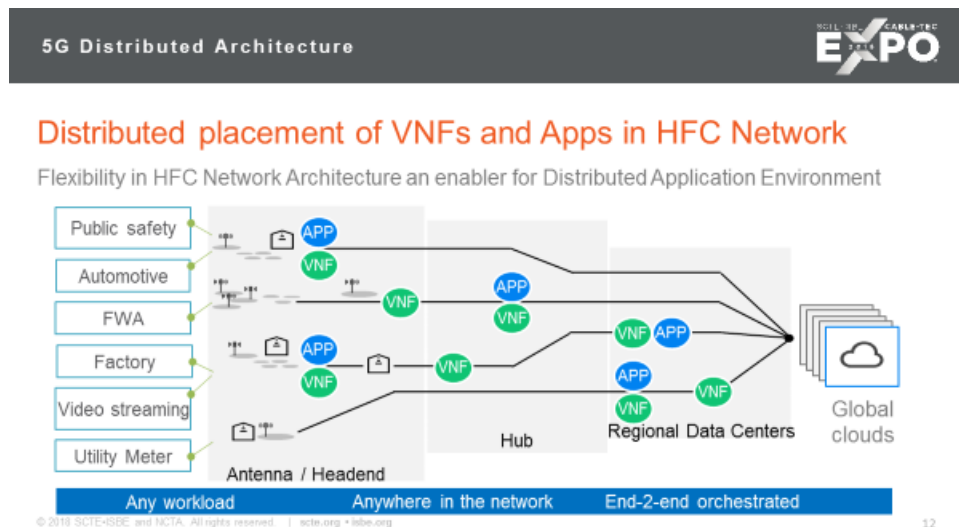


Figure 9 - 5G Distributed Architecture

Next, we consider the requirements of the distributed architecture on the HFC transport network. As shown in the figure below, for the case of critical IoT (also called Machine Type Communication, or MTC), where a radio processing function at a hub site serves a cluster of radio sites, the one-way latency between the radio site and the hub site would need to be 30 micro seconds and the bandwidth requirement would be 10-25 Gb/sec per radio. With eCPRI, ethernet support would be feasible for such an interface.

In the case of fixed or mobile broadband, each site would have the radio and the processing function collocated at the site, with the Virtual Controller at the Headend. In this scenario, the latency requirements would be less stringent and on the order of 75 micro seconds. As control signaling and payload is sent back to a server, the bandwidth requirement on the link is directly related to the payload and can be 10 Gbps or less. In such an architecture, DOCSIS / HFC network can serve as the access transport with a lesser degree of impact.

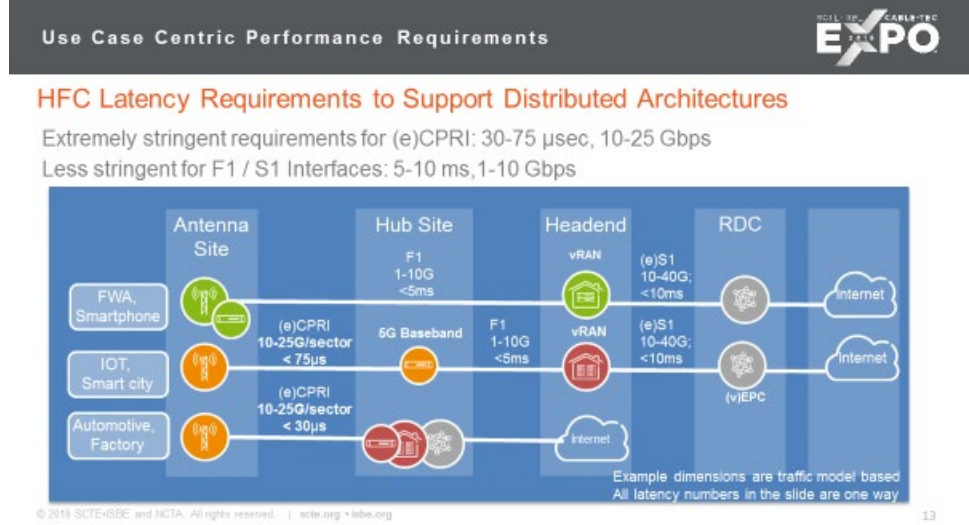


Figure 10 - HFC Throughput and Latency Requirements

In summary, an optimized 5G deployment requires multi-dimensional coordination with a DOCSIS/HFC network. To begin with extending site access, that presently supports hundreds of thousands of WiFi strand/pole mount radios, to 5G is a key factor. These sites also require extending the HFC transport and DOCSIS power to 5G pico base stations.

For several use cases, including Video Streaming to residential customers and autonomous automotive control, servers at Headend and Hub Sites capable of hosting VNFs and applications are expected to be deployed in scale.

Finally, an intelligent, fully orchestrated, Software Defined, Secure E2E network covering 5G RAN, Core and HFC are key to a fully automated and seamless service experience for the users.

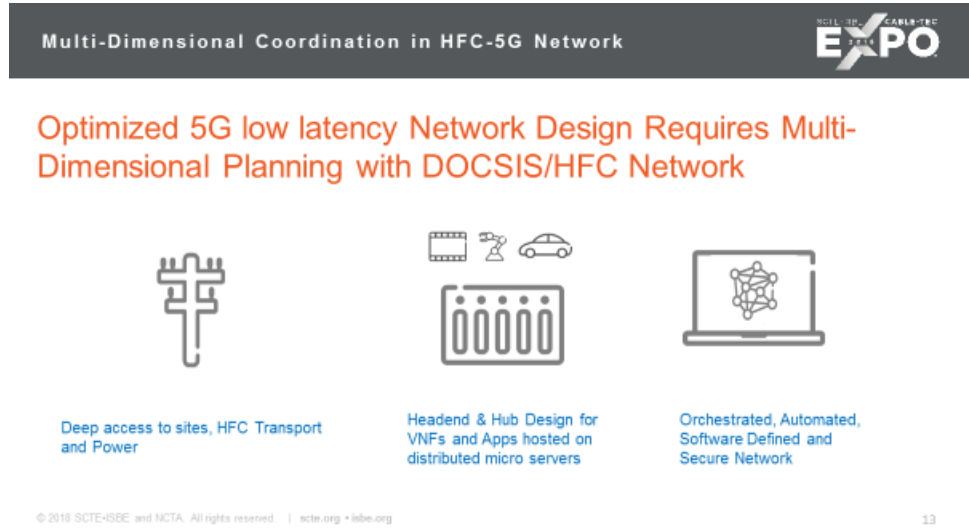


Figure 11 - Coordination Considerations between HFC and 5G

Conclusion

Summary



5G Architecture Flexibility Drives Seamless HFC Integration

- 5G drives performance boost in 8 dimensions
- New Architecture considerations with 5G to enable Network slicing
- 5G Scheduling optimized to deliver gigabit throughput with ultra low latency
- Strand/Pole optimized 5G Radio with DOCSIS Interfaces enable seamless integration for dense deployments



HFC Networks enabling deeper penetration for Stringent 5G performance

© 2018 SCTE•ISBE and NCTA. All rights reserved. | scte.org • isbe.org

15

In summary, 5G is designed to deliver multi gigabit throughput experience, sub 5 ms latency, 99.9999% reliability and 3 ft positioning accuracy. This is achieved via highly optimized a network slice-aware distributed architecture complemented with microsecond level radio scheduling.

From a radio site perspective, for several of the use cases, the DOCSIS / HFC can be directly connected into an integrated Radio-Baseband unit and provide backhaul to a Headend or a Regional Distribution Center site. This will enable a deeper penetration of 5G small cells into urban and residential areas, thereby enhancing the coverage of IoT and broadband services.

Abbreviations

5G NR	5G New Radio
AAS	Adaptive Antenna System
BFF	Beam Forming Function
BPF	Baseband Processing Function
cMTC	Critical Machine Type Communication
CU	Central Unit
CU-C	Central Unit Control Plane
CU-U	Central Unit User Plane
DU	Distributed Unit
eCPRI	Packet based Common Public Radio Interface
FWA	Fixed Wireless Access
IUA	Instantaneous Uplink Access
KPI	Key Performance Indicator
MBB	Mobile Broadband
mMTC	Massive Machine Type Communication
mmWave	millimeter Wave
PPF	Packet Processing Function
RF	Radio Function

RPF	Radio Processing Function
RCF	Radio Control Function
TTI	Transmit Time Interval

Improving the Customer Experience with Network Automation and with AI-Powered Voice

A Technical Paper prepared for SCTE•ISBE by

Pravin Mahajan
Infinera

PMahajan@Infinera.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content – Cognitive Networks	3
Conclusion.....	8
Abbreviations.....	8

List of Figures

Title	Page Number
Figure 1 - Path to a Cognitive Network	4
Figure 2 - Evolution from OSI to Layer T and Layer C Model.....	5
Figure 3 - High Capacity through Super-channels.....	6
Figure 4: Software Defined Capacity	7

Introduction

Think about the ways we've been navigating around computers and information systems for the past four decades: keyboards, mice, touchscreens. What do they have in common? They are all designed for people who can use their hands as part of their information task. But what about the workforce that needs to keep their hands on their tools and equipment?

Fast-paced migration to the cloud, forthcoming 5G deployments, and the proliferation of internet-connected devices driven by Internet of Things (IoT) including voice enabled devices like Alexa & Siri are fueling the migration toward cognitive networking. Cognitive networks use advanced analytics, machine learning, and artificial intelligence techniques to help build self-optimized, self-healing and highly autonomous transport networks, setting new benchmarks in scalability, agility and automation.

This paper describes how Network automation, AI powered technology & Voice can help build a cognitive network which will greatly improve the end customer experience.


Content – Cognitive Networks

Cognitive networks use advanced analytics, machine learning and artificial intelligence techniques to help build self-optimized, self-healing and highly autonomous transport networks, setting new benchmarks in scalability, agility and automation. This article describes the journey to cognitive networking by explaining its key building blocks, such as software defined capacity (SDC), multi-layer software-defined networking (SDN) control and transport technology breakthroughs that make the creation and the deployment of cognitive networks a fast-approaching reality.

Cognitive networking is the ultimate goal for the intelligent transport layer that underpins all cloud-based digital communications. By definition, a cognitive network is multi-layer, self-aware, self-organizing and self-healing, and can take predictive and/or prescriptive action based on what it has gleaned from its collected data and experience. Realistically, no network can completely plan or run itself; however, cognitive networking will dramatically reduce the number of manual tasks required across a multi-layer network. This can be achieved by leveraging advanced software, streaming telemetry, big data with machine learning and analytics to autonomously conduct network operations to meet the demand for connectivity, maximize network resources and increase reliability. There are multiple important elements in a cognitive network, such as:

- **Advanced analytics** designed to parse streams of machine data to monitor network health and raise awareness of any anomalies
- **Machine learning** software tools that leverage advanced analytics to understand and identify trends in operations
- **Autonomous** hardware and software capable of executing various tasks and conducting required maintenance
- **Predictive intelligence** tools capable of identifying potential problems before they happen
- **Prescriptive** software tools designed to proactively recommend new solutions for maximizing capacity, enhancing reliability and optimizing assets

Leading the way to Cognitive Networking



12 | © 2017 Infinera

Figure 1 - Path to a Cognitive Network

Cognitive networking is the result of seamless and highly dynamic interaction between software and hardware assets across network layers and brings optical networking to a new level of scalability, flexibility and automation. The following section describes how to build the foundation for cognitive networking by using the latest technology breakthroughs.

The journey to cognitive networking starts by building the foundation of a highly scalable, flexible and programmable network architecture. The building blocks for this foundation are described in the following bullets:

- 1. Evolve the network architecture:** A well-defined architecture dictates how networks are planned, operated, and evolved. Today, when content is king and must be accessible anywhere, anytime, and on any device with the highest level of quality, it is clear that the 1980s-era seven-layer Open Systems Interconnection (OSI) model has reached a tipping point. It needs to support the transformation in networks (e.g. network functions virtualization [NFV], SDN, etc.) and the new service delivery model based on cloud applications, service virtualization, etc. The OSI model's heritage of function- and layer-specific network definitions led to closed and proprietary protocols, rigid networking capabilities and high operational costs. This is sparking an urgent need to evolve toward a simpler, more efficient and agile architectural model to accelerate the adoption of cloud-based networking. Thus, the first step in paving the way for cognitive networking consists of evolving the network architecture to a simpler cloud-powered model: a cloud services layer, Layer C, and an intelligent transport layer, Layer T, as depicted in Figure 2. This new model consolidates and simplifies cloud service delivery and networking into two layers, wherein all the OSI networking layers (Layer 3 and below) are represented by Layer T, while all the application layers (Layer 4 and above) are grouped under Layer C. Layer T sets the guidelines and principles for the transport of data streams, whether between end users and data centers, or multiple data centers with bursty and often unpredictable traffic patterns. It also

defines the features and capabilities that increase network agility and performance and sets new benchmarks for service delivery and cost-effectiveness, key ingredients to the successful deployment of any cloud application. Layer C contains all the applications, functions and services that run in the cloud, including consumer and business applications, SDN-based service creation and orchestration tools, software frameworks and applications for big data and machine learning, virtualized network functions (VNFs) and many others, to enable the large-scale task automation and programmability that streamline operations, eliminate human error, and reduce operating costs.

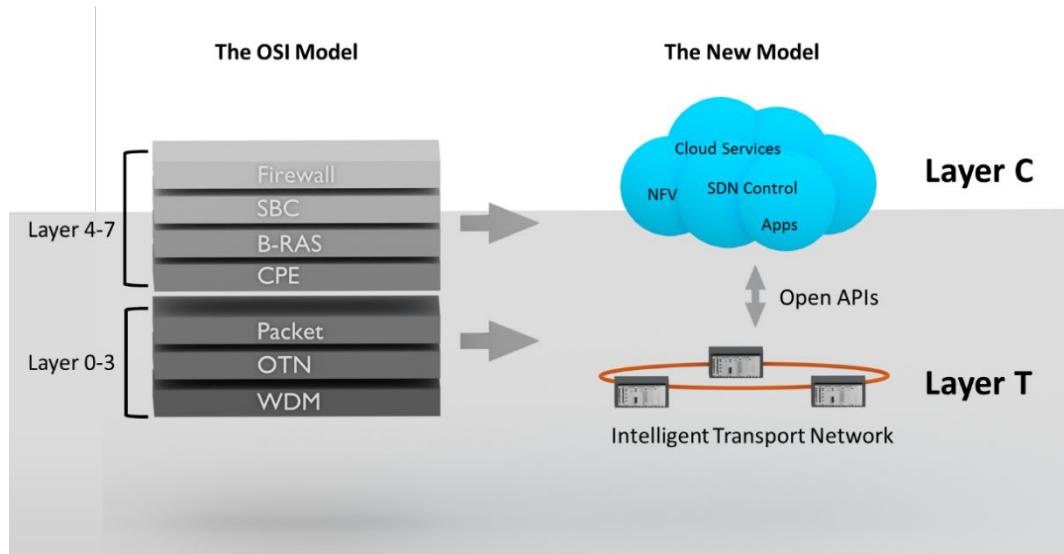


Figure 2 - Evolution from OSI to Layer T and Layer C Model

2. **Unlock the network's full capacity through super-channels:** Dense wavelength-division multiplexing (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, the current growth in internet traffic and enterprise migration to the cloud is demanding a new level of scalability. An innovation called super-channel, evolved to take DWDM networks to a new era of high capacity and optical performance, all without increasing 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, as depicted in Figure 3. Super-channels overcome three fundamental challenges: optimizing DWDM capacity and reach, scaling bandwidth without scaling operational procedures and supporting next-generation high-speed services such as 100 Gigabit Ethernet (GbE), 400 GbE, etc. The use of super-channels increases spectrum efficiency and thus network capacity by reducing spectrum waste due to guard bands. It also allows seamless capacity growth without the need for network re-engineering or major disruption to current operating processes.




Figure 3 - High Capacity through Super-channels

3. **Leverage software defined capacity (SDC):** A key stepping stone toward cognitive networking is to break away from the current methods of optical capacity planning, engineering and hardware-based deployment that require numerous truck rolls, extensive manual labor, and human interaction at multiple points in the network. The road to cognitive networking 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 (Layer C), as depicted in Figure 4. SDC provides instant software activation of additional capacity, creating a pool of bandwidth that can be dynamically allocated based on traffic demand. SDC extends the principles of SDN, which has primarily focused on Ethernet and packet layers, to the optical transport layer. With intelligent software tools, a network can become an integral part of the rest of the information technology infrastructure, enhancing service turn-up and management. SDC is a true game-changer from both business and operational perspectives. It allows a perfect match between the timing of capital expenditure (CapEx) and service revenue, thus accelerating time to revenue from month to minutes, while reducing operational expenditure (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.


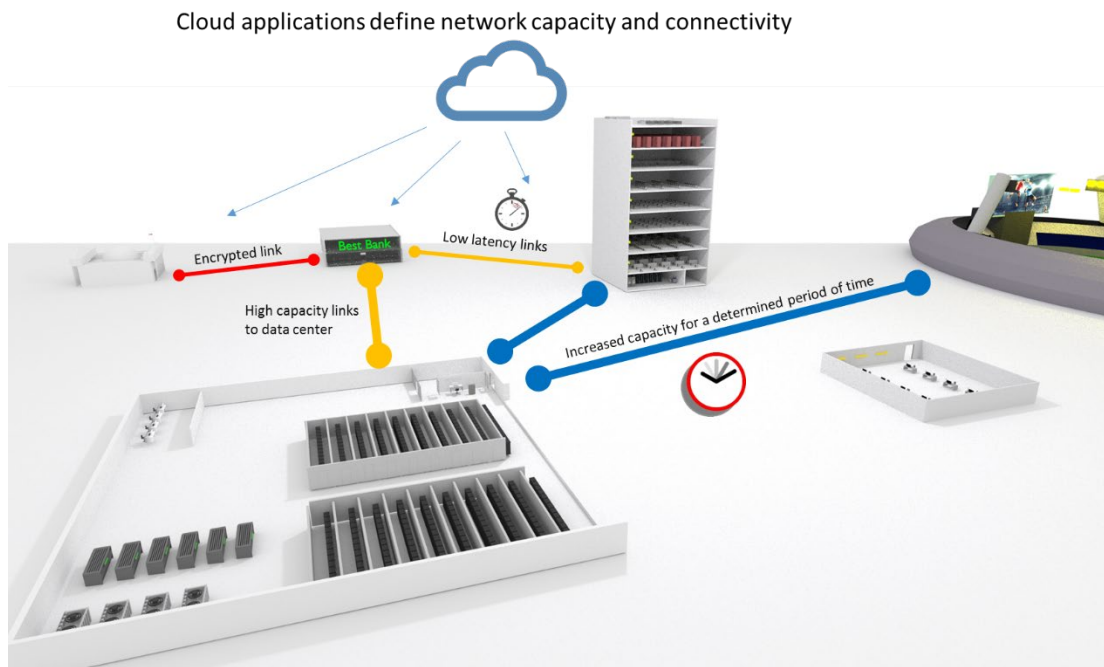


Figure 4: Software Defined Capacity

4. **Automate network operations through software-defined networks:** Building automation and intelligence across all network layers and various operating tasks is central to cognitive networks. Multi-layer SDN controllers and frameworks span both Layer C and Layer T to act as the brain of the network by providing cloud-based intelligence to plan, monitor and conduct various network operations without human intervention. SDN frameworks also provide various ready-to-deploy network applications, such as real-time capacity planning, bandwidth on demand, network virtualization and many others, to unlock full network potential and take advantage of the dynamic transport layer (Layer T). For example, very sophisticated algorithms and data models can be used to build a microservices-based path computation engine (PCE). The PCE replaces manual offline route and capacity planning processes with highly automated, real-time service

planning and activation over optimal routes across both layers, multiple paths and often-challenging fiber impairments. As cognitive networking relies heavily on streaming telemetry, big data with machine learning and analytics to learn and adjust, it is crucial to ensure a seamless flow of information between the various parts of the network elements and the upper layer software tools and SDN controllers. This flow is enabled by open application programming interfaces (APIs) such as RESTCONF, NETCONF/YANG and other northbound interfaces that connect to upper-layer orchestration systems or same-layer intelligent tools and scripts to coordinate and optimize resources across the network as well as to orchestrate VNFs. This bidirectional flow of data serves as the bloodstream of the network, delivering predictive and prescriptive real-time recommendations and taking actions to enable maximum performance.

Conclusion

The characteristics for a modern, cognitive network which can handle today's cloud-based applications was described. Cognitive networks use advanced analytics, machine learning and artificial intelligence techniques to help build self-optimized, self-healing and highly autonomous transport networks, setting new benchmarks in scalability, agility and automation. In this, the network shares the modern interfaces of today's applications. The four characteristics of a cognitive network were described: a different way of looking at the network architecture, greater capacity through super-channels, the concept of software defined networks is extended to be software defined capacity, and the management of these software defined networks is enhanced with automated network operations.

Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
CapEx	Capital Expense
DWDM	Dense Wavelength Division Multiplexing
GbE	Gigabit Ethernet
IoT	Internet of Things
NETCONF	Network Configuration Protocol
NFV	Network Function Virtualization
OpEx	Operations Expense
OSI	Open Systems Interconnection
PCE	Path Computation Engine
SDC	Software Defined Capacity
SDN	Software Defined Network
RESTCONF	RESTful Configuration protocol
VNF	Virtualized Network Function
YANG	A data modeling language

Increasing Cable Bandwidth Through Probabilistic Constellation Shaping

A Technical Paper prepared for SCTE•ISBE by

Patrick Iannone

Member of Technical Staff
Nokia Bell Labs
791 Holmdel Rd., Holmdel NJ 07733
+1-732-285-5331
Pat.Iannone@nokia-bell-labs.com

Yannick Lefevre

Member of Technical Staff
Nokia Bell Labs
Copernicuslaan 50, 2018 Antwerp
Yannick.Lefevre@nokia-bell-labs.com

Werner Coomans

Head of Department, Copper Access
Nokia Bell Labs
Copernicuslaan 50, 2018 Antwerp
Werner.Coomans@nokia-bell-labs.com

Dora van Veen

Distinguished Member of Technical Staff
Nokia Bell Labs
600 Mountain Ave., Murray Hill NJ 07974
Dora.van_Veen@nokia-bell-labs.com

Junho Cho

Member of Technical Staff
Nokia Bell Labs
791 Holmdel Rd., Holmdel NJ 07733
Junho.Cho@nokia-bell-labs.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
1. Probabilistic Constellation Shaping	4
2. PCS for Digital Optical Links in DAA HFC Networks	5
3. PCS Over Coaxial Links	7
4. Potential Applications of PCS to the PON Segment in Future FTTH Variants	9
Conclusion.....	11
Abbreviations	11
Bibliography & References.....	12

List of Figures

Title	Page Number
Figure 1 – Examples of PCS constellation with various entropy rates	4
Figure 2 – 4-PAM and PCS 8-PAM constellations with same entropy rate of 2 bit/symbol	4
Figure 3 – Rate adaptation via uniform and PCS QAMs	5
Figure 4 – Legacy HFC access network	6
Figure 5 – Distributed access architecture with digital optical link terminated on an Ethernet switch.....	6
Figure 6 – Distributed access architecture using DWDM wavelengths to connect to RPDs.....	7
Figure 7 – SNR gap to capacity of different coding schemes as a function of information rate	8
Figure 8 – Distributed access architecture with PON connections to end users	10

List of Tables

Title	Page Number
Table 1 – Characteristics of the first five PCS-LCM modulation indices defined in [20].....	8

Introduction

In the mid 1990s, multiple system operators (MSOs) began replacing coaxial cable trunks with optical links, creating the first hybrid-fiber coaxial (HFC) networks [1]. This new architecture reduced capital and operations costs, but required highly linear optical transceivers capable of transporting a full spectrum of analog RF signals while satisfying stringent electrical signal-to-noise ratio (SNR), composite triple beat (CTB), and composite second order (CSO) requirements. In the ensuing decades, wavelength-division multiplexing (WDM) technology was applied to these linear optical networks to aggregate and route headend-to-hub traffic and to split fiber nodes (FNs) in the access plant, thereby decreasing the number of households served per FN and thus increasing the available bandwidth per subscriber. Nonlinear optical impairments, primarily four-wave mixing (FWM) and cross-phase modulation (CPM), limit WDM channel counts and optical launch powers in linear HFC networks [2], adding further constraints as compared to the digital optical links used for telecom and datacom networking.

The distributed access architecture (DAA) [3], a widely accepted vision for the future evolution of HFC, includes remoting the digital-to-RF interface from the headend to the DAA node. In this scenario, the legacy analog optical links are replaced with digital links, opening a vast array of existing, high-performance, digital optical technologies to MSO network designers. Here, we describe a spectrally efficient and flexible modulation technique that has been recently developed for digital optical transmission links, probabilistic constellation shaping (PCS), that has several potential applications in DAA networks:

- High-speed digital optical links from the headend to the FN (or DAA node);
- Coaxial cable links from the FN (or DAA) to the user (modem);
- Future flexible-rate passive optical network (PON) systems for the next generation evolution.

PCS has been known for decades [4] as an essential element of communications to approach the capacity of a Gaussian channel, known as the Shannon limit [5]. It reached large scale adoption in the mid 90's inside dial-up and fax modems [6], but has been omitted in subsequent higher rate technologies that adopted a paradigm shift from single-carrier to multi-carrier modulation. The invention, in 2015 [7], of an efficient implementation of PCS for optical transmission systems has led to the rapid commercialization of this technology for core optical networks using coherent optics [8].

Although commercial optical coherent systems typically use square quadrature amplitude modulation (QAM) constellations with a uniform probability distribution, optimally performing constellations for a fixed average transmit power should follow a Gaussian probability distribution, yielding a “shaping gain” of up to ~1.5 dB in SNR compared to square QAM. New PCS-based coherent optical networks will benefit from this SNR improvement to increase the aggregate data rate to within a fraction of a dB of the Shannon capacity of optical fiber [9], and enable fully flexible control of transceiver rates, thereby realizing systems with optimized and consistent operating margins in any network configuration.

In this paper, we summarize the basics of PCS for high-speed optical links and describe how PCS-enabled coherent optics may be applied to aggregated digital links in DAA systems. We also report current progress toward transferring this technology to wired copper networks leveraging data over cable service interface specification (DOCSIS) and digital subscriber line (DSL) technologies. In the context of fiber-deeper MSO networks, we explore the potential impact that PCS can have on the optical transport and access networks, including applications to future flexible-rate passive optical network (PON) systems for fiber-deeper architectures.

1. Probabilistic Constellation Shaping

The basic idea of PCS is to transmit a constellation symbol of a smaller energy with a higher probability than that of a larger energy (hence called “probabilistic”), so as to mimic continuous Gaussian signaling that consumes the smallest transmit energy to achieve a desired data rate [5]. In practice, PCS creates a Gaussian distribution not on a continuous symbol set but on a finite and discrete symbol set, as shown in Fig. 1, which is often referred to as the *Maxwell-Boltzmann (MB)* distribution¹. The variance of the MB distribution determines the *entropy rate*, which is the maximum number of information bits that can be carried by a symbol. For example, as the variance of the MB distribution of the two-dimensional PCS 64-QAMs in Fig. 1 increases, the entropy rate also increases. When a desired entropy rate is given, it is the MB distribution that consumes the smallest average energy among all possible probability distributions for the discrete symbol set [10]. To see the energy efficiency of PCS by an example, the two-dimensional PCS 64-QAM in Fig. 1(b) is projected onto one dimension, such that a PCS 8-ary pulse amplitude modulation (PAM) is produced as shown in Fig. 2(b) with the probability mass function (PMF) of $\mathbf{p} = [p_1, \dots, p_8] = [0.0006, 0.0147, 0.124, 0.3607, 0.3607, 0.124, 0.0147, 0.0006]$. This PCS 8-PAM creates an entropy rate of $H(\mathbf{p}) = -\sum_{i=1}^8 p_i \log_2 p_i \approx 2$ bit/symbol, by consuming an average energy of $[(-7)^2, (-5)^2, (-3)^2, (-1)^2, 1^2, 3^2, 5^2, 7^2] \times \mathbf{p}^T \approx 3.75$. On the other hand, for the same entropy rate of 2 bit/symbol, and the same minimum distance between constellation points, a uniform 4-PAM shown in Fig. 2(a) uses a much higher average energy of $[(-3)^2, (-1)^2, 1^2, 3^2] \times [0.25, 0.25, 0.25, 0.25]^T = 5$.




Figure 1 – Examples of the probabilistically shaped square 64-QAM constellation, with entropy rates of (a) 3.2 bit/symbol, (b) 4 bit/symbol, and (c) 4.8 bit/symbol.




Figure 2 – One-dimensional constellations with the same entropy rate of 2 bit/symbol: (a) uniform 4-PAM, and (b) PCS 8-PAM.

¹ It may more commonly be called the Boltzmann distribution in today’s nomenclature, but has historically been called the Maxwell-Boltzmann distribution in the context of PCS.

Although the capacity-achieving performance of PCS has been known since information theory was first established [5], and despite intense effort to advance its practical implementation in the mid 1980s to the early 1990s [4, 10-12], PCS did not initially make it into many applications. One notable exception was the V.34 voice band modem technology over telephone lines standardized by the International Telecommunication Union (ITU) in 1994, that adopted PCS in the form of shell mapping [6]. At the time, V.34 became a popular modem technology for dial-up internet and was also commonly implemented in fax devices. It is only three years ago that PCS began to draw enormous attention again with the focus on optical communications, when an implementation method was reported that is both capacity-approaching and practical [7]. From that time onward, PCS has set numerous new transmission records in optical fiber communications across virtually all transmission distances from 50 km to 11,046 km, with spectral efficiencies ranging from 17.3 b/s/Hz down to 5.7 b/s/Hz [13-17]. This is attributed largely to the flexible rate adaptability, and partly to the optimal energy efficiency. Traditional uniform QAMs permit only a handful of integer-valued entropy rates, hence cause an excessive margin in most optical links where rate adaptation by forward error correction (FEC) codes is limited. On the other hand, PCS can create an arbitrary real-valued entropy rate (cf. Fig. 1) that matches any given optical link, thereby leaving only an intended margin, as illustrated in Fig. 3. With the recent commercialization of PCS for coherent digital optical networks [8], this new technology has gone from the invention of an efficient and near-optimal implementation to a laboratory demonstration to commercialization in a short period of three years. The speed with which PCS has been commercialized is a testament to its potential impact on optical network performance.




Figure 3 – Rate adaptation via uniform and PCS QAMs.

2. PCS for Digital Optical Links in DAA HFC Networks

Legacy HFC networks, as shown in Fig. 4, include analog RF optical links on a pair of fibers between the head-end (HE) or hub node and the fiber node (FN). The FN is the optical-to-electrical interface between this analog optical link and the coaxial cable distribution plant, that serves several hundred households via cascades of RF amplifiers.

Over the past several years, the cable industry has coalesced around a vision for the evolution of HFC networks that leverages the technical and economic advantages of digital optics to replace the analog optical link with a more robust, higher performance digital link. CableLabs recently released a pair of specifications describing point-to-point (P2P) 100-Gb/s coherent optics [18] and related architectural use cases [19]. The P2P specification is based on coherent transceiver technology using differential quadrature phase shift keyed (DQPSK) modulation, which has been widely deployed in long-haul and metro networks. Since HFC hub-to-FN links are relatively short (typically substantially less than 100 km) and the aggregated data rates are moderate, the specification defines relaxed requirements as compared to commercial dense WDM (DWDM) long-haul transceivers, which should result in reduced costs. These include lower output power, lower optical signal-to-noise ratio (OSNR), relaxed number and spacing of

ITU DWDM channels, and limited chromatic dispersion compensation. Assuming these 100-Gb/s P2P transceivers gain acceptance, PCS-enabled transceivers would be a potential next-generation upgrade capable of providing the maximum data rate allowed by the link given the transceiver's launch power.




Figure 4 – Legacy HFC access network.

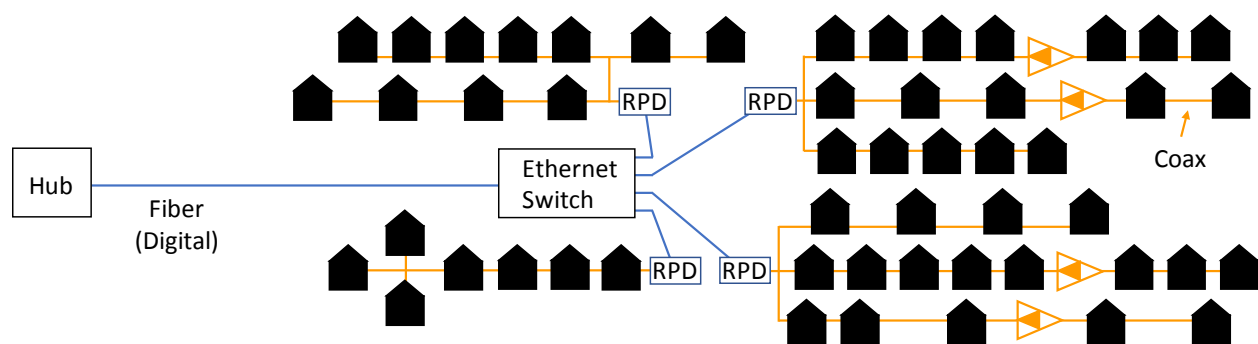


Figure 5 – Fiber deeper Distributed Access Architecture (DAA) with the aggregated digital optical link terminated on an Ethernet switch. Lower rate optical Ethernet links connect the switch to RPDs, located at the sites of legacy RF amplifiers.

The most relevant architectural use cases described in [19] are variations of the distributed access architecture (DAA), for which the analog optical link in a traditional HFC network is replaced with a digital link, and fiber is deployed deeper into the network, significantly reducing the number of households passed per fiber termination and creating new opportunities such as backhauling wireless small cells. One instantiation, shown in Fig. 5, replaces the analog RF optical link with a digital coherent link terminating on an Ethernet switch. The Ethernet switch is in turn connected via lower rate point-to-point optical Gbit or 10-Gbit Ethernet links to remote PHY devices (RPD) or remote MACPHY devices (RMD), that convert these bidirectional digital optical signals to RF signals sent to end users over the coax. Another version of this architecture, shown in Fig. 6, uses multiple DWDM wavelengths passively routed via DWDM optics, or a time division multiplexed (TDM) passive optical network (PON) with a passive splitter, between the hub and the RPDs/RMDs. These options are compatible with PCS-enabled transceivers to optimally adapt the data rate to the particular link. Consider, as an example, an upgrade for the hub-to-switch architecture (Fig. 5) for which a single fixed-rate 100-Gb/s wavelength according to the P2P standard is replaced by PCS-enabled transceivers having a range of finely adjustable data rates from 100 Gb/s up through 600 Gb/s. In this case, the PCS transceivers would automatically optimize the data rate to the link, while maintaining a constant minimum system margin (i.e. separation between the signal SNR and the Shannon capacity), thereby minimizing launch power.

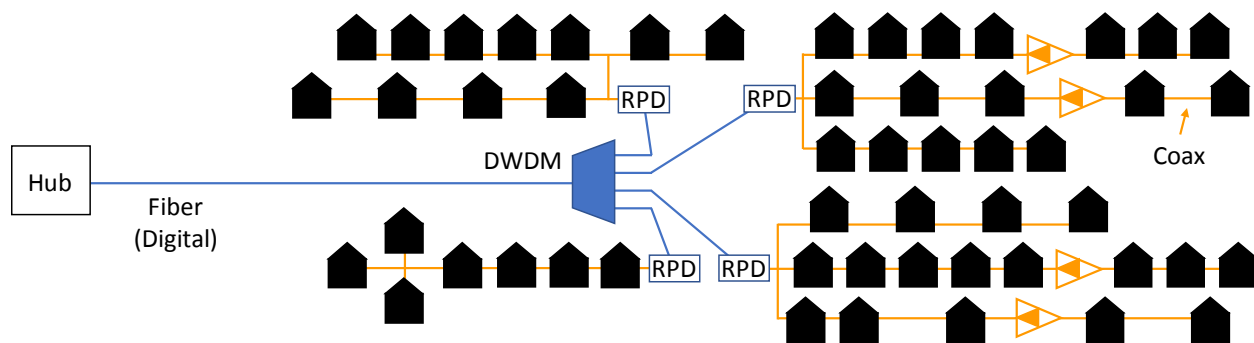


Figure 6 – Fiber deeper Distributed Access Architecture (DAA) using distinct DWDM wavelengths to connect to multiple RPDs. {Note that a PON variant of this architecture is also possible, for which TDM or TWDM optical signals connect the Hub to OLTs (sec. 4).}

3. PCS Over Copper Links

PCS can also be beneficial for transmission over copper access networks, consisting of coaxial cables or telephone wires, but requires several changes to the scheme as reported in [7]. These are explained in detail in [20] and are summarized below.

One major difference with digital optical transmission is that copper access technologies, such as DOCSIS 3.1 for cable and G.fast for DSL, use multi-carrier modulation formats like orthogonal frequency division multiplexing (OFDM), as opposed to single carrier modulation in digital optics. So instead of shaping a single communication channel (or a relatively small number of independent channels in the case of WDM) with a single SNR, thousands of small subchannels have to be shaped simultaneously, each with a potentially different SNR. A solution to this problem consists of using a set of pre-defined shaping codes labeled by so-called “modulation indices”, each related with a particular QAM modulation order [20]. Similar to bit-loading in DSL, or modulation profiles in DOCSIS 3.1, one can then assign a modulation index (and the corresponding shaping code) to each subcarrier based on its SNR value. This implies that the algorithms and protocols used for assigning modulation orders in current technologies can be re-used to realize shaping based on modulation indices, and thus that this approach requires minimal standard changes.

Another important difference is that the maximum obtainable SNR on copper networks can be extremely high (40 to 50 dB), which supports very large constellation sizes (e.g., up to 2^{14} -QAM). A characteristic of the PCS scheme as discussed above is that all parity bits from the low-density parity-check (LDPC) code must be transmitted on the sign bits, so as not to alter the occurrence probability of the modulated symbols (i.e., the shaping). This approach introduces a constraint on the modulation orders than can be achieved with a given code rate. For instance, a code rate of 3/4, as used in upstream DOCSIS 3.1, would only allow shaping for constellation sizes up to 2^8 -QAM. Hence, the constraints due to this coding approach effectively limit the maximum system performance. To remove this constraint, a set-partitioning coding technique, known as LDPC-coded modulation (LCM) can be exploited [21]. Compared to regular LDPC as used in DOCSIS 3.1, which encodes all bits, LCM only encodes the least significant bits, i.e., a subset of the bitstream. A modified LCM scheme, which maps the parity bits to the sign bits, can be combined with PCS [20]. With this approach, the constraint is shifted from imposing a maximum modulation order to imposing a maximum number of coded bits per symbol (size of the subset). LCM has the additional advantages that it achieves a better net data rate for the same error-correction performance compared to regular LDPC, and a significantly less complex encoder/decoder since only a subset of the bits have to be encoded.

Table 1 – Characteristics of the first five PCS-LCM modulation indices defined in [20].

Modulation Index	QAM modulation order	Information rate [bits/QAM symbol]	SNR [dB]
1	4	2.6	9.06
2	6	3.4	11.81
3	6	4.4	14.86
4	8	5.5	18.17
5	8	6.4	21.17

Table 1 shows as an example the 5 first modulation indices defined in [20], using a rate-28/33 LCM encoding with 6 coded bits. The shaping codes can be designed to obtain any desired distribution of SNR operating points. In the example, the codes have been designed to obtain operating points that lie about 3 dB apart, leading to a uniform coverage of the useful SNR range. As can be seen from the table, modulation indices 2 and 3 both use the same modulation order (2^6 -QAM), but realize a different information rate by using a different shaping code. Figures 7(b) and (c), respectively, show histograms of the received frequency-domain IQ samples of modulation indices 2 and 3, i.e. the PCS distribution plus a zero-mean Gaussian noise cloud at the receiver. With this approach, PCS removes the need for odd-bit constellations (i.e., 2^3 -, 2^5 -, and 2^7 -QAM and higher) while still obtaining the same granularity of SNR operating points. Since odd-bit constellations require more complex mapping and demapping than their even counterparts, PCS effectively reduces the QAM mapping/demapping complexity.




Figure 7 - (a) Comparison of the SNR gap to capacity of different relevant coding schemes as a function of the information rate. The distribution of received samples for PCS-LCM modulation indices 2 and 3 are shown in respectively (b) and (c).

Note that in upstream DOCSIS 3.1 and DSL, the main advantage of using PCS is the shaping gain, and not the matching of capacity and line rate, which can already be done by tailoring the transmit power and/or bit-loading on a per-subcarrier basis. The shaping gain is illustrated in Figure 7(a), where we compare the SNR gap to capacity of the proposed PCS-LCM scheme (blue diamonds) with that of three other schemes:

- inner trellis coded modulation (TCM) with an outer Reed-Solomon (RS) code as used in contemporary DSL systems (cyan circles);
- LDPC encoding with the rate-28/33 code used in upstream DOCSIS 3.1 (purple plus signs);
- LCM encoding with 6 coded bits with the same LDPC code (red crosses).

Here the PCS-LCM scheme also uses 6 coded bits, with the same LDPC code. Compared with the non-shaped LCM, results show that this specific PCS-LCM scheme achieves a shaping gain of 1.1 dB, but even larger shaping gains can be achieved by improved schemes (the theoretical maximum being 1.53 dB). In addition, the SNR gap to capacity of the PCS-LCM scheme is very flat across all the modulation indices, mostly due to the LCM coding scheme.² The total gain of the PCS-LCM scheme is roughly 1.5 dB over conventional LDPC as used in DOCSIS, and roughly 2.6 dB over conventional TCM-RS as used in DSL. Using a hardware proof-of-concept platform, we have observed that the PCS-LCM scheme gives a consistent data rate increase of around 9% over TCM-RS in several twisted pair cables (European operator cables and CAT5e) where the entire range of usable SNR values are sampled.

4. Potential Applications of PCS to the PON Segment in Future FTTx Variants

As described above, probabilistic shaping has been shown as an attractive method to introduce flexible data rates and improve spectral efficiency in both coherent optical transport systems [13, 16, 22] and over copper links [20].

Flexibility and optimizing overall throughput are also considered attractive features for future optical access systems. The most commonly used network architecture in Optical Access is a Passive Optical Network (PON) based on time-division-multiplexing (TDM) in downstream and time-division-multiple-access (TDMA) in upstream to share the bandwidth between all users. A PON is a point-to-multipoint (p2mp) architecture which enables higher equipment density in the optical line terminal (OLT), typically located at the central office (CO), and shares fiber among multiple users to reduce cost. A PON also has a stringent optical power budget due to the passive splitting of the signals between the OLT and users. Another specific PON characteristic is the asymmetrical cost sensitivity, where, in the case of fiber to the home (FTTH), the user equipment is more cost-sensitive because it is not shared over the users of the PON.

Even though the OLT equipment is less cost-sensitive relative to the user side, the power consumption requirement is more stringent putting a limit on complexity. Also, signal reception at the OLT is in burst mode to enable TDMA. This often results in further complexity limitations for optimizing data overhead in the upstream due to burst recovery.

Finally, the p2mp nature of PON results in variations of channel quality for each user mainly depending on their distance from the OLT. Modulation flexibility can make advantageous use of this variation to increase overall throughput and thus optimize cost per bit in a PON.

In the literature, techniques have been proposed to achieve a flexible PON. For example, in [23] a flexible PON is implemented by adaptive bit loading in orthogonal frequency division multiplexing (OFDM), whereby an appropriate QAM modulation format is assigned to each OFDM subcarrier based on its SNR.

² In case of full LDPC encoding, the number of encoded bits increases with the modulation order. Since, for practical non-ideal codes, the information loss increases per coded bit, implying that the gap-to-capacity increases with the information rate. By using LCM, with or without PCS, the number of coded bits stays limited to for instance 6 bits, and there is no additional loss for higher modulation indices.

Reference [24] proposes the use of so-called non-uniform multilevel pulse amplitude modulation (PAM) to enhance flexibility and the aggregated capacity of a PON. Non-uniform multilevel PAM is a combination of applying PAM with unequally spaced levels and interleaving of PON users using the individual PAM-symbol bits which was earlier proposed in [25] as a method to implement PAM modulation in PON with significantly reduced hardware complexity. Optimizing throughput of a PON based on non-uniform PAM with multilevel interleaved users was demonstrated successfully with an actual PON deployment data in [24].

Although probabilistic constellation shaping has to date mostly been applied in coherent networks [13, 22] to introduce flexible data rates, it has recently also been shown to be advantageous to optimize the capacity of direct-detection links based on PAM for data centers [26]. PCS makes use of existing PAM hardware, so it would not add hardware complexity (and thus cost) relative to a flexible PON based on PAM as described above but does have the potential to further optimize spectral efficiency (throughput) and refine data rate granularity compared to the flexible PAM-based PON. The authors of [26] assume soft-decision FEC in their analysis, but [27] shows that hard-decision FEC, which has much lower complexity can also work outstandingly well with PCS. Moreover, PAM-4 has been shown to be technically feasible and cost-effective for a line rate of 25 Gbps for residential PON (which has the most stringent optical power budget) when optical amplification is applied at the OLT [28].

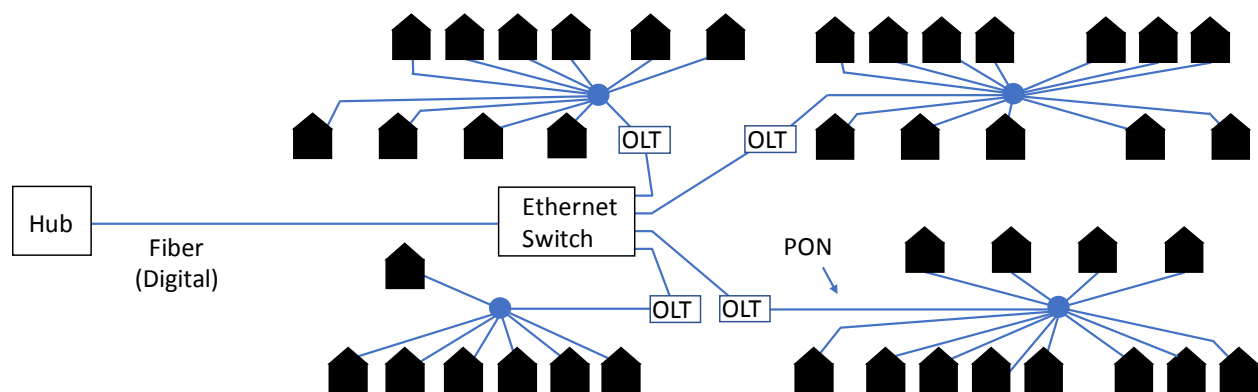


Figure 8 – Fiber deeper Distributed Access Architecture (DAA) with the aggregated digital optical link terminated on an Ethernet switch. Lower rate optical Ethernet links connect the switch to PON OLTs, that are in turn connected to user ONUs (not shown).

Applications beyond FTTH, such as mobile backhaul, mobile fronthaul, and the backhaul of DSL nodes and DOCSIS RPDs (see Figure 6) will benefit from a PON technology that adapts to a range of network topologies, including shorter reach and lower split-ratio and therefore smaller optical budgets. For these shorter and/or smaller PONs, modulation orders beyond PAM-4 will then also become feasible which further enlarges the total achievable data rate [26]. Figure 8 shows an example architecture from the recent CableLabs architectural specification [19], where the coaxial network between the RPDs and users, similar to that shown in Fig. 5, is replaced with PONs.

In summary, flexible PON based on using different modulation orders of PAM, combined with PCS and multi-level user-interleaving with non-uniform PAM, while applying direct-detection and hard-decision FEC, is a valid solution for future flexible PON use cases like high speed residential access and next generation x-haul. It has been shown that these techniques can also optimize overall throughput of the PON, which reduces the cost per bit of this very cost-sensitive network.

Conclusion

In recent years, probabilistic constellation shaping has become a very promising digital signal processing technique thanks to crucial advances in implementation architectures. It allows a reduction of the gap to reach the Shannon capacity limit and the elimination of excess margins (capacity loss) introduced by the coarse granularity of QAM constellations used for signal modulation, increasing the achievable data rates on communication links. We have demonstrated its potential in coherent optics and copper access communication technologies, and see further potential for PCS in PON technologies as used in FTTH networks and/or for backhaul of access technologies.

4. Acknowledgements

The authors gratefully acknowledge Andrew Chraplyvy, Vincent Houtsma, Marty Glapa, Jochen Maes, Laurent Schmalen, and Peter Winzer for helpful contributions.

Abbreviations

CAT5e	Category 5e twisted copper cable
CO	Central office
CPM	Cross phase modulation
CSO	Composite second order
CTB	Composite triple beat
DAA	Distributed access architecture
DOCSIS	Data over cable service interface specification
DQPSK	Differential quadrature phase shift keying
DSL	Digital subscriber line
DWDM	Dense wavelength-division multiplexing
FEC	Forward error correction
FN	Fiber node
FTTH	Fiber-to-the-home
FWM	Four-wave mixing
HE	Head-end
HFC	Hybrid fiber-coaxial
ITU	International Telecommunication Union
LCM	LDPC-coded modulation
LDPC	Low-density parity-check
MB	Maxwell-Boltzmann
MSO	Multiple system operator
OFDM	Orthogonal frequency division multiplexing
OLT	Optical line terminal
ONU	Optical network unit
OSNR	Optical signal-to-noise ratio
P2P	Point-to-point
P2MP	Point-to-multipoint
PAM	Pulse amplitude modulation
PCS	Probabilistic constellation shaping

PMF	Probability mass function
PON	Passive optical network
QAM	Quadrature amplitude modulation
RF	Radio frequency
RMD	Remote MACPHY device
RPD	Remote PHY device
RS	Reed-Solomon
SNR	Signal-to-noise ratio
TCM	Trellis coded modulation
TDM	Time-division multiplexing
TDMA	Time-division multiple access
WDM	Wavelength-division multiplexing

Bibliography & References

1. T. E. Darcie and G. E. Bodeep, "Lightwave Subcarrier CATV Transmission Systems," *IEEE Trans. on Microwave Theory and Techniques*, Vol. 38, No. 5, pp. 524-533, 1990.
2. S. L. Woodward and M. R. Phillips, "Optimizing Subcarrier-Multiplexed WDM Transmission Links," *J. Lightwave Technol.*, Vol. 22, No. 3, pp. 773-778, 2004.
3. Cable Television Laboratories Specification, "Data-Over-Cable Service Interface Specifications, DCA-MHAv2, Remote PHY Specification," May 2018
4. A. R. Calderbank and L. H. Ozarow, "Nonequiprobable Signaling on the Guassian Channel," *IEEE Trans. on Inf. Theory*, Vol. 36, No. 4, pp. 726-740, 1990.
5. C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Tech. Journal*, Vol. 27, No. 3, pp. 379-423, 1948.
6. ITU-T Recommendation V.34, "A modem operating at data signaling rates of up to 33 600 bit/s for use on the general switched telephone network and on leased point-to-point 2-wire elephone-type circuits," Feb. 1998.
7. G. Böcherer, F. Steiner, and P. Schulte, "Bandwidth Efficient and Rate-Matched Low-Density Parity-Check Coded Modulation," *IEEE Trans. on Communication*, Vol. 63, No. 12, pp. 4651-4665, 2015.
8. <https://networks.nokia.com/photonic-service-engine-3>
9. R. J. Essiambre, G. Kramer, P. J. Winzer, G. J. Foschini, and B. Goebel, "Capacity Limits of Optical Fiber Networks," *J. Lightwave Technol.*, Vol. 28, No. 4, pp. 662-701, 2010.
10. F. R. Kschischang and S. Pasupathy, "Optimal nonuniform signaling for Gaussian channels," *IEEE Trans. Inform. Theory*, vol. 39, no. 3, pp. 913-929, 1993.
11. G. D. Forney Jr., R. G. Gallager, G. R. Lang, F. M. Longstaff, and S. U. Qureshi, "Efficient modulation for band-limited channels," *IEEE J. Select. Areas Commun.*, vol. SAC-2, no. 5, pp. 632-647, 1984.

12. G. D. Forney Jr., "Trellis shaping," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 281–300, 1992.
13. F. Buchali, F. Steiner, G. Böcherer, L. Schmalen, P. Schulte, and W. Idler, "Rate Adaptation and Reach Increase by Probabilistically Shaped 64-QAM: An Experimental Demonstration," *J. Lightwave Technol.*, Vol. 34, No. 7, pp. 1599-1609, 2016.
14. S. Chandrasekhar et al., "High-spectral-efficiency transmission of PDM 256-QAM with parallel probabilistic shaping at record rate-reach trade-offs," *Proc. ECOC*, 2016, paper Th.3.C.1.
15. A. Ghazisaeidi et al., "65Tb/s transoceanic transmission using probabilistically shaped PDM-64QAM," *Proc. ECOC*, 2016, paper Th.3.C.4.
16. J. Cho et al., "Trans-Atlantic field trial using high spectral efficiency probabilistically shaped 64-QAM and single-carrier real-time 250-Gb/s 16-QAM," *J. Lightw. Technol.*, vol. 36, no. 1, pp. 103-113, 2018.
17. S. L. I. Olsson et al., "Probabilistically shaped PDM 4096-QAM transmission over up to 200 km of fiber using standard intradyne detection," *Opt. Exp.*, vol. 26, no. 4, pp. 4522-4530, 2018.
18. Cable Television Laboratories Specification, "P2P Coherent Optics Physical Layer 1.0 Specification," June 2018.
19. Cable Television Laboratories Specification, "P2P Coherent Optics Architecture Specification," June 2018.
20. Nokia Corporation, "G.mgfast: Probabilistic amplitude shaping (PAS) for G.mgfast" ITU-T Q4/15 Contribution T17-SG15RGM-Q4-171127-C-0025, November 2017.
21. E. Eleftheriou and S. Ölçer, "Low-Density Parity-Check Codes for Digital Subscriber Lines", *IEEE International Conference on Communications 2002*, New York, USA, April 2002.
22. T. Fehenberger, et al., "LDPC coded modulation with probabilistic shaping for optical fiber systems," *Proc. OFC*, Th2A.23 (2015).
23. L. Zhou et al., "Demonstration of software-defined flexible-PON with adaptive data rates between 13.8 Gb/s and 5.2 Gb/s supporting link loss budgets between 15 dB and 35 dB," 2014 The European Conference on Optical Communication (ECOC), Cannes, 2014.
24. R. van der Linden, N. C. Tran, E. Tangdiongga and T. Koonen, "Optimization of Flexible Non-Uniform Multilevel PAM for Maximizing the Aggregated Capacity in PON Deployments," in *Journal of Lightwave Technology*, vol. 36, no. 12, pp. 2328-2336, June15, 15 2018.
25. V. Houtsma, D. van Veen, and H. Chow, "Demonstration of symmetrical 25 Gb/s TDM-PON with multilevel interleaving of users," in *Journal of Lightwave Technol.*, vol. 34, no. 8, pp. 2005–2010, Apr. 2016.
26. T. A. Eriksson, M. Chagnon, F. Buchali, K. Schuh, S. ten Brink and L. Schmalen, "56 Gbaud Probabilistically Shaped PAM8 for Data Center Interconnects," *2017 European Conference on Optical Communication (ECOC)*, Gothenburg, 2017.

27. A. Sheikh, A. G. i. Amat, G. Liva and F. Steiner, "Probabilistic Amplitude Shaping With Hard Decision Decoding and Staircase Codes," in *Journal of Lightwave Technology*, vol. 36, no. 9, pp. 1689-1697, May 1, 2018.
28. D. T. van Veen and V. E. Houtsma, "Symmetrical 25-Gb/s TDM-PON With 31.5-dB Optical Power Budget Using Only Off-the-Shelf 10-Gb/s Optical Components," in *Journal of Lightwave Technology*, vol. 34, no. 7, pp. 1636-1642, April 1, 2016.

Internet of Things Dynamics: Opportunities and Challenges for Broadband Network Operators

A Technical Paper prepared for SCTE•ISBE by

Tim Johnson

Director, Global Product Management
Alpha Technologies
360-392-2234

Tim.Johnson@alpha.com

Arun Ravisankar

Sr Engineer
Comcast
215-286-7558

Arun_Ravisankar@cable.comcast.com

J. Clarke Stevens

Principal Architect
Shaw Communications
587-393-0605

Clarke.Stevens@sjrb.ca

Chris Bastian

SVP/CTO
SCTE-ISBE
610-594-7304

cbastian@scte.org

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
1. The IoT Service Opportunity	3
1. IoT Use Case Descriptions	4
1.1. Home Monitoring.....	5
1.1.1. Home Automation	5
1.1.2. Home Security.....	5
1.2. Connected Healthcare	6
1.3. Smart Cities/Mobility	9
2. IoT Perspectives from Network Operators: Developing, Operating and Maintaining Consumer IoT Services	13
3. Technical factors	14
3.1. Protocols and Standards.....	14
3.2. Security threats and security solutions/best practices	15
3.3. Operational factors.....	16
3.3.1. Training the workforce to install and operate a network of home-based IoT sensors and objects	16
3.3.2. IoT supporting IoT	16
4. Conclusion.....	17
5. Abbreviations.....	17
6. Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1: Evolution of Cable Beyond the Box	3
Figure 2: Internet of Things Value Add by 2020	4
Figure 3: Devices in a Home Monitoring Use Case	5
Figure 4: Home Security System Flowchart.....	6
Figure 5: Activity Monitoring Applications	8
Figure 6: Biometric Devices used for Remote Patient Monitoring	9
Figure 7: Trend Capture and Monitored Data	9
Figure 8: Network architecture supporting Security Cameras	10
Figure 9: Network architecture supporting LoRaWAN.....	11
Figure 10: Demonstrating Vehicle to Infrastructure (V2I) and subsequent Vehicle to Vehicle (V2V) communications via DSRC and the network of Road Side Units (RSUs).....	12
Figure 11: Traditional RSU deployment.....	12
Figure 12: IoT connected device growth forecast.....	14
Figure 13: Network Protocols supporting IoT applications	15

Introduction

Cable network operators are always looking for ways to add services for their customers, especially so since the 1990s. To name a few: Data over cable, then voice, then DVR evolving to nDVR, were added to the service bundle. More recently home monitoring and security services have also been offered by many operators. [1]




Figure 1 - Evolution of Cable Beyond the Box

Source: calcable.org

The Internet of Things is loosely defined as Internet-connected sensors in homes, businesses and public spaces, as well as the data analytics monitoring of those sensors back in the data center. With the Internet of Things, there is an opportunity to rapidly open up entirely new service opportunities that can differentiate cable network operators from their competition. However, the primary challenge will be to smoothly install, operate and integrate these new services with the operator's existing service bundle.

Cable network operators are uniquely positioned to offer IoT services to new and existing customers. They have four characteristics that industry start-ups and OTT service providers covet:

- Existing service location in millions of homes, businesses, and public spaces
- High speed and reliable network connectivity
- Power for sensors and gateways
- An existing and localized/in-market fleet of fulfillment technicians

Cable network operators have a well-established presence in the home including cable modems, home gateways, set top boxes, Wi-Fi extenders and home security hubs, however the evolution to new services - such as connected healthcare, and smart homes - will require new devices and sensors, as well as increased care to ensure the highest network performance while preventing security breaches.

1. The IoT Service Opportunity

Gartner, Inc. forecasts that 8.4 billion Internet-connected things were in use worldwide in 2017, up 31 percent from 2016, and will reach 20.4 billion by 2020. Total spending on endpoints and services will reach almost \$2 trillion by 2020. [2]

Internet of Things Value Add by 2020




Figure 2: Internet of Things Value Add by 2020

How big of this market share will the cable network operators capture? The service cases are expanding, as are the number of customers. Will the cable industry focus on these service sectors? Recent press releases seem to indicate that they will. Comcast launched its low power, wide area network, machineQ, in July 2017 [3], and Cox followed suit with Cox2M in March 2018 [4], each focused on rapidly introducing IoT services to their customers.

The SCTE Data Standards Subcommittee, and Internet of Things Working Group



Established in 1996, the Data Standards Subcommittee (DSS) develops standards for the delivery of digital service supporting high-speed data, video, VoIP, and other services over cable networks.

The Internet of Things working group was announced on October 31, 2016 as an entity under DSS, and conducted its first meeting on November 16, 2016. [5]

The IoT working group's charter is to facilitate communication between service providers and industry partners to standardize new IoT-based services. The working group aims to make standards and operational practices deployable and manageable for service providers, as well as focusing on the vast use cases available in the IoT community to support service providers' business objectives. The early focus of the working group was to develop use case descriptions, which will be outlined in the following sections.

1. IoT Use Case Descriptions

One of the leading categories in IoT use cases is in home monitoring, and specifically home automation, home security, and connected healthcare. This section details those use cases.

1.1. Home Monitoring

The Home Monitoring use case mainly spans two broad sets of applications, described here.

1.1.1. Home Automation

Home automation applications provide services that augment the capabilities of devices and sensors in a home, and help customers to seamlessly access these services. Automation is intended to provide ease of use and access to products and services, while simultaneously helping with overall energy conservation.

Home Automation may include managed devices, like Door/Window Sensors, Motion Detectors and the building of rules engines that could operate and control other devices. For example, some use cases could be to detect motion, operate lights/STB/TV, or to open/close a garage door.

Home automation also plays an important role in providing safety-related features, like smoke detectors and flood sensors. These sensors can trigger actions which include raising alarms for appropriate help as needed.

Home automation applications include integration with personal voice assistants, and other smart devices in the home that are capable of connecting to a network and exposing APIs to control their actions. Examples include smart speakers, thermostats, washers, bulbs and many more. Figure 3 captures many home monitoring devices.



Figure 3 - Devices in a Home Monitoring Use Case

1.1.2. Home Security

Home Security, as the name implies, is designed to secure the premise with the use of IoT sensors. Customers could opt in to receive alerts when certain anomalies detected. Additionally, a third-party service can provide verification of any events or activities, and possibly contact the user or law enforcement based on the indications/events.

The evolution of IoT and smart home applications has increased the demand for residential security products. Apart from traditional security requirements, IoT security is also gaining importance to ensure

data and device integrity in a smart home. Various research findings suggest an increased need for home security systems to reduce burglary-related emergencies. In most cases where burglaries have been reported, one of the major causes is that the home owner has forgotten or neglected to close doors or windows. A typical home security system, when set in “armed” mode, will detect any such anomalies and alert the user. This is a huge relief for the home owner, and illustrative of the tacit benefit that is peace of mind.

It is important for any Home Security system to have redundancy in backhaul connectivity in case there is an attempt to sabotage the primary internet connection, which is the broadband cable internet. Home Security systems tend to have backup cellular connectivity in case there is a drop in broadband connectivity.

Figure 4 shows the behavior of a home security system when an event occurs. The system, when in “Armed” mode, detects if any door/window is opened and alerts the user. However, there are efforts underway to use AI and Machine Learning tools to analyze data from all sensors in a home to determine if there is no one in home and if the customer has forgot to “Arm” the system.

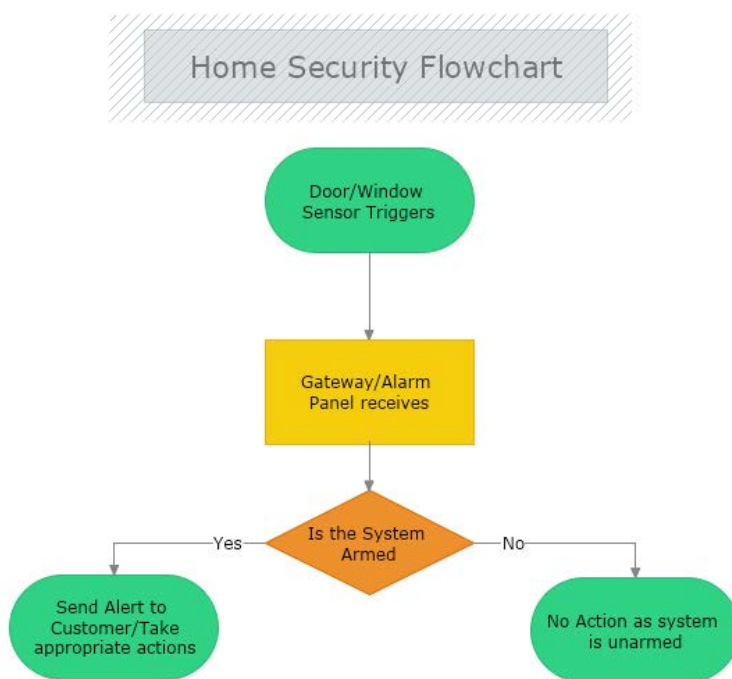


Figure 4: Home Security System Flowchart

1.2. Connected Healthcare

IoT technologies can impact a lot of healthcare use cases and can bring peace of mind to customers. These technologies could also assist Medicare professionals by augmenting them with information when providing medical care to patients.

Health and wellness applications, generic or specific, could target a large set of population. Examples include fitness-based applications that help consumers track and organize their health data, and to plan daily activities and exercise routines.

Another major area where IoT could play a major role is in eldercare. Studies show that there is significant growth in the senior population in the coming years, as the Baby Boomer generation enters its twilight years. IoT is poised to offer applications that facilitate aging-in-place -- because personal independence is a high priority for Senior citizens.

Healthcare use cases include the following:

- *Home Health and Tele-Medicine Applications:* Bridging between FDA-approved home health devices (blood pressure cuffs, glucometers, asthma dispensers, etc.) and medical professionals/caregivers, via Bluetooth-to-LPWAN adaptors.
- *Health and Wellness applications:* To meet and extend people's health/wellness goals.
- *Remote Patient Monitoring:* To extend the reach of Internet-connected devices by bridging between short- and long-range networks.
- *Aging in Place applications:* To extend the range of ADL (Activities of Daily Living) that are often stressed by age-related consequences, such as falling.
- *Responding to Emergencies:* To make it faster and easier to get urgent care.

Various studies indicate a gradual increase of Internet users among people aged 75 years and older, as well as an increased level of smartphone ownership. This growing segment of connected Seniors similarly indicates that IoT could play a major role in enhancing lifestyles. Notably, the population of those aged 65 and over has increased from 36.2 million in 2004, to 46.2 million, in 2014 -- a 28% increase. The Senior population is projected to more than double, to 98 million, by 2060. [6]

As Internet adoption increases within the Senior community, it represents a major tool for providing value-based services:

- Applications like ADL monitoring and remote patient monitoring would help reduce the number of visits to a care provider. This would also help reduce the burden on conventional healthcare systems (such as hospitals and clinics.)
- Medicare has begun implementing incentives to reduce hospital re-admissions, which has stimulated the growth of remote patient monitoring. Efforts like the Hospital Readmission Reduction Program (HRRP) actually penalizes hospitals, financially, if they exhibit high rates of Medicare readmissions. [7]
- Remote patient monitoring could be used to source and convey health and wellness information, so as to:
 - Provide first-hand information to users and care providers (family members, medical personnel)
 - Encourage patient adherence to medical protocols (medications, exercise)
 - Enable caregivers and medical providers to plan and adjust the course of action
 - Help individuals to make lifestyle choices fueled by individualized data
- Apart from eldercare, these technologies could also be tailored to assist patients or otherwise vulnerable family members with chronic conditions.

Connected health applications offer services to consumers and caregivers (professionals & personal/family members.) These applications provide a platform on which the patient and caregiver could interact, exchange data and configure alerts. Such services would provide appropriate information and alerts to a family member or care provider, so that corrective action could be taken. For instance, it is an invaluable peace of mind for a son or daughter who no longer lives near an aging parent, to know that

medications are being taken as scheduled, or that something abnormal or problematic is happening (or, preferably, not happening!)

Remote patient monitoring typically includes:

- **Activity Monitoring:** Tracking activities and detecting abnormal or emergency situations, like a fall event or an incapacitation. Fall detection or incapacitation could be used to trigger a PERS (Personal Emergency Response Systems) event. Figure 5 shows how activity monitoring applications could be used to monitor activity of elders for an aging-in-place application. These could also be used for fall detection and raise alerts for help when such an incident occurs.
- **Biometric Monitoring:** Measuring body vitals, like blood sugar, BP (blood pressure), weight; establishing a secure health data record which is monitored constantly; predicting future anomalies; reducing risks. Figure 6 shows examples of devices that are used to monitor body vitals. These devices are BLE- (Bluetooth Low Energy) enabled and data could be sent and analyzed by medical care providers.
 - Data collected could be put into analytic engines and trends could be analyzed. Figure 7 shows a sample trend related to Blood Pressure.
 - These trends indicate the progress and well-being of the patient.
 - Use of AI/ML combined with IoT technologies could help bring *Sensors to Insights*
- **Patient Adherence:** Ensuring that patients follow their doctor's orders; providing appropriate reminders to patients and family members, such as for medication refills.
- **Virtual Visits:** Meeting doctors or care providers over a video conference, rather than a face-to-face meeting, which with Seniors often involves collapsible wheel chairs and a considerable amount of extra effort for everyone involved.
- Other applications include access to electronic health records (EHRs), to help doctors and care providers optimize a patient's health with vital information.

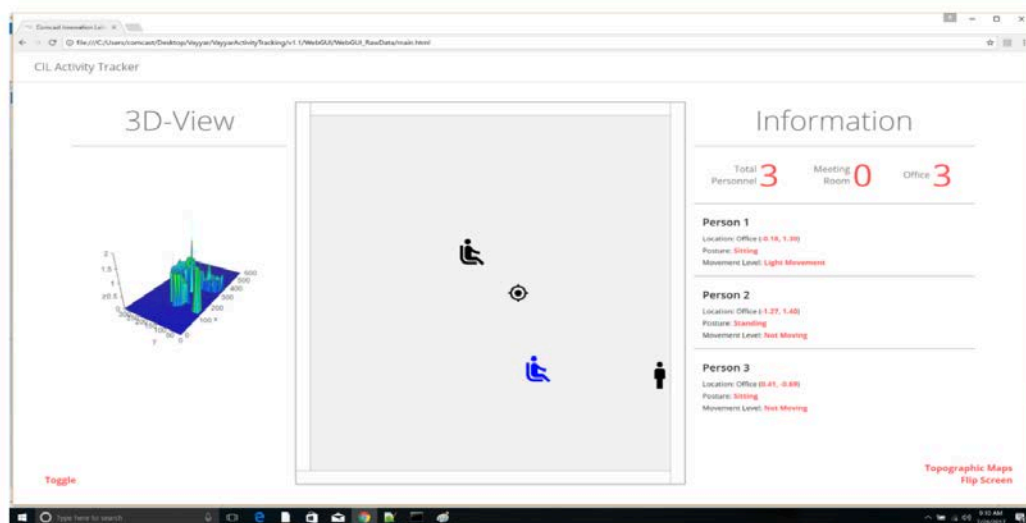


Figure 5: Activity Monitoring Applications



Figure 6: Biometric Devices used for Remote Patient Monitoring



Figure 7: Trend Capture and Monitored Data

1.3. Smart Cities/Mobility

In December 2015, the U.S. Department of Transportation (U.S.D.O.T.) launched the *Smart City Challenge*, asking mid-sized cities (200,000 – 800,000) across America to share their ideas for how to create an integrated, first-of-its-kind smart transportation system that would use data, applications, and technology to help people and goods move faster, cheaper, and more efficiently. [9]

By challenging American cities to use emerging transportation technologies to address their most pressing problems, the Smart City Challenge aimed to spread innovation through a mixture of competition, collaboration, and experimentation. But the Smart City Challenge was about more than just technology. The U.S. Department of Transportation (USDOT) called on mayors to define their most pressing transportation problems and envision bold new solutions that could change the face of transportation in U.S. cities by meeting the needs of residents of all ages and abilities; and *bridging the digital divide* so that everyone, not just the tech-savvy, can be connected to everything their city has to offer.

Of the approximately 100 cities which fit this criteria, 78 quickly activated cross-organizational (City; State; Private Enterprise; NFP, etc.) teams to submit comprehensive applications for the \$40M which the USDOT put forth for the winning bid. [9]

The powerful two-fold message of this fact:

1. The USDOT recognizes that the “Smart Transportation/Mobility” is at the core of the Government’s perspective on what is foundational for a “Smart City”.
2. The vast (78%!) majority of cities either have, or are rapidly seeking to have, leadership and/or funding in place to immediately enact Smart City initiatives.

The relevance (“so what”) for the Cable Industry, MSOs and supporting eco-system?

In order to enact many of the use cases which are associated and envisioned for “Smart Mobility/Smart Cities”, there are fundamentally three core infrastructure attributes which are required:

1. Power
2. Communications Backhaul
3. Real Estate (site)

Enter the HFC Network. Given the near ubiquity of this network throughout America (either aerial or subterranean) the enablement of strand-mounted “gateway” devices provides the opportunity for rapid and cost-effective deployment of several types of Smart Mobility-affecting IP-based devices, including, but not limited to: small cells (NB-IoT); Wi-Fi access points; IoT (LoRa), and Security Cameras.

The diagrams below demonstrate the architecture, based on actual deployments, of two of these areas:




Figure 8: Network architecture supporting Security Cameras




Figure 9: Network architecture supporting LoRaWAN

Question:

In addition to these “ready to deploy” Smart Mobility/Smart City technologies, what could be the next BIG thing that the HFC could enable across the country?

One Answer:

Digital Short Range Communications (DSRC) to promote safer, more intelligent transportation systems.

DSRC is a two-way short-to-medium-range wireless communications capability that permits very high data transmission critical in communications-based active safety applications. The Federal Communications Commission set aside 75 MHz of spectrum around the 5.9 GHz band (5.850-5.925 GHz) band in 1999 to be used for vehicle-related safety and mobility systems. [9]

The USDOT has identified more than 40 use cases for vehicle to infrastructure (V2I) technologies, such as:

- the ability to pay for parking and tolls wirelessly
- identifying when a car is approaching a curve too quickly and alerting the driver
- adjusting traffic signals to accommodate first responders in an emergency; and
- alerting drivers of conditions such as road construction, among many others




Figure 10: Demonstrating Vehicle to Infrastructure (V2I) and subsequent Vehicle to Vehicle (V2V) communications via DSRC and the network of Road Side Units (RSUs)

The current major roadblock in deploying the DSRC wireless network? The need for efficiencies in deploying Roadside Units, or RSUs, relative to three core attributes: Power, Backhaul, and Real Estate.

Current RSU deployment topology is focused on “traditional” sources of these needed ingredients, as portrayed below using a combination of Utility Power; available Street Furniture (such as traffic lights), and, more often than not, Wireless Backhaul:

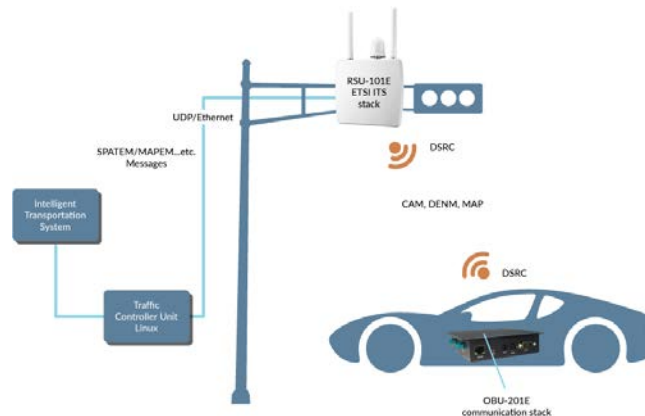


Figure 11: Traditional RSU deployment

This conventional wisdom is largely driven by a lack of general understanding of the value proposition offered here:

1. The HFC Network provides a high-speed, QoS-based, communications network, with diverse, backed-up power, provided via strand-mounted gateways which are deployed on “real estate” in an almost unlimited footprint of deployment choices.
2. RSU Devices are generally powered by PoE+ (under 50W load), and transmit approximately 300 Yards.

The *new message* which could/should be promoted to all stakeholders:

The HFC Network provides an ideal solution to quickly and efficiently deploy RSUs in the scale required to effectively realize the dream the USDOT envisioned in 1999...*Safer and Smarter* (more Intelligent) transportation systems for citizens across the country.

In summary, in order to enact a significant core of the most cutting-edge technologies and resulting Smart Cities/Mobility use cases, one of the most iconic and legacy infrastructures of the last century could and should be quickly re-energized to perform a vital 21st century service:

The over 1.5 million miles of America's HFC Network.

2. IoT Perspectives from Network Operators: Developing, Operating and Maintaining Consumer IoT Services

Estimates for the market size of IoT vary greatly. Most estimates show strong growth of around 20% year-over-year. Other predictions aren't nearly so conservative. One thing is for sure, almost all analysts see the market as big and growing. The variability is due in part to the virtually unlimited perspectives of the market. The smart home is perhaps the most visible manifestation of the IoT market to most people. (People have always talked to their appliances, but it's hard to avoid the wonder of having the appliances listen.) It seems that the home market is probably one of the less profitable sectors. Corporations have a financial incentive to use IoT to increase productivity and they have more cash to invest. It's true that knowledge is power and by that measure IoT is a revolution. It's not only feasible to know virtually everything about your product and its production instantaneously, the sheer volume of data and the automated analysis of that data allow for insights that a human is never likely to discover.

The good news is that this information and control is available to operators at a reasonable cost. The bad news is that it's available to your competitors at that same cost. So how can IoT be advantageous to cable? The key is to look at the differentiation cable has built through decades of intense investment in a service industry:

- Cable has a monthly financial relationship with the customer.
- Cable has a fleet of skilled technicians who visit customers in person at their homes.
- Cable has equipment in customers' homes and the physical plant to reach those homes. That equipment is connected to the headend twenty-four hours a day.

These advantages, however, come with challenges. A monthly bill has never been attractive to the customer paying it and increasingly, they are presented with options that are based on actual consumption.

Our skilled technicians are often not qualified in the skills that customers want. Engineers are not always great salespeople. The sheer volume of IoT devices makes it almost impossible to know which options will best serve the customer.

The physical plant is always in need of an upgrade and increasingly customers are choosing the freedom of wireless infrastructure that is also striving for continuous reliability improvements.

In order to understand how the cable infrastructure can be effectively leveraged, it is critical to look at customer pain points. Some of the most important include:

- *Lack of interoperability* – Customers are looking for solutions to their challenges. Those solutions often come from different vendors who rely on different ecosystems. This diversity is never going away, but that is a problem that providers must address. Consumers can't.
- *Security* – There is legitimate fear around the security of the network and IoT solutions. As IoT devices become more ubiquitous and easier to use, they pose an increasingly attractive resource for bad actors.
- *Management* – Most customers are unqualified to be system administrators and have no desire to assume that role, just as most drivers are unqualified to be auto mechanics. Operators need to find a way to economically provide this service.

The market is indeed big and growing. [10] Cable does have some intrinsic assets that provide an advantage. However, that advantage is a head start, not a reservation. MSOs need to quickly establish the efficient infrastructure to provide the best and most responsive service to customers while leveraging the very features of IoT to run that infrastructure efficiently.




Figure 12: IoT connected device growth forecast

Source: Statista

3. Technical factors

3.1. Protocols and Standards

IoT applications use a wide variety of signaling protocols and standards. Most IoT devices today are based on the following communication protocols:

- Zigbee
- Bluetooth Low Energy (BLE)
- Wi-Fi-based devices

Figure 13 shows a typical representation of protocols that are generally associated with IoT applications. [11]

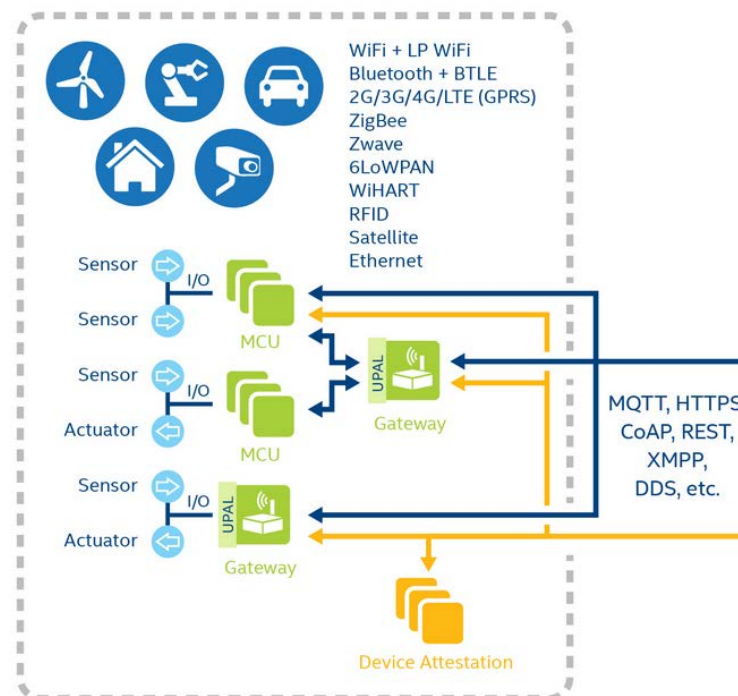


Figure 13: Network Protocols supporting IoT applications

As part of the application design, it is important that the service provider infrastructure supports these protocols and hence serve a large number of devices that connect to the network.

One of the major design aspects is the coexistence of all these protocols and how to best bridge them on to a network, as well as reliably transfer data for analytics and processing. Challenges include working with multiple protocols, hardware and software implementations. Signal interference is another challenge which leads to data loss and may impact some critical functioning of the system.

The action of “rules engines” and devices depends heavily on how each device in the system (home) functions, irrespective of its type, power source, or data model. It is very important to monitor all devices and to ensure all of the devices have the right configurations (software, settings, and battery level).

3.2. Security threats and security solutions/best practices

Systems and applications need to be secured from unauthorized access, and to protect sensitive information. Connected devices lead to an increased risk of being exposed to attacks, and as such need elaborate security measures to provide these safe and secure services to the customers.

Incidents are routinely reported where peripheral devices in an IoT network have been compromised by DDoS attacks.

The attack surface of IoT devices are greatly increasing as operating systems and communication stacks are becoming more complex. [12] Further, the IoT-related firmware typically includes open-source components, and code written by large distributed teams. It is the service provider’s goal to ensure that

the number of penetration points and attack vectors are minimized for its deployed IoT devices. However, if the attacker manages to penetrate the device, the impact to other devices in the network should simultaneously be minimized.

Trusted devices, with a Root of Trust, provide additional security. Unmanaged devices could be monitored for the data patterns being sent and received by the device. For example, a thermostat which is expected to receive settings and report status could be suspicious if the device is suddenly observed to be generating heavy data traffic and communicating with unintended destinations.

Industry groups like OCF (Open Connectivity Foundation) help in specifying device attributes and helping monitor the network status.

3.3. Operational factors

3.3.1. Training the workforce to install and operate a network of home-based IoT sensors and objects

It is very important to effectively train the customer-facing workforce to address issues relating to IoT services. Conventionally, the service provider's call center receives calls when there is an issue with any of the devices (including unmanaged devices) in the premise. Our customer care professionals need to be equipped with the necessary tools to understand the problems associated with IoT services, so as to authentically help the customer in resolving the problem.

Another important aspect is installing and configuring devices on the network. Professional installation involves finding the right spots in the home to install devices like motion sensors and other IoT devices. The technician needs to consider such things as sensor range, number of sensors needed, and interference issues. The technician needs to be equipped with the tools which can help them effectively troubleshoot issues with both installation and maintenance.

Device configurations and software also need to be managed that include new feature updates and bug fixes.

3.3.2. IoT supporting IoT

One of the best ways to support a robust IoT infrastructure is to use IoT in that support. IoT can be viewed from two sides: One side looks at the benefits of connecting real-world devices to the Internet so that they can be monitored and controlled. The other side looks at the data that a constellation of IoT devices generates and sees that as a rich resource for understanding and improving IoT service.

There are the obvious observations of knowing when equipment is on and off, and the direct information that can be queried from an individual device. This sort of information is good for customer service personnel and can be directly leveraged on a customer call.

What may be more valuable, however, are inferences that can be made on an aggregate basis of massive amounts of data that are statistically evaluated to get insights that are less obvious. This information is yet another tool that can be used in the ongoing field of proactive network maintenance. While there are many benefits to the science of using tuned signals to isolate the location of a network fault, that information becomes even more valuable if the traffic traversing that network can be dynamically understood. If active elements in the network can be remotely controlled to instantly repair or avoid the breach automatically and without an immediate truck roll, the savings quickly accrue. If data can be used to provide a predictive diagnosis, potential problems can be avoided altogether.

IoT makes this possible by providing intelligence to all equipment connected to the network with a big data analytics infrastructure to collect, analyze and actuate the IoT network elements. Network operators are unlikely to be the creators of the various IoT devices on the network. However, they can define a common network infrastructure that will improve manageability for the operator and provide a target platform for vendors. For maximal efficiency, this platform should have the following features:

- *Interoperability* – While the different advantages of various network technologies ensure that there will never be one network that everyone agrees upon, the Internet has shown us that different networks can be made highly interoperable. The cable industry should adopt an upper-level IoT network that allows for these different technologies underneath, while providing the commonality required to make the technologies work together.
- *Security* – It is clear that operators will be expected by customers to insure the integrity of the network and a safe environment for their data. If MSOs will have to address any breaches, it makes far more sense to prevent those breaches in the first place.
- *Standards* – The challenge of providing a reliable, interoperable, and secure IoT platform is not something that can be done in isolation. Just as DOCSIS is responsible for the industry's broadband vibrancy, the next generation of the cable industry will rely on a standardized IoT infrastructure that can support consumers, businesses, government and industry. It will also benefit itself by providing the information and equipment to significantly maintain itself autonomously. The Internet of Things must be supported by international standards that address not only the needs of customers, but also the needs of operators, equipment makers, chip providers and every participant in the IoT chain.

4. Conclusion

The Internet of Things is viewed as many different things by different participants. However, one thing is clear: For cable operators, it is not simply a new product category, it is a new platform for products and services. Cable has been through this before. While many predicted the demise of cable, they neglected to acknowledge cable's evolution. Indeed, if cable were still based on analog television services, it would be irrelevant. But that's not what happened. Cable evolved to digital services, changed the basic structure of its network and became the premier operator of Internet access. Now it's time to evolve again to be the premier platform for the Internet of Things – a common platform that will serve consumers, businesses, government, industry and itself.

5. Abbreviations

ADL	Activities of Daily Living
AI	Artificial Intelligence
AP	Access point
BLE	Bluetooth Low Energy
DOCSIS	Data over coax service interface specification
DSRC	Digital short range communications
DSS	Data Standards Subcommittee
EHRs	electronic health records
HFC	Hybrid fiber coax
HRRP	Hospital Readmission Reduction Program
IoT	Internet of Things
ISBE	International Society of Broadband Experts
LoRaWAN	Long range wide-area network
LP-WAN	Low-power wide-area network

ML	Machine Learning
NB-IoT	Narrowband Internet of Things
nDVR	Network digital video recorder
OCF	Open Connectivity Foundation
OTT	Over The Top
PERS	Personal Emergency Response Systems
QoS	Quality of Service
RSU	Road side unit
SCTE	Society of Cable Telecommunications Engineers
STB	Set-Top Box
USDOT	United States Department of Transportation
V2I	Vehicle to infrastructure
V2V	Vehicle to vehicle

6. Bibliography & References

- [1] <https://www.cable.org/>
- [2] <https://www.gartner.com/en/newsroom/press-releases/2017-02-07-gartner-says-8-billion-connected-things-will-be-in-use-in-2017-up-31-percent-from-2016>
- [3] <https://corporate.comcast.com/news-information/news-feed/machineq-comcasts-enterprise-internet-of-things-service-expanding-to-12-major-us-markets>
- [4] http://newsroom.cox.com/cox_launches_cox2m_for_smart_cities_smart_businesses
- [5] <https://www.multichannel.com/news/scte-is-be-eyes-iot-standards-408241>
- [6] <https://www.census.gov/content/dam/Census/library/publications/2015/demo/p25-1143.pdf>
- [7] <https://www.cms.gov/medicare/medicare-fee-for-service-payment/acuteinpatientpps/readmissions-reduction-program.html>
- [8] <https://www.transportation.gov/smartcity>
- [9] https://www.its.dot.gov/pilots/pilots_thea.htm
- [10] <https://www.statista.com/statistics/764051/iot-market-size-worldwide/> and <https://www.statista.com/study/27915/internet-of-things-iot-statista-dossier/>
- [11] <https://newsroom.intel.com/news-releases/intel-unifies-and-simplifies-connectivity-security-for-iot/>
- [12] <https://www.techrepublic.com/article/as-iot-attacks-increase-600-in-one-year-businesses-need-to-up-their-security/>

Internet of Things Security: Implement a Strong, Simple & Massively Scalable Solution

A Technical Paper prepared for SCTE•ISBE by

Ron Ih

Director of Business Development, Kyrio Security Solutions
Kyrio
r.ih@kyrio.com

Table of Contents

Title	Page Number
Table of Contents	2
1. Executive Summary	3
2. Scope of the Problem.....	4
2.1 Device Manufacturers.....	4
2.2 Semiconductor Manufacturers	5
2.3 Cloud Service Providers	6
2.4 Security Infrastructure Providers	6
2.5 Scope of the Problem—Summary	7
3. Understanding the Components of the Solution	7
3.1 Foundations of Cybersecurity.....	7
3.1.1 Encryption vs. Authentication.....	7
3.1.2 Public Key Infrastructure (PKI).....	9
3.2 Considerations for PKI.....	11
3.2.1 Infrastructure Complexity	11
3.2.2 Private Key Generation and Storage	11
4. Repackaging as a Scalable Security Solution for IoT	13
4.1 Reducing PKI Implementation Complexity	13
4.2 Certificates in Secure Chips	13
4.3 Designing Products for Security	14
4.3.1 Automated Secure Device Commissioning.....	15
4.3.2 Access Control	15
4.3.3 Access Denial	16
4.3.4 Executable Code Verification	17
4.4 Bringing It All Together	18
5. Conclusion.....	19

List of Figures

Title	Page Number
Figure 2 - Asymmetric Cryptography	8
Figure 2 - Creation and Verification Process for a Digital Certificate	9
Figure 2 - Digital Certificate Authentication.....	9
Figure 2 - Example Public Key Infrastructure Hierarchy	10
Figure 2 - Example Revocation in a PKI Ecosystem	10
Figure 2 - Public/Private Key Authentication Example	12
Figure 2 - Server-to-Server Authentication	15
Figure 2 - Device-to-Server/Cloud Authentication	16
Figure 2 - Device-to-Device Authentication	16
Figure 2 - Using Managed PKI to Create Separate Sub-CA Branch to Sign Code	17
Figure 2 - Code Signing with Managed PKI.....	17

1.Executive Summary

The industry has been chasing its tail for the past 5 to 7 years on the issue of Internet of Things (IoT) security, and it seems every week brings a new article about the need for device security or details about yet another security vulnerability exploit. The Internet is teeming with articles about issues and after-the-fact bandages, but very few of them get to the heart of the problem, which is how to secure network ecosystems that include interoperable autonomous devices.

IoT adoption continues to grow—but at the expense of good network and cybersecurity practices. Industrial and commercial IoT had previously been characterized by isolated networks that allowed devices within the network to communicate, but there was no connection to the outside world. In these use cases, it was possible to get away with weak security because it was more difficult to execute a wide-scale attack from the Internet. However, there is growing demand by utilities and builders to enable external communication and control of commercial devices to improve energy efficiency and provide better power grid management. This requires that commercial, industrial and even residential IoT devices be connected to the Internet so that they can be reached by utilities and state energy regulators. These would include lighting control systems, smart meters, solar inverters and home appliances. In fact, network connectivity is already starting to be mandated in some states (e.g., California Rule 21).

However, as critical electric power infrastructure is being network-connected, there needs to be an economical solution that adequately addresses security concerns as well as the logistics surrounding its implementation.

Companies that can provide strong security at scale will be able to use that as a key differentiator for their products, protect their brand and future-proof their products—which can have lifespans of 10 to 20 years or more—as calls for stricter requirements regarding device security loom on the horizon. Even as more wired control systems get connected, wirelessly connected devices are seeing exploding growth. Wireless devices are much easier to install and often reduce deployment time from several weeks to just a few days—or even hours. Easier installation reduces the amount of time installers need to spend on a job, thus reducing costs and increasing revenue by enabling them to do more jobs in the same amount of time. However, expanded wired and wireless connectivity accelerates the need for a more scalable security solution for these networked devices.

This paper covers the fundamentals of security architecture, best practices and new processes that can vastly simplify the implementation of strong, enterprise-grade security into small resource-constrained IoT devices. The goal is to enable deployment of security on the massive scale needed for IoT, while not sacrificing security robustness, and provide a workflow that can be implemented in hardware across a highly fragmented, embedded system.

This paper will not cover any hacks or exploits; those have been covered quite sufficiently to date. What is needed are more articles that cover the “how” of IoT security, not just further descriptions of new problems. In the new IoT reality, users need to know how to apply security that is strong, simple and massively scalable to tens of billions of hardware devices.

2.Scope of the Problem

The first question that is often asked is: “Why bother? A small IoT device possesses limited data and limited capabilities, so why is it worth securing?” Consider the situation from another perspective. When you are issued an access badge at your company, is the company securing you as a person? No, the company is securing access to its assets. This security measure is a way to control access based on the verification of your identity. The same concept applies to the IoT device: It should be verified before it is permitted access to the network. The device itself is not nearly as important as what the device potentially has access to. Just as unverified people should not be allowed to wander through secure buildings, unverified IoT devices should not be allowed access to your networks. With advances in technology and new logistical processes, it is now possible to provide manageable secure identities to IoT devices by default on a massive scale.

Resolving the issue of IoT security has been a complex problem to address. The reason the problem remains is that people have been trying to address it within the limited scope of their own market position and place in the value chain. The importance of IoT security spans multiple interrelated but very different market constituents. Before we can understand how to solve the overall problem, we need to understand each part of the IoT value chain and the respective concerns and issues surrounding cybersecurity implementation.

2.1 Device Manufacturers

General-purpose devices such as PCs, mobile phones and servers are what consumers are accustomed to thinking about when it comes to computer hardware. These devices have large processors, contain lots of memory and storage, can execute many types of software applications and can perform many different tasks.

Embedded systems, by contrast, are purpose-built to perform a specific task. Examples of embedded systems are lighting control systems, Wi-Fi-enabled thermostats and networked security cameras. Each of these systems has a microprocessor, memory and storage, but they are generally much smaller and optimized to perform specific tasks and no more.

Most of the IoT devices hitting the market are small embedded systems that have microcontroller-class processors with far more limited compute power and resources available to them. These devices include temperature sensors, light switches, cameras and so on.

For the most part, manufacturers of these devices use network connectivity in their products to improve functionality, features and ease-of-use/installation. Network connectivity allows manufacturers to build value with improved functionality, and wireless network connectivity allows further enhancement by simplifying deployment.

Cost reductions and advanced miniaturization have made it technically possible and economically feasible to network-connect very small devices and even put wireless radios in the smallest, simplest devices. However, device network connectivity (as well as Internet connectivity in itself) requires greater emphasis on security to control access to the networks and ecosystems those devices are connected to. Wireless-enabled IoT further compounds the need for scalable security because the ease of connectivity and deployment of such devices makes them not only abundant but easily accessible.

The problem is that most of these companies typically do not have or cannot afford a dedicated team of cybersecurity specialists. So far, companies have not been held accountable for producing IoT products with inadequate security, but that is changing with greater government scrutiny.

The final link in the device manufacturer chain consists of those involved with the installation of these devices. Whether they are professional installers on commercial construction sites or end users in their homes, IoT devices are deployed by the thousands every day. Device authentication to services needs to move beyond the username-and-password paradigm that assumes every machine has a human being behind it. In addition to being a weak form of security, usernames and passwords are not scalable. They work adequately for 10 devices, but they do not work for a company that requires stronger cybersecurity and will ship 10 million devices to a global sales and installation channel.

Thus, IoT device makers need a security solution that is inexpensive on a per-unit basis, uses minimal computational resources in the device and does not require a cybersecurity specialist to implement. Most IoT devices will use small microcontrollers that do not have a lot of compute power, and it does not make economic sense to put a large System on a Chip (SoC) in place merely to crunch cryptographic math for an operation that is only used to establish the authenticated secure session. How can you accomplish strong cryptographic security using a small, inexpensive microcontroller that doesn't require a cybersecurity expert to implement and doesn't resort to using weak aftermarket network security?

It turns out that there are solutions on the market that can provide strong, scalable device security for a broad range of devices with a price point of less than \$1 in moderate volumes.

2.2 Semiconductor Manufacturers

The use of secure element chips is becoming more prevalent in the market. These application-optimized cryptographic chips provide a pre-packaged solution for securely storing private keys and they also provide crypto-math acceleration, which is very useful when used in conjunction with small, low-power microcontrollers.

However, an issue arises when promoting these chips to the embedded systems companies that design and manufacture the devices. When the time comes to complete the sale, the companies face the issue of setting up Public Key Infrastructure (PKI). Specifically, the challenge is establishing the chain of authority for the certificates and the associated cryptographic keys that are provisioned into the secure elements. Today, it is up to the device manufacturer to find a PKI provider and to define the security requirements around the company's security domain and its policies. It requires the chip manufacturer's customer to be quite knowledgeable about cybersecurity and cryptography. On top of that, it is up to the customer to coordinate the deployment of PKI with the manufacturing flow of the secure elements. This all adds complexity and friction to the sales and design processes for a secure chip manufacturer.

Chip manufacturers need a security infrastructure that is well integrated into their production flow and that also abstracts from their customers most or all of the technical complexities behind establishing a certificate chain. Their customers need something that is as simple as adding a component to their bill of materials (BOM).

2.3 Cloud Service Providers

Due to the massive scale of device management for IoT product lines, most device vendors will use some form of IoT cloud-based management service (e.g., Microsoft Azure, Google Cloud, Amazon Web Services IoT). However, cloud service providers need a scalable provisioning process so that:

- Devices authenticate with the cloud provider's servers to prove that a given device originates from an authorized supplier and establish secure Transport Layer Security (TLS) communication sessions.
- Devices are automatically assigned to the proper company accounts. This must be able to work for thousands of manufacturers, each with many product lines, and must happen in a way that is seamless to the end user.

Cloud service providers have a similar problem as chip manufacturers because their customers are device manufacturers who need a simple, scalable way to strongly authenticate devices from many different manufacturers to their online services. These providers need a security solution that coordinates between the device makers, chip manufacturers and the cloud provider's online authentication processes.

2.4 Security Infrastructure Providers

Elliptic Curve Cryptography (ECC) is rapidly becoming the algorithm of choice for IoT because of the smaller key sizes compared with RSA. A 256-bit key in ECC is of roughly equivalent cryptographic strength to a 2048-bit RSA key. This makes ECC a more practical way to implement strong asymmetric authentication in small networked devices. Asymmetric authentication—arranged in a hierarchy known as PKI—is the framework on which managed secure infrastructure is constructed. As discussed in more detail later, authentication is the cornerstone of network security because it must be able to identify whom you are communicating with. If you cannot verify who you are sending data to, nothing else really matters.

PKI is the cornerstone of most enterprise cybersecurity plans, but for IoT it needs to be repackaged so that it fits seamlessly into the IoT hardware supply chain. In addition, the packaging needs to make PKI simple enough for implementation by device manufacturers who may have limited or no knowledge of cryptography, while not weakening its underlying secure authentication processes.

PKI has been in use for about 30 years in various forms, beginning with RSA in the 1980s and recently migrating to ECC. PKI has many advantages:

- It is highly scalable because credentials/certificates are digitally signed by an authority that has very rigidly controlled access. Any credential signed by an authority's private key can be validated using the authority's public key. As a result, millions of device credentials can be validated using a single key, which makes key management very simple.
- Even though a single public key is used to validate many credentials, each credential is unique. This allows for very granular management of devices because each one can be uniquely identified.
- PKI has been rigorously tested over the decades and is the basis of security used in most enterprise datacenters and servers that house sensitive data. It is not perfect, but it is robust enough for the vast majority of applications including IoT.

For all of its benefits, PKI's main weakness is the cost and complexity of its deployment. PKI was originally deployed in enterprise software environments where customers typically had security teams that understood network security and how to properly employ it. In addition, because deployments were in server and browser software applications, there has been substantially more flexibility in making changes or applying updates and patches, compared with deployments in embedded systems.

Finally, in most web browser applications, there was no need to authenticate the computer/device itself; users needed only to authenticate the server. As a result, most network sessions were only authenticated one way because the device didn't matter. It was the user that mattered. So, users authenticated themselves in the web application using their username and password to verify their identity.

For IoT, the identity of the device matters and the deployment model needs to adapt to that.

2.5 Scope of the Problem—Summary

The challenge is that for IoT, typically no active user is behind the device, unlike with PCs and mobile phones. The device logs in on its own and sends data on its own. For all intents and purposes, today an IoT device is a user on the network and now the task of authenticating the device itself becomes a concern.

The situation demands a practical and economical way to deliver private keys and certificates that belong to hundreds of PKI domains and thousands of manufacturers making billions of devices. This is something that PKI can do in theory but is not something that has been done scalably in practice. Therefore, we need to consider designing in a verifiably authentic digital identity as a basic capability in networked devices.

3. Understanding the Components of the Solution

Properly addressing the IoT security issue requires addressing the needs of each of the constituents in the supply chain. Otherwise, the solution will likely not be adopted or fit in the deployment flow. If a solution addresses only a specific constituent of the value chain, it is easier to deploy because it can be done independently—but is also likely to not be as effective because the other parts of the value chain are not cooperating in the solution. To properly address the IoT security problem, the different parts of the value chain must collaborate and coordinate with a solution that is continuous throughout the entire design and production flow.

3.1 Foundations of Cybersecurity

For many, particularly those not involved directly with cybersecurity, the popular term that comes to mind when the word “security” is mentioned is encryption. Although encryption is certainly necessary, it is by itself *insufficient* to provide a proper level of meaningful security in network communications. If encryption does not provide sufficient security, what does?

3.1.1 Encryption vs. Authentication

First, we need to understand that there are several components comprising proper security protocol and each of those components plays a key function in the process.

Encryption is an important part of any security solution, but it has a specific purpose: to prevent eavesdropping on transmissions. If you send data from one place to another, you do not want unauthorized people to intercept and read that data in transit. Modern encryption schemes, known as ciphers, are quite strong and effective at preventing interception of data transmissions. However, encryption by itself lacks a critical capability: It is unable to verify *whom* you are communicating with.

If you cannot verify whom you are sending data to, encryption by itself becomes less meaningful. Even though people cannot eavesdrop on your transmissions, you may be communicating with one of the entities you were trying to avoid in the first place because their identity cannot be established with reasonable certainty.

The ability to exchange verifiable credentials and validate them is known as *authentication*. The two most common forms of cryptographic authentication are symmetric and asymmetric authentication.

Symmetric authentication is fairly simple: The sender and receiver share the same key, and you can do a relatively simple random number challenge (or *nonce*) to determine whether the entity you are communicating with has the same key without actually transmitting any keys between you. Thus, for every entity you want to authenticate with, you have a unique symmetric key.

The issue with symmetric authentication is that scalability becomes an issue when the size of the ecosystem becomes very large. If an ecosystem contains millions of members, you will need something that manages millions of keys, and soon key management becomes a complex and costly system to maintain. In addition, key security must be maintained by both the sending and receiving parties. This is further complicated by the problem of securely provisioning keys on both sides— especially if devices are introduced to the ecosystem at different times and you need to authenticate new devices with older ones that are already deployed. If either side leaks its private key, the identity is compromised. If you add the complexity of a multi-vendor open ecosystem, the issue gets even more difficult.

Asymmetric authentication, as the name implies, involves a system in which the sender and receiver have different keys. The two keys are mathematically related to each other and are generated in pairs. The private key is protected and is used to “sign” digital data. The public key is often included in what is called a digital certificate and is used to verify the signature on its certificate. In the case of asymmetric authentication, only the signing (private) key needs to be protected. This accomplishes authentication in which the public key verifies that the device holding the paired private key sent the message. The strength of a public key cryptography system depends on the impractical amount of computation required to derive the private key from its paired public key. This means that effective security only requires keeping the private key private; the public key can be openly distributed without compromising security.



Figure 1 - Asymmetric Cryptography

This scheme also has the critical benefit that devices can easily be added at different times by different suppliers. Because the public keys can be openly distributed, enabling new devices to authenticate with devices that have already been deployed is easy.

The operation of asymmetric authentication involves three main operations: hash, sign and verify. A cryptographic hash is a one-way algorithm that produces a 256-bit number (in the case of SHA-256) that is consistent but not reversible. If you put in the same data, it always comes out with the same number. However, it is computationally infeasible to work backwards to determine the original data from the hash alone. So, to prove that data has not been altered, you provide someone with both the data and its hash. That person re-runs the hash with the data, and the results should match if the data has not been altered.

The sign operation is performed with the private key. Assuming that the private key is protected and has restricted access, data/hashes signed by a private key can originate only from an authorized source with access to that private key. So, a successful verification by the public key associated with that private key ensures the origin of that data. The combination of hashing and signing data forms the certificate. The creation and verification process for a digital certificate are shown below.




Figure 2 - Creation and Verification Process for a Digital Certificate




Figure 3 - Digital Certificate Authentication

A common scenario is to use asymmetric private keys as signing or “certificate authorities” (CAs) and arrange them into a hierarchy that generates, signs and organizes digital certificates. This security hierarchical structure is referred to as PKI. The CA certifies ownership of the key pairs and provides proof that a certain public key is authentic, belongs to the entity claimed and has not been tampered with.

3.1.2 Public Key Infrastructure (PKI)

As mentioned in the previous section, asymmetric authentication is often arranged in a hierarchy of CAs that sign and issue digital certificates/credentials. An example of a PKI hierarchy is shown below:

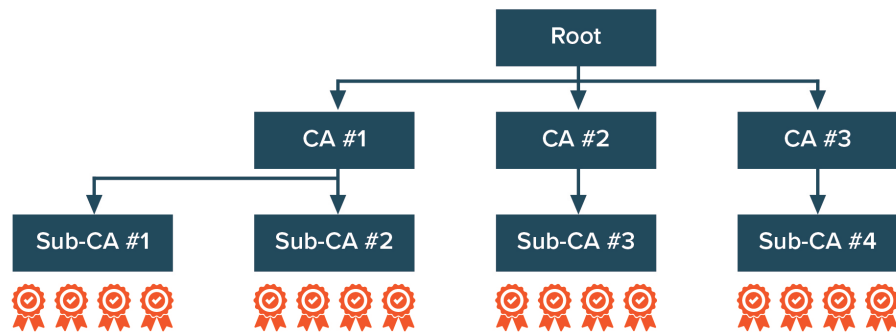


Figure 4 - Example Public Key Infrastructure Hierarchy

Each CA grants authority to sub-CAs, which then sign digital certificates for devices. The digital certificates at the bottom are carried by end devices and are authorized by the sub-CA above them that generated and signed them. These are generally called *device certificates*.

The sub-CAs that generate the device certificates possess their own certificate, authorized by the digital signature of CA above them, and so on. PKI eventually terminates at the root, which is the foundation on which this particular PKI ecosystem domain is constructed.

The reason PKI ecosystems are arranged in hierarchies is that it allows for selective levels of revocation or access denial if it is later determined that a leak or compromise of a private key in the ecosystem has occurred, as shown below:

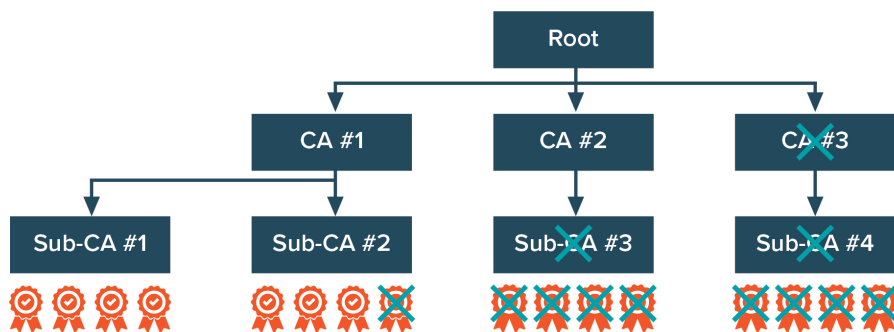


Figure 5 - Example Revocation in a PKI Ecosystem

As the figure shows, you can revoke the certificate of any element within PKI from the device up to a high-level CA, depending on the nature of the security compromise. However, it should be noted that if you revoke a certificate, you also revoke anything below that element in the hierarchy. It is possible to specify the revocation to be effective after a specific date noted in a PKI certificate, or be fully retroactive. It should also be noted that revocations cannot be undone— so they should not be initiated lightly.

This illustrates why PKI implementations are set up in tree-like hierarchies: This design allows for the ecosystem owner to perform selective damage control in the event of a compromise. It is also the reason why it is not a good practice to issue any device certificates off the root CA, which limits flexibility. If something goes wrong in that case, you may need to invalidate and revoke the entire PKI and all deployed devices in the field, which is generally not a good option. This is why device certificates are almost always issued by sub-CAs below the root.

Using this hierarchical structure enables the recipient of a digital device certificate to validate its pedigree or “chain of trust.” To do this, you use the public keys recursively up the chain to verify each signing authority above the device certificate to validate the device’s origin and the validity of its certificate. This is known as a chain validation. To avoid excessive computational requirements and time, it is recommended to not make PKI hierarchies too deep, because each level of validation requires additional certificate storage as well as computation.

The final and potentially largest benefit is that with a single sub-CA generating potentially millions of device certificates, you can authenticate millions of devices with one sub-CA public key, making key management far more scalable and manageable than other options.

If properly designed, PKI can be a compelling solution for IoT security. PKI is highly scalable and allows management of credentials and access control, so why hasn’t it been adopted more widely for IoT?

3.2 Considerations for PKI

Although PKI certainly has many attractive aspects, it still has some considerations that need to be properly addressed for it to be an acceptable security solution for IoT devices. Below are a few complications with PKI as it relates to IoT applications in hardware and large-scale device deployments.

3.2.1 Infrastructure Complexity

One of the main drawbacks of PKI and asymmetric cryptography is the solution’s complexity and the relatively intensive mathematical computations needed to perform operations.

When PKI was first developed, its target applications were servers, datacenters and web hosting. In these environments, security was implemented in software and companies that used PKI typically had a dedicated team of experts.

Setting up a PKI requires that the customer understands how to define the architecture needed based on the use cases. In addition, the customer needs to define the format of its certificates and the security policy around the management and protection of the PKI ecosystem. Next, the customer needs to implement the PKI hierarchy (or hire contractors to do it) and make sure that all the proper security protocols and cipher suite modules are used. Finally, because the private signing keys are the authorities that enable access to the ecosystem, they need to be securely hosted and periodically audited for policy compliance. Even with a sophisticated customer audience, the complexity of PKI makes implementation from scratch non-trivial and costly in both time and money.

3.2.2 Private Key Generation and Storage

Although asymmetric cryptography requires the protection of only the signing private key, not the verifying public key, any entity that needs to prove (authenticate) its identity needs to hold a private key.

A digital certificate is similar to a passport in that they both include two critical authentication components:

- Something that proves the credential originated from an authorized source and was not altered
- Something that proves the credential actually belongs to the bearer

On a passport, the embedded holograms, graphics and other physical security features make it very difficult to fake and alter the document. Examining the physical security features helps prove that the document is authentic and originated from the Government. The picture on the passport proves that the bearer of the passport is the owner.

For digital certificates, the asymmetric cryptographic signature on the certificate allows the recipient to use available public keys to verify the origin of the certificate and also prove that the data within it has not been modified. It proves authenticity and is akin to a passport's holographic security features.

The digital certificate contains a public key in addition to its certificate data. The public key is mathematically related to a specific private key. So, when a digital certificate is presented, the public/private key pair relationship proves that the device that presented the certificate actually possesses the unique private key associated with that certificate. By doing what's called a random number challenge, the recipient of the certificate gives the sender a number to hash and sign with its private key and send it back. If that private key corresponds to the public key in its certificate, then the recipient can decrypt the response using the public key and compare it to its own hashed number. If the two hashes match, the authentication passes. If not, the authentication fails.




Figure 6 - Public/Private Key Authentication Example

Through this example, you can see why protecting any private key is vitally important— because it is used to prove rightful ownership of a digital certificate. Every member of an ecosystem needs a certificate to gain access, and therefore every member of an ecosystem needs a private key that is well protected. Consider the signature stamp of the company CEO: If that gets out, anyone can start signing things on behalf of the CEO and you could not tell which signatures are real or fake because you would not be able to determine the rightful owner.

Furthermore, the generation/creation of the public/private key pairs needs to be tightly controlled to prevent unauthorized access or leakage of private keys. In the best case, private keys are generated within a secure storage chip and are never exposed outside that device.

In the days when only servers hosted in protected datacenters needed to authenticate themselves, key generation and storage was typically performed in software because those systems were generally not physically accessible and (if properly secured) did not permit access to protected key storage.

In addition, the scale was quite different: The number of websites at the time was relatively small and easier to manage. Even today, there are approximately 1.8 billion websites in total around the world, although only about 250 million of them are active. That number may seem large, but the scale of IoT and small embedded systems is measured in billions of microcontrollers sold per year. The IoT future will involve generating and properly storing private keys on a scale at least an order of magnitude larger than web server certificates.

The current technology of PKI can theoretically support that future, but the logistics of making it happen at that scale is a whole new problem. IoT devices are meant to be deployed in the field and in places that are physically accessible to the public. Yet, the devices must hold a private key so that they can authenticate themselves to gateways, servers and cloud service software. For PKI to work for IoT, it needs a method to provide strong protection for private keys stored in small devices that have limited compute power— and be able to do it economically on a massive scale.

4.Repackaging as a Scalable Security Solution for IoT

The industry needs a solution that can provide the security robustness of PKI but that abstracts away the technical and logistical complexities. Most previous attempts at this were done by companies trying to address the security issue from within their own market vertical, with limited collaboration across the separate parts of the supply chain. As a result, these solutions are only partially successful at best.

4.1 Reducing PKI Implementation Complexity

Although a typical PKI implementation has its own certificate policy, profile and other security specifications unique to the requirements of that ecosystem, the reality is that most IoT applications do not require this level of customization.

Enterprise PKI solutions tend to provide the “Cadillac” version of software and systems, which is highly customized and tuned to a specific client’s needs. When you are dealing with a multi-billion dollar bank or global corporation, that makes sense. However, for IoT, that approach is neither economically or technically feasible. Small devices are highly cost-sensitive and require any security solution to fit into their manufacturing flow— not vice versa.

As such, a pre-established PKI that encompasses a multi-vendor ecosystem is more appropriate. In this case, a standard is established that covers the needs of a broad range of adopters. A PKI implementation is created based upon this standard, and all adopting manufacturers simply need to use certificates issued from that PKI. This creates much greater efficiencies and economies of scale because the costs of creating the PKI implementation are spread out among many companies who share the benefits of the common ecosystem PKI.

4.2 Certificates in Secure Chips

Advances in semiconductor miniaturization have made it possible to create chips specialized in cryptographic key storage and mathematical operations. In addition to being physically very small to save

circuit board space, they include physical and electronic protection mechanisms to prevent unauthorized access to private keys. They are commonly referred to as secure elements, although the term is something of a misnomer because they often perform many functions and are more like a cryptographic co-processor.

Private keys stored in secure elements are literally inaccessible once locked. You cannot electronically access the memory slots, and the chips themselves include features such as extra metal layers that prevent chip de-capping and micro-probing of memory cells to extract keys.

The co-processing capability is critical for small IoT devices because the math behind cryptographic authentication can be computationally intensive. By offloading the cryptographic operations to the secure element, the use of a very simple main microcontroller will save cost and simplify the design. In addition to the ability to perform advanced crypto functions in a few tens of milliseconds (orders of magnitude faster than if performed in firmware), this technique also saves code space and power. The crypto functions are hard-coded in the secure element, which means those functions do not need to be included in the main firmware code, thereby reducing memory requirements. Because the hardware is optimized for performing cryptographic math, it completes these functions far faster. Less compute time equals fewer clock cycles, and fewer clock cycles equal less power used.

Once the authentication and crypto operations are complete, the secure element goes to sleep and draws current only on the scale of nanoamperes.

Drawing upon the pre-established PKI ecosystems referenced in the previous section, it is now possible to pre-provision keys and certificates into these secure elements. By doing so, you have now reduced the implementation of PKI and digital certificates in a small device to a line item in a BOM. The cryptographic math is baked in, as are securely stored keys and the digital certificate. You add the secure element with its digital certificate on the I²C bus next to your host microcontroller, add a small library/SDK to your firmware, and you are done.

Once the implementation is complete, you have now placed enterprise-grade security into an IoT device. Each device has a unique, cryptographically verifiable identity. You can verify that the certificate is authentic because the digital signature prevents the data from being altered. The securely stored private key proves that the certificate corresponds to the device that sent it. On top of that, this is a method that can be rapidly scaled to billions of units, if needed.

4.3 Designing Products for Security

The pre-packaged certificate-in-a-chip can now be used by the device manufacturer or system integrator. The key to implementing a trusted ecosystem is to treat IoT devices like users on a network. An IoT device logs in by itself and it sends data by itself. Certificates-in-a-chip allow for the ecosystem owner to embed verifiable unique identities in devices so that they can be managed almost like human users. The difference is that the device identity is authenticated with the ecosystem certificate as opposed to a username and password.

With the turn-key certificate-in-a-chip solution, different product lines can very easily adopt a common authentication method that is based on strong cryptography and security methodologies, yet simple enough to be implemented by engineering teams that are not specialized in cybersecurity. Implementation is a matter of adding one additional component on the device/system BOM and integrating a small software module into the main firmware to call the secure element when it is needed.

Now that there is a cryptographic identity embedded in every device, what can you do with it?

4.3.1 Automated Secure Device Commissioning

One of the trickiest aspects of IoT is adding a new device to a network or ecosystem. How do you know that the device is authorized to join? For IoT, a device's physical location typically matters because it is controlling something or taking data in a specific area. How do you assign a device to a location on a large scale?

With embedded certificates securely stored in IoT devices, you can use that unique identifier with mobile device software to rapidly deploy devices and assign physical locations to them and do it securely. A technician's mobile devices, IoT devices, gateways and servers will all have ecosystem certificates, and anything that wants to connect to them must authenticate using a certificate and private key.

With pre-provisioned keys and credentials embedded in secure chips that are manufactured into the products (as opposed to added in after manufacture), a cryptographically verifiable, securely stored, unique device identity is an integral part of each device. As such, it is able to authenticate locally in the case of isolated networks, or authenticate over a WAN to public or private cloud services and other network resources right out of the box.

Although this all sounds complicated, it can be completely transparent to the installer and end user, if done properly. The mutual authentications occur as part of standard protocols such as TLS. Assuming that all certificates are valid, the secure connection, device commissioning and other data transfer happens without any additional user intervention. It just happens, and the installer can move onto the next installation.

4.3.2 Access Control

With a verifiable identity and certificate data, you can now effectively control who has access to what on the network. Because each device can be individually identified by its certificate, you can put devices in groups and determine which devices can access certain parts of the network and which cannot.

Certificates can be used for server-to-server authentication:




Figure 7 - Server-to-Server Authentication

They can also be used for device-to-server/cloud authentication:




Figure 8 - Device-to-Server/Cloud Authentication

And they can be used for device-to-device authentication:




Figure 9 - Device-to-Device Authentication

Most access breaches occur not as a result of breaking the encryption but as a result of private key leakage or identity spoofing. An example of identity spoofing is when something like a serial number or MAC address is used as the device identity. If the intruder knows the proper format of that number and presents one that fits the profile, there are very limited (and often cumbersome) ways to verify whether the number is legitimate. Spoof attacks do not actually break any encryption or security. Spoofing is essentially identity theft and allows the unauthorized entity to pretend that it is someone/something that has legitimate access to the network. By spoofing the identity, it gains the same access as that identity.

Providing access control for devices is similar to providing access privileges to human users, but certificates are used to provide their identity. If the certificate is issued from a well-managed PKI and private keys are securely stored in the device, a digital certificate identity is very difficult to spoof because the signature prevents alteration and the private key proves ownership. Even in the case of the compromise of a private key, the use of a managed PKI allows for the revocation of a certificate. This prevents use of the compromised certificate even if all of the cryptographic math works out and helps provide layered security.

4.3.3 Access Denial

Although a digital certificate is difficult to spoof, it is not impossible. It is possible to potentially steal chips or compromise PKIs that are not well managed or implemented. In this case, it is possible to revoke a certificate, as described in Section 0. Revocation allows the ecosystem owner to disallow access by

someone/something bearing a certificate even if the cryptographic authentication succeeds. This can happen when a private key is leaked along with its associated certificate. Because the private key is what proves legitimate ownership, legitimate ownership can no longer be proven once they private key is stolen.

As shown in Section 0, you can revoke a single device certificate or an entire branch or chain of a PKI implementation. These revocations are logged in what is called a Certificate Revocation List (CRL). These are posted on servers that are called Online Certificate Status Protocol (OCSP) responders. Some techniques, such as OCSP stapling, help improve response time and efficiency, but details about those are beyond the scope of this white paper. Suffice it to say that when something wants to authenticate and sends its certificate, you can look it up to see whether that certificate has been revoked.

4.3.4 Executable Code Verification

Another benefit of having a managed PKI is that you can create a separate sub-CA branch that can be used to sign code.



Figure 10 - Using Managed PKI to Create Separate Sub-CA Branch to Sign Code

This is often referred to as a Code Verification Certificate (CVC) sub-CA. Authorized developers use the CVC sub-CA to sign code so that it carries a signature from the ecosystem PKI that can be verified. For example:




Figure 11 - Code Signing with Managed PKI

With code signing, you can now secure the trust in your ecosystem even further by programming your devices to only permit execution of code that has a verified signature. When combined with a Trusted Execution Environment (TEE), available in many of today's microcontrollers, you use PKI, digital certificates and secure semiconductors to provide a trusted ecosystem not just within a single device but across a global network of devices.

4.4 Bringing It All Together

For all entities in communication with one another within a trusted ecosystem, you need to:

- Verify the identity of the device
- Verify the identity of the server
- Verify the signature of any executed code

This means that in the ideal case, any communications over the network must be mutually authenticated using each endpoint's certificate.

The best approach is to make it so that implementing good ecosystem security meets the following requirements:

- **It must be economical in terms of cost and time.** The large benefit of ecosystem security must be bolstered by requiring only an incremental increase in BOM and development time cost. It must be inexpensive and simple enough so that there is no reason not to do it.
- **The solution must not require in-depth cybersecurity expertise to implement.** Security engineers are not cheap, and not every company can afford to have its own. The solution must be simple enough for a non-security person to implement, but not compromise the security itself. Storing keys and credentials in secure silicon provides a pre-packaged solution that offers both secure key storage and strong cryptography.
- **It must fit within the existing manufacturing flow for hardware devices.** If a solution requires a substantial change to the established hardware design and manufacturing flow, the probability that it gets adopted decreases dramatically.

To address these points, the security infrastructure must be part of the design and manufacturing flow—not bolted on after the fact. If security is implemented from the outset, it can be fed into the manufacturing flow for the chips and subsequently the IoT devices themselves. From there, once it is known that new devices are carrying certificates, you can start to apply the methods outlined in Section 0.

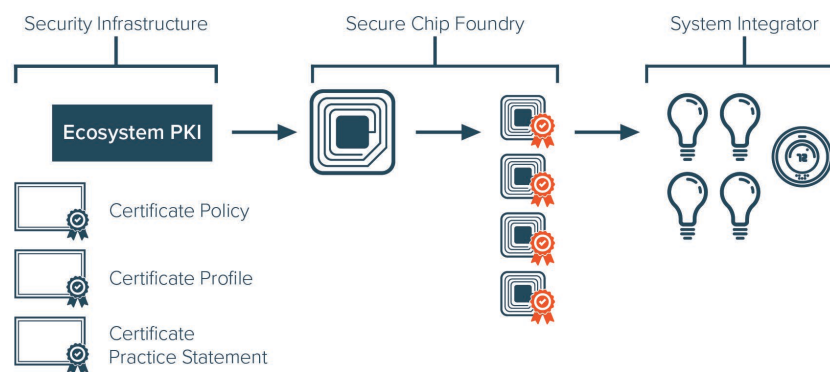


Figure 12 – Example Security Infrastructure in Design and Manufacturing Workflows

5. Conclusion

It is possible to implement highly scalable device security for IoT without compromising the strength of that security. If security is worked into the front-end of the design process, it makes it far easier to implement later on. With a managed PKI implementation backed by a trusted CA providing the rigor and process behind issuing and revoking authentic digital certificates, this secure back-end can feed directly into the manufacturing flow for semiconductor components used in everything from larger devices (e.g., gateways) to small IoT devices (e.g., lighting, sensors). With a cryptographically verifiable identity securely embedded in each device, you can maintain the integrity of the ecosystem and enable secure software updates by providing control over network access and code execution throughout the lifetime of your devices.

The value of IoT security lies in the protection of the greater ecosystem, not in individual devices. However, it is the sum of all the individual devices that collectively contribute to the security of that ecosystem which allows the value to be realized.

Internet Scale Blockchain Architecture

Akamai Technologies – Akamai Labs

A Technical Paper prepared for SCTE•ISBE by

Akamai Labs
Michael Fay et al.
Akamai Technologies, LLC
150 Broadway Cambridge MA
mifay@akamai.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. An Overview of Blockchain	3
2. Challenges and Limitations of Blockchain	4
3. A Blockchain Architecture for Scale & Speed	5
3.1. Core Tier	6
3.2. Consensus	6
3.3. Edge Tier.....	7
4. Blockchain as a Service	7
5. Benefits	8
Conclusion.....	8
Abbreviations	8
Bibliography & References & Copyrights.....	9

List of Figures

Title	Page Number
Figure 1 – Examples of Decentralized Permissions and Permissioned	5
Figure 2 – Tiered Architecture.....	6
Figure 3 – Transaction Lifecycle	7

Introduction

Blockchain has quickly become a disruptive technology enabling whole new business models and ecosystems including the rise of cryptocurrencies such as Bitcoin and Ethereum. It is changing how people think about recording interactions, transparency and auditability.

Applications can extend beyond just providing an auditable ledger for tracking payments, financial services transactions, digital assets and currency balances. Blockchain can be leveraged for proof of existence and proof of integrity for data and/or state and establish provenance via an auditable trail of interactions. This could be applied to many industries ranging from insurance, advertising, health care, shipping and logistics, commerce and defense.

However, many popular blockchain implementations have real challenges and limitations around scalability, performance, reliability and security. A new blockchain architecture is required to address enterprise and institutional use cases that need scale and speed.

Content

1. An Overview of Blockchain

A simple definition of blockchain is a digital record of interactions added over time and protected from alteration. A more detailed definition would be a permanent, append-only distributed ledger that addresses data provenance and provides transparency and auditability by embodying the canonical history of transactions leading up to a given exchange of value. Essentially, it's a database that enables sharing authority amongst participating entities and only allows information to be written once, preventing deletions and modifications.

So what's so special and useful about a blockchain? Since the information in a blockchain is not centralized and cannot be altered, it provides a great mechanism for use cases that require a tamperproof record. Examples of popular applications for blockchains include providing an auditable ledger for tracking payments, financial services transactions, digital assets and currency balances. In these cases, a blockchain can represent and manage real value and/or digital assets and thus a transaction includes both the clearing and settlement phases of a financial transaction. Another popular blockchain application is the use of Smart Contracts to provide asset custodianship services and scripted policy execution. Other potential applications include using blockchain to provide auditable trails of transactions including, everything from parts to repairs to medical records to warranties to insurance claims.

Blockchain is built upon several key principles, which include decentralized authority, transparency, and immutability. The first principle is that there is no central authority to approve transactions or set rules. Authority is distributed through consensus across multiple participants, usually represented by various networked computers. Transparency means that the records in the blockchain are self-verifiable, containing all information required for auditing by any participant. Lastly, immutability refers to the inability to alter or forge data once it has been committed to the blockchain.

Many of these applications are served by a centralized database, so what additional advantages does a blockchain provide? With centralized databases, clients must depend on the trustworthiness and reliability of the database operator. Database records are not inherently transparent, immutable, tamper-proof or self-verifiable. In addition, the cryptographic foundations of blockchain provide superior reliability and security in the face of failures or adversarial conditions.

In order to achieve these advantages, blockchains employ cryptographic techniques as a foundation of their design. At regular intervals, a decentralized group of nodes (networked computers) add a new block to the ledger. Each block contains a sequence of transactions organized into a tree of cryptographic hashes called a Merkle Tree. The root hash of the tree is recorded in the block header, and the hash of the block header uniquely identifies the block. Each block is cryptographically linked to its predecessor by including the previous block hash in its header, forming a “chain”. Clients of the blockchain submit transactions to the network that must be digitally signed by a valid private key. This digital signature makes the transaction computationally intractable to forge or alter. Furthermore, the use of cryptographic hash functions in the block structure make it intractable to modify any block. Therefore, both transactions and blocks are tamper-proof. This property makes blockchains applicable to many use cases that require trustworthy exchanges between multiple, independent parties (e.g., auditable records that are shared between insurance and medical providers and their patients). The first block in the chain is called the “genesis block”. Each block has a link to a single previous block. Following these links and verifying each of the transactions within a block allows any eligible party that has a copy to independently verify the entire blockchain.

2. Challenges and Limitations of Blockchain

Blockchain systems are typically limited by some of these characteristics: **scalability**, **performance**, **reliability** and **security**. These limitations are a barrier to leveraging blockchain for use cases that require a high-performance solution to quickly, securely, and efficiently process transactions with almost limitless scale.

Today, many of the limitations of blockchain implementations are around **scalability** and **performance**. These two characteristics are closely related and highly dependent on the architecture and implementation of the system. Some of the impacting factors include the number of nodes, the number of users, the number of transactions, and the number of connections or network traffic.

In terms of **scalability**, many implementations have limitations to volume and rate of transactions processed. A decentralized permissionless network (such as Bitcoins) has challenges with **scalability** (or transaction volume). This type of implementation usually encompasses a large geographic area, potentially resulting in unpredictable latencies and unreliable timing assumptions and therefore is difficult to scale horizontally (adding more machines does not increase performance, but rather may deteriorate performance). Even current implementations of permissioned networks haven’t achieved high transaction volumes in terms of scalability.

While permissioned networks, without the burden of resource intensive computation, would seem to scale, they have not yet achieved the scale required. In addition, they have the disadvantage of limited geographic presence as compared to a distributed platform.

Performance limitations are primarily related to the latency in confirmation time or the time required to commit a transaction. To support massive numbers of users and transactions, considerations must be made for the nodes and the connections. Adding more nodes with more connections and traffic between nodes can negatively impact scale and performance (or confirmation times) because propagation times for both transactions and blocks will increase. This is a critical factor in why large decentralized permissionless implementations such as Bitcoin have slow block creation and transmission

Reliability is challenged by availability as it relates to distribution. Adding more nodes can increase reliability. While systems like Bitcoin are highly reliable, they suffer scale and performance issues as

previously mentioned. Controlling the number of nodes and connections by centralizing them in few data centers will reduce the reliability of the system because the nodes are less distributed.

Lastly, **security** and trust has been limited by unproven trust models and key management at scale. Decentralized permissionless approaches are protected against a single point of failure but have inherent risk due to lack of security and governance in the participating machines. Conversely, a permissioned approach has secure access and governance, but usually introduce a single/limited points of failure in terms of an attack surface.

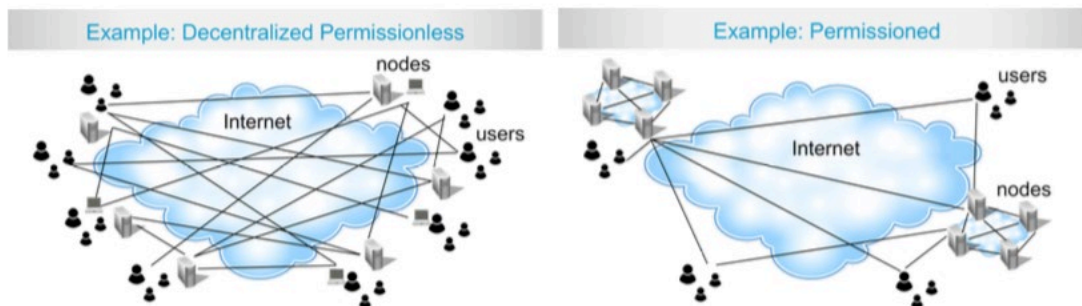


Figure 1 – Examples of Decentralized Permissions and Permissioned

3. A Blockchain Architecture for Scale & Speed

Akamai has built a new unique high-performance blockchain architecture that will enable institutions and enterprises to quickly, securely and efficiently process transactions with almost limitless scale. Akamai has focused on significantly improving the **scalability** and **performance** weaknesses of existing blockchain implementations, while maintaining and improving the **reliability** and **security** characteristics. Akamai's blockchain system is capable of processing 10M+ onchain transactions per second, with each transaction committed and confirmed in the system in under 2 seconds*.

So what is unique about Akamai's approach and the potential benefits? The Akamai blockchain architecture combines innovations in blockchain technology with Akamai's globally distributed platform. It provides the benefits of decentralization (geographic and network diversity) with a permissioned approach, offering improved scalability, reliability, performance and security. Akamai's globally distributed platform enables fast and secure access from anywhere around the world to provide a solution for enterprise customers that is superior to permissionless networks (like Bitcoin) and current permissioned network implementations. In addition, Akamai's globally distributed platform provides inherent critical security features to protect the blockchain network against a variety of attacks, including DDoS.




Figure 2 – Tiered Architecture

Akamai has architected a tiered approach for integrating blockchain technology with its globally distributed platform. Incoming transactions are accelerated and secured by the globally distributed Akamai platform (edge tier), then handed off to a highly secure set of network nodes executing the blockchain transactions (the core blockchain tier). Combining an Internet-scale core blockchain tier with advanced content acceleration and security capabilities in the edge tier, makes the architecture unique in providing an end to end flexible, performant and secure blockchain platform.

** Adaptable, represents end-to-end latency in a user generated payment processing transaction.*

3.1. Core Tier

The core tier provides unparalleled scale and speed for transaction processing, at the same time, adhering to all of the blockchain principles such as transparency, immutability, reliability, and self verifiability. To achieve scale, Akamai has applied years of experience in Internet-scale distributed computing into each node that processes transactions in the core. Other approaches to scale blockchains employ off-chain (or layer 2 approaches) to achieve scale. While these approaches have their merits, a system that processes and commits all transactions on-chain adhere better to the core blockchain principles. The layer 2 approaches could in fact be leveraged to scale the Akamai architecture even further. An innovative consensus algorithm powers the high speed transactions and block processing, with all nodes participating actively in finalizing transactions. In addition, the core tier provides high reliability, leveraging a node deployment that spans multiple, disparate data centers and geographies, combined with resilient network connectivity to handle disruptions.

3.2. Consensus

Akamai's unique consensus protocol is far more efficient in terms of scale, performance and cost than both current blockchain and traditional consensus mechanisms. The protocol ensures that node selection for block generation is both unpredictable and non-influenceable while remaining self- verifiable. In addition, current block propagation and finalization mechanisms have inherent limitations in scalability and resilience. To achieve block and transaction finalization at high speed, the Akamai protocol features configurable quorum requirements, with automated resolution in response to network partitions and attacks to overcome these limitations. In short, Akamai's innovative consensus protocol lays the foundation for a robust blockchain architecture suitable for Internet-scale adoption.

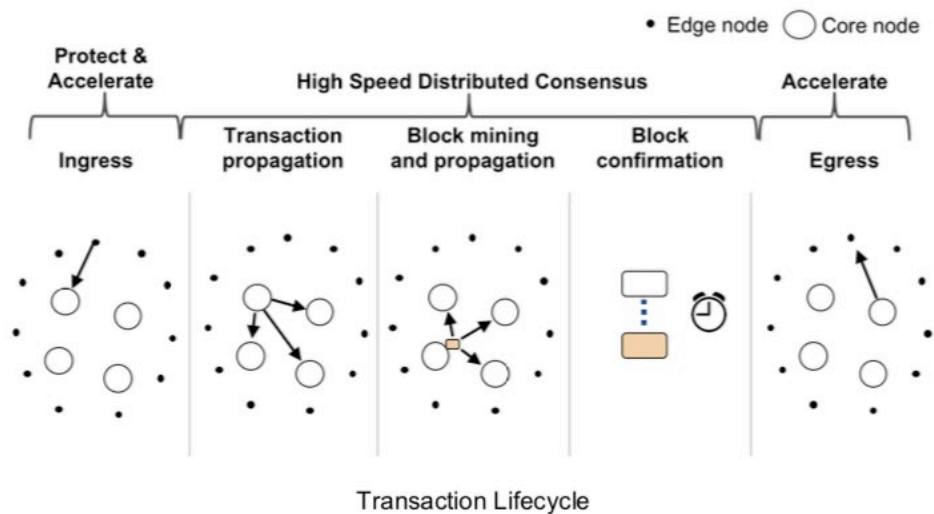


Figure 3 – Transaction Lifecycle

3.3. Edge Tier

Akamai’s globally distributed intelligent platform forms an integral part of the the blockchain architecture serving as its edge tier. This tier provides sophisticated capabilities by ensuring content is accelerated across the internet, solving for issues such as congestion, unreliable connectivity and sub-optimal routing. The blockchain architecture leverages these capabilities to ensure that transactions are ingested close to their origination and reliably accelerated to the core tier for further processing.

In addition, advanced security capabilities such as a cloud based distributed firewall provide attack resilience from a myriad of attacks on the Internet, including DDoS. Such advanced capabilities add a critical layer of protection for the core blockchain tier and ensure its security. Moreover, end to end content security and strong cryptography for the core blockchain tier is ensured by leveraging Distributed Key Management Infrastructure (KMI) features in the edge tier.

Apart from these market-leading capabilities leveraged by the most prominent brands on the Internet, the edge tier has innate flexibility to adapt the blockchain platform for varied use cases, by providing configurable workflows and data processing capabilities as part of the transactional flow. Further, the edge tier provides flexibility for applications with robust APIs and the ability to host application logic, close to transaction origination.

The capabilities of the edge tier are renowned in the industry as part of Akamai’s current services, serving a large portion of global Internet traffic. The Akamai blockchain architecture builds upon years of battle-tested capabilities to create a robust end-to-end blockchain platform for institutional and enterprise consumption.

4. Blockchain as a Service

The Akamai blockchain architecture anticipates supporting multi-tenancy, allowing it to safely host multiple customers on the platform. Akamai’s globally distributed platform is multi-tenant and currently supports thousands of mission-critical customer properties. In addition, given Akamai’s global network deployment, instantiating unique blockchain networks for different needs, addressing privacy and locality requirements is now achievable.

The combination of these capabilities allows Akamai to offer blockchain as a service, thereby unlocking tremendous potential for hyper-scale blockchain applications. Building, deploying, maintaining and securing a blockchain platform can be risky, time consuming and costly. The complexity and challenges of DIY blockchain deployments may be cost prohibitive and difficult to scale. As a service company, the value of Akamai's globally distributed platform offers any business or organization an on-demand service for security, delivery, acceleration and now blockchain.

5. Benefits

Akamai's blockchain architecture provides substantially improved benefits over other implementations. First, in terms of **scalability**, Akamai has chosen a highly concurrent scalable distributed node architecture. For **reliability**, Akamai's distributed platform provides superior reliability with diversity in geographies/networks with nodes that have smaller standard deviations in terms of compute/connectivity. To address **performance**, Akamai has innovated on a low latency distributed consensus, which lowers computation overhead and reduces confirmation time. Lastly, **security** leverages Akamai platform and expertise to improved risk mitigation. . In addition, Akamai reduces security and operational risks with its robustness via its round the clock NOCC, heterogeneity of network architecture, secure servers, and global reach. Finally, Akamai benefits from a virtualized management layer that manages deployment of and, communication between nodes and provides API that hides the complexity of managing network of blockchain nodes.

Conclusion

Akamai has built an innovative blockchain design, from ground up, that leverages years of distributed computing principles in each node of its network, with nodes deployed in heterogeneous networks to provide fault tolerance. Moreover, the system also implements an innovative consensus algorithm to ensure that key blockchain principles aren't sacrificed for scale or speed. This system inherently leverages Akamai's market leading advanced security and performance capabilities to enhance its robustness. The results in terms of **Scalability** have shown that Akamai's blockchain system is capable of processing 10M+ on chain transactions per second. In terms of **performance** transactions are committed and confirmed by Akamai's blockchain system in < 2 seconds*. Akamai has demonstrated **reliability** by coupling it highly available blockchain network with Akamai's global edge platform. Lastly, in terms of **security**, Akamai's blockchain transactions are protected by strong cryptography, and our platform network is protected by our cloud security.

Abbreviations

CDN	Content Delivery Network
tps	Transactions Per Second

Bibliography & References & Copyrights

Akamai secures and delivers digital experiences for the world's largest companies. Akamai's intelligent edge platform surrounds everything, from the enterprise to the cloud, so customers and their businesses can be fast, smart, and secure. Top brands globally rely on Akamai to help them realize competitive advantage through agile solutions that extend the power of their multi-cloud architectures. Akamai keeps decisions, apps and experiences closer to users than anyone — and attacks and threats far away. Akamai's portfolio of edge security, web and mobile performance, enterprise access and video delivery solutions is supported by unmatched customer service, analytics and 24/7/365 monitoring. To learn why the world's top brands trust Akamai, visit www.akamai.com, blogs.akamai.com, or [@Akamai](https://twitter.com/Akamai) on Twitter. You can find our global contact information at www.akamai.com/locations. Published September 2018.

Is the Smart Assistant Mutually Inclusive with IoT?

A Technical Paper prepared for SCTE•ISBE by

Charles Cheevers

CTO CPE Solutions
ARRIS International plc
3871 Lakewood Drive Suwanee, Georgia 30024
+1 678 474 8507
Charles.Cheevers@arris.com

Jonathan Wu

VP Product Management, Retail
ARRIS International plc
2450 Walsh Ave, Santa Clara, CA 9505
(408) 235-5500
Jonathan.Wu@arris.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
1. Smart Assistants and the making of the smart home	3
1.1. What is inside a Smart Assistant.....	3
1.2. The undeniable trend of consumer pull for Smart Assistants	5
1.3. Where should Smart Assistants be located and in what physical form?	9
1.4. The STB as the Smart STB with a splash of BLE IoT	12
2. Which IoT protocol and which Assistant or can there be more than one?.....	14
Conclusion.....	19
Abbreviations	19
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1 – The basic anatomy of the Smart IoT Hub	4
Figure 2 – The 4 for 1 insertion device	4
Figure 3 – The growing trend in Smart Speakers	6
Figure 4 – Smart Speaker growth vs other Home Devices.....	7
Figure 5 – What people do with their Smart Assistant devices.....	7
Figure 6 – What people do with their Smart Assistants	8
Figure 7 – What IoT services people most use with Smart Assistants.....	8
Figure 8 – IOT – its not just about turning on the lights	9
Figure 9 – What drives the design of the Smart IoT Hub.....	10
Figure 10 – Where it makes sense to put Smart Assistant and IoT in the home.....	11
Figure 11 – Smart Rooms not Smart Home.....	12
Figure 12 – Why the STB is a key device to make Smart	12
Figure 13 – Its not just the phone app – it's the remote and TV too.....	13
Figure 14 – The 5 for 1 opportunity with the Smart Media Device.....	14
Figure 15 – The Simple Cable IOT Stack	15
Figure 16 – Multiple Smart Assistants are possible.....	16
Figure 17 – The flow of voice from Mics to ASR.....	17

Introduction

It seems that the smart assistant can do anything these days – including be the pivotal application to enable the service provider IoT solution. As service providers wrestle with how to deploy their smart home solutions, it's clear that the smart assistant is a key part of enabling the solution.

This paper takes the reader through some of the ideas circulating the industry and reviews:

- What is the correct inclusion of mic and speakers in gateway, access point, extender, set-top, and standalone devices?
- Inclusion of IoT radios in the same devices – what are the pros and cons?
- What is the role of the smart phone, TV, and smart assistants?
- Amazon, Google, and Apple smart assistants: friend, foe, or must-have partner?
- What is the role of AI, advanced speech recognition systems, and NLU's in the service provider arsenal?
- Can there be more than one voice assistant per home?
- Could there be multiple smart assistant solutions in a single service provider device?

This paper also reviews the role of IoTivity and the OCF as the potential basis for any service provider IoT hub and connection solution.

1. Smart Assistants and the making of the smart home

1.1. What is inside a Smart Assistant

The anatomy of the Smart Speaker has typically been the following

1. Wi-Fi subsystem that connects back to the home Access Point and onto the internet
2. Far field microphones – Amazon Echo debuting with 7 far field mics – but improvement in mic and DSP technology has now seen 4 far field mic's being sufficient to recognize human voice
3. Speaker subsystem – has the biggest influence on the size of the device – and typically ranges from 2W at the lowest level of voice only feedback to potential full high end soundbar instances of Smart Assistants at 50Watt and above
4. IoT and Low Power Radios – most of the Smart assistants focus on just adding BLE for pairing and authentication during onboarding of the device and to allow streaming of audio from and to other audio sources and sinks – most notably streaming music from smart phones.




Figure 1 – The basic anatomy of the Smart IoT Hub

The main components of Smart IoT

Wi-Fi

- Connection to network or extension of network

IoT Elements

- H/W – 802.15.4, BLE, Z-Wave – other ?
- S/W – Iotivity/ZigBee/BLE/Z-wave – other ?

Voice and Audio

- Typically 4 far field microphones
- Speaker quality – 2W small – 12W larger

Multiple Voice Trigger words

- Support Alexa
- Support Ok Google
- Support – Cortana/Watson – other ?
- Support Ok <Service Provider>



The 4 for 1 device concept – with leverage of GAFA directions

Figure 2 – The 4 for 1 insertion device

The other elements of the smart assistant that are key are

1. The “wake” or “bargh word” – this is the spoken word that invokes the smart assistant and opens the logical connection to issue voice commands to the device
 - a. Typically, this word is coded locally into the DSP – to ensure that the device – while always in listen mode is not sending audio from the home to the cloud/internet.
 - b. Sometimes more than one “wake” word can be programmed. This will be the topic of a later section in this paper and one of the opportunities for the Service Provider.
2. The DSP engine which has a couple of functions
 - a. Most importantly, to be able to pick out the “wake” word from background noise. There are several techniques for doing this that are outlined later in the paper.

- b. Digitizing the vocal discussion and sending this digitized waveform to the cloud
- 3. The Cloud Voice Processing Engine
 - a. While there can be local voice processing in the Smart Assistant – typically to be able to discern single word actions – all Smart Assistants typically rely on a cloud connection to the Automated Speech Recognition system. This ASR service processes the Voice commands or conversation through a number of levels
 - i. Discerns language and then uses the appropriate language context to try and figure out what the person is saying.
 - ii. Can discern who is talking to separate out speakers – using voice pattern as passwords or parental control or limiting access to certain skills and actions
 - iii. Applies context to the voice command or discussion – as it tries to figure out what is spoken – it can use previous dialogs or specific learned words or items in dialog to help understand what is being said.

For example if “Whats my Wi-Fi password?” is the spoken text , the ASR function may discern that this is potentially one of the following requests

“Whats my Wifes password”


“Whats my Wi-Fi password”
- 4. The cloud based Natural Language Understanding Engine
 - a. One the ASR has tried to parse the digitized voice input to text strings – the NLU system tries to understand what the strings actually mean – using context and other elements to feed the NLU engine.
 - i. The NLU process can be as simple as processing one sentence at a time or with the introduction of more sophisticated engines can now even engage in an interactive discussion to ensure that the assistant fully understands the request.
- 5. The Skills or Actions that are executed by the NLU when it understands the request
 - a. This has been one of the most exciting part of the emergence of Smart Assistants and the core of this papers title. The ability to create a skills ecosystem that can be added to the smart assistant. These skills often run on devices that are not directly connected to the smart assistant but part of the IP network that can be accessed from the Smart Assistant. This could even include issuing a command to an Assistant in your primary dwelling home and executing the command in a Summer home 1000’s of miles away to start video recording on a security camera for example. The ability for third party companies to include their products to be controlled by a smart assistant – is both hugely powerful but also makes the usability of IoT and Smart Devices work even better and breaks down the technical use inertia for the device.
 - b. Today Smart Assistants add skills for more and more solutions and devices. This is primarily done by implementing primarily cloud API’s that allow for example a smart device like a light to be controlled by a separate smart assistant. In this simple example a smart light bulb that is IP addressable creates a skill that is coded to support the cloud API’s to the Smart Assistant. These are typically invoked by a voice command sequence like “Ok Jarvis Tell <Product Name> to turn on light bulb”. The Product owner is given the command string and it then executes the Skill. Skills typically must be enabled and authenticated to the Smart Assistant in a preliminary step to allow the connection between the assistant and the out of band skill.

1.2. The undeniable trend of consumer pull for Smart Assistants

One in three homes now has a smart assistant device – Figure 3. This number does not include Smart Phone Smart Assistants like Apple IOS Siri or Google Android Assistant devices. The standalone smart speaker has become a popular addition to US and global household. This growth has been fueled by

strong pushes from companies like Amazon and Google making it a fundamental part of their core business (Shopping in the case of Amazon and Search in the case of Google) and their desire to understand the consumer and their behavior more through more data gathering devices. The success of the Smart Assistant is also down in part to the continued investment in the Voice processing that has improved substantially over the last number of years to ensure that the assistant gets it right at 9 in 10 voice commands. This level of reliability keeps the consumer engaged with the device and the persona of the assistant!

Smart speakers ownership among US BB HH



Clear consumer interest in Smart Speakers with this impressive growth !

Figure 3 – The growing trend in Smart Speakers

The smart assistant is also the fastest growing home device with a projected 32% compound annual growth rate from 2017-2020 (see Figure 4) fueled by significant investment in Voice Processing, Natural Language Parsing and Understanding and alignment of the voice device with other growing services like Home Automation and Security solutions. Therefore, there is an undeniable link between the 2 worlds of Smart Speaker and IoT and why strong consideration should be given to rolling out both services together and potentially leveraging a single Speaker and IoT device in one solution. There are certain home factors – where people are mostly to engage in voice commands with Smart Assistants – which drive the decision on their type, location and effectiveness. This will be discussed later.

32% CAGR – the highest from all smart devices

Product Category	2017 Value (US\$M)	2022 Value (US\$M*)	CAGR, 2017 – 2022*
Video Entertainment	\$133,091.48	\$201,063.36	9%
Home Monitoring/Security	\$4,271.30	\$12,136.50	23%
Smart Speaker	\$4,401.39	\$17,431.00	32%
Lighting	\$1,120.53	\$3,511.32	26%
Thermostat	\$1,774.35	\$3,875.91	17%
Others	\$17,532.54	\$38,963.93	17%
TOTAL	\$162,191.59	\$276,982.02	11%


Source: IDC Worldwide Quarterly Smart Home Device Tracker, March 2018

A key device and function for the Service Provider to leverage !

Figure 4 – Smart Speaker growth vs other Home Devices

If we look at the trend of usage and skills in the 2 most popular smart speakers – Alexa and Google Assistant, we see from Figure 5 that Music and Audio still tends to dominate with Games and Fun growing fast too. However, you also now see the relevance to Home Control with 10% of the applications for Google Assistant being for Home Control. Additional services like IFTTT also help to ensure that thematic control can be affected across multiple different smart devices and those even developed by different companies. For example issuing a command like “Ok Jarvis Lock the House” could call an IFTTT script to lock every door lock, turn off and dim certain lights and briefly show the 4 outdoor cameras on the TV – before alarming the security system.

Smart Assistants work well with IoT and Smart Home




Skills growing to allow voice control of smart things !

Figure 5 – What people do with their Smart Assistant devices

Figure 6 below shows that Smart Assistants are still primarily being used for immediate audio feedback skills and tasks – like playing music and getting information. However, 15% of consumers are now using them with the Smart Devices – and this aligns almost 1:1 with the consumers who have smart devices in their homes. As the number of consumers who automate their home rises – so will the Smart Assistant’s – as its clear that they are coupled in terms of similar consumer types.

How consumers are using Smart Assistants today?




Applications to interact with other devices and services still emerging !

Figure 6 – What people do with their Smart Assistants

You can see from Figure 7 below that the projected uses of Smart Assistants with Home Control - ranges from Light control to being able to use voice commands with TV navigation and video content selection.

How they would like to use Smart Assistants tomorrow



Natural overlap to Video watching ; Natural use to control certain smart devices !

Figure 7 – What IoT services people most use with Smart Assistants

And there are future opportunities for revenue as part of Service Providers ecosystem to aggregate the different smart devices/services into one cohesive User Experience. New areas like Health and Education

also come into focus for the Service Provider to mine out the opportunities – with partnerships with Insurance, Public and Private Hospitals and digital education solutions

The increase ARPU beyond triple and quad play
Will IoT services generate new ARPU?

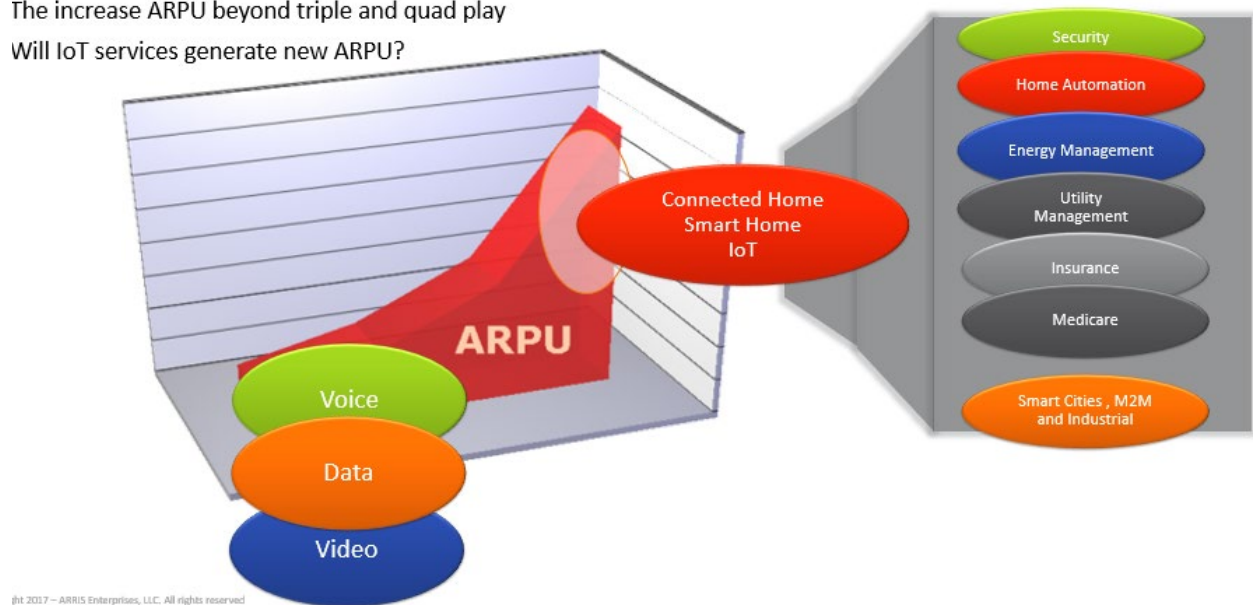


Figure 8 – IOT – its not just about turning on the lights

1.3. Where should Smart Assistants be located and in what physical form?

This is the most interesting and potential best opportunity for a Service Provider to leverage their existing presence in the home. When we think of Smart Assistants – we typically think of an Amazon Alexa device. This device is typically a standalone smart speaker that is added as another device to the home. Its typically designed as shown in Figure 9 and varies in size driven particularly by the Speaker power/volume. Current Smart Assistant devices don't typically perform the radio and protocol functions of an IoT hub and instead focus on working with additional consumer added IoT hubs that speak ZigBee or Z-wave. However, there is a trend growing to add ZigBee and even Z-wave additionally to the existing BLE radio – to allow the Service Provider Smart IoT hub – to support the onboarding customers bought IoT devices. By having all of ZigBee, Thread/Dotdot, Z-wave, IoTivity and BLE all in the smart IoT hub – this pretty much covers the vast majority of consumer bought IoT devices and gives them an IoT hub point as well as a voice assistant/smart speaker capability.

Smart Assistant or IoT Hub

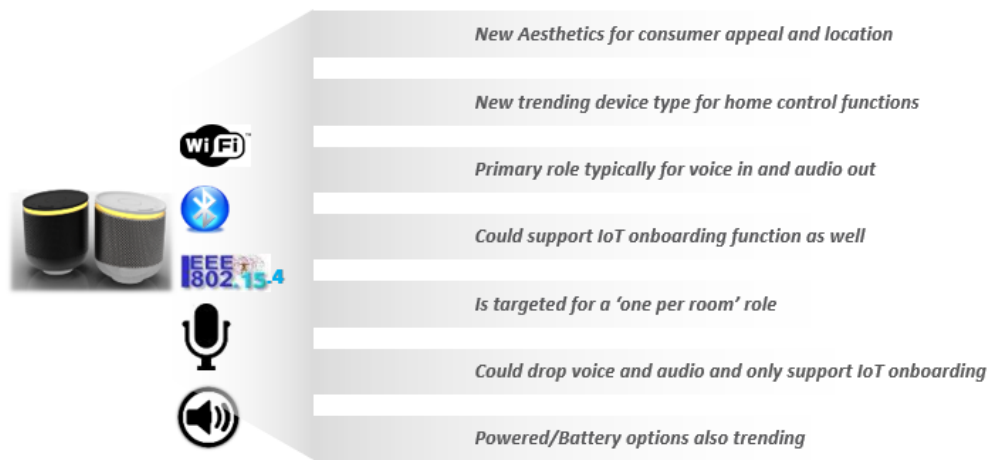


Figure 9 – What drives the design of the Smart IoT Hub

Figure 10 below shows the other potential options open to the Service Provider who already has devices in the consumers home. The Service provider has fought hard to get devices positioned in consumers' homes

- a. Gateway – typically 1 per home
- b. Wi-Fi AP – typically 1 per home
- c. Wi-Fi Extender – emerging now as typically required for at least 40% of US homes – particularly those over 2k sq ft
- d. STB – typically 2.3 STB per US home – all typically located in high footfall rooms

From a STB perspective it also offers unique additional elements of the Smart Assistant potential. In particular

- (i) Use of remote as push to talk voice input and TV speakers as Audio out. Great use cases for voice navigation of video content – but also increasingly being used for other Smart functions.
- (ii) The TV screen - Smart Assistant is typically associated with Smart Speaker where audio feedback is the main output UX. However, there is a growth towards Visual UX feedback that will be key to drive new services.
 - a. You can already see this being introduced with Amazon Dot and Amazon Show – where they have added screens for Visual UX output. A simple reason for doing this is that if a consumer is to buy products with voice – they really want to “see the product” to verify or confirm that it’s the right one before letting the purchase complete.

You can see from what we have discussed that there are some practical considerations for Smart IoT devices

1. The location of the device.
 - a. It needs to be in rooms where people are and can engage with the smart assistant on a regular basis.
 - b. This for example usually means that Gateways don’t make good Smart Assistant investments because many of them are biased towards outside walls and floor because of the location of the Coax outlet or Fiber drop.

2. Making the device too big – particularly with the addition of a high power / fidelity speaker. As we want to get devices like Access Points to be put on table tops – there has to be a balance in size and speaker performance to keep their footprint ergonomic for the consumer.
3. Associating them with TV – this allows the Smart IoT device to either potentially forgo the addition of speakers (at least powerful ones) and leverage the TV speakers -or- to leverage the TV screen for visual UX output and consumer feedback.
4. The number of devices (the capex investment) that are required to fulfill the voice input points in the home and the extension of an IoT hub mesh
 - a. Smart Speaker – could be practical or at least somewhat used in every room of the house. Far Field Mic is typically tuned to be useful to 6-10ft of distance from the device.
 - b. IoT hub – For ZigBee, Z-wave or BLE – it is not required to have a radio hub per Room as the range of these low frequency low power IoT radios spans a single room as well as the IoT end devices themselves typically mesh and add to the range of the network.
 - c. For the typical home of 2.6 people and 2,500sqft it could be an optimal setup to have
 - i. 2 specific Smart IoT hubs for 2 main use rooms in the home – Kitchen and Living Rooms typically
 - ii. Leverage of potentially the STB with additional push to talk or Smart STB with Mic/Speaker added to extend to other media or TV rooms. In particular BLE in the STB for both Remote control usage and IoT hub/presence detection is a very usable feature.

The importance of far field microphones and speaker

Device potential for SA/IoT	Comments	Pros and Cons
Gateway	Tends to bias towards outer wall location	Con : Typically not in best position for voice interaction Pro : One device for smaller MDU/Apartments
Access Point	Has higher chance of being Table top and in room with high footfall	Con : Can make device bigger Pro : One device for smaller homes
Wi-Fi Extender	Typically put deeper in home and in room with people	Con : Trying to make extenders smaller and this adds to size Pro : Can be a 3 for 1 device to service
STB	2.3 devices per home Already many Remote Control 'push to talk' implementations in place	Con : makes the STB bigger Pro : In 2.3 rooms with high footfall
Smart Assistant/IoT device	New device type to	



Figure 10 – Where it makes sense to put Smart Assistant and IoT in the home

Nodes in each room

- Containing
 - Voice assistant
 - Wi-Fi
 - IoT radios
 - Environmental sensors
 - Wireless HDMI

New Services in room

- Potential Services
 - Wellness/Health/Aging in Place
 - Education
 - Home Control Center
 - Environmental sensors
 - Security features




Figure 11 – Smart Rooms not Smart Home

1.4. The STB as the Smart STB with a splash of BLE IoT

As described above the STB has the distinction of being in typically at least 2 high traffic rooms in the US. And the STB is typically, 6-8ft from most people. It also controls the TV as the largest screen(s) in the house – a new canvas for adding visual UX output from voice engagement with home services.

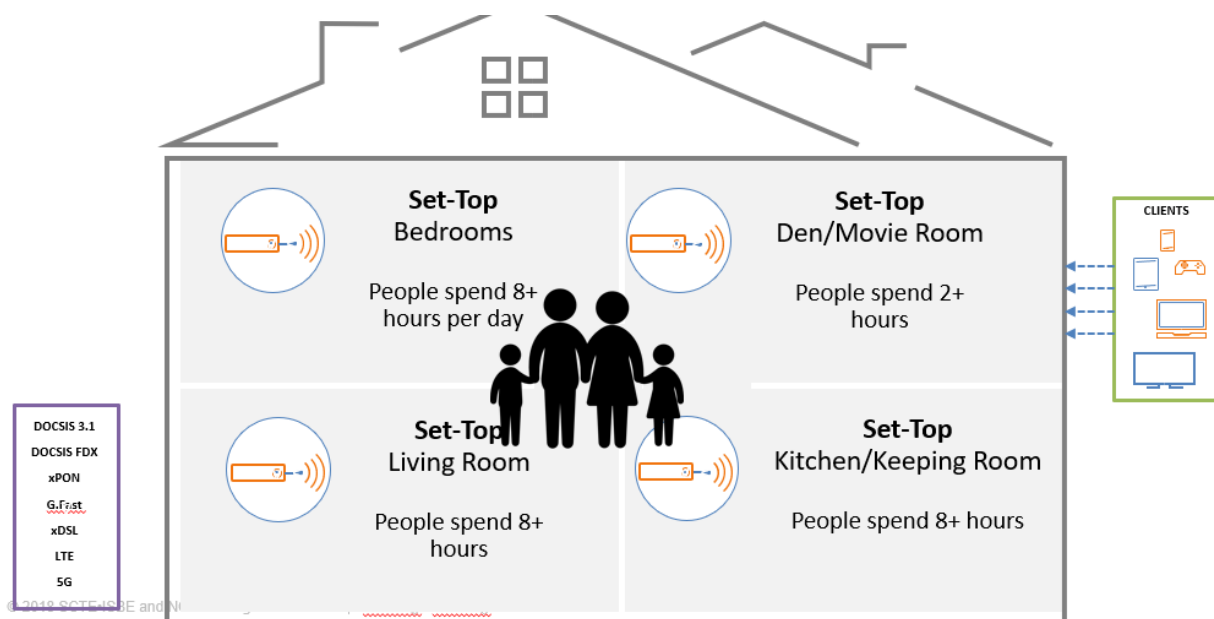


Figure 12 – Why the STB is a key device to make Smart

From a Service Provider perspective – it also has the ability to leverage as a 5 for 1 device which is a powerful capex reducing option as well as helping the consumer with both ergonomics and aesthetics (less devices) and overall power utilization in the home.

The 5 for 1 elements are depicted below in Figure 13

- i. The fundamental function – the STB video function. However, this video display function now also extends to support visual UX feedback from Smart Assistant inputs and Skills
- ii. Smart Assistant – the addition of far field Mics or even push to talk remote with near field Mic – can turn a STB into a smart speaker. The TV providing the speaker functionality. Using the TV as the speaker to the voice input – has one problem – when the TV is off – the smart assistant audio feedback is also off. There are potential workarounds to this
 - a. The TV could be on a smart switch that could power it up. This is not a good user experience waiting the 10+ seconds for the TV to power on
 - b. Using HDMI CEC input – to wake a TV from standby to generate the audio feedback. A better experience at typically under 10 seconds to sound output.
- iii. IoT hub – with the addition of BLE for Remote control and audio streaming the STB could also then provide BLE hub functions to BLE based IoT devices. Consideration could also be given to add either/or both ZigBee and Z-wave to cover as many consumer owned smart devices
- iv. With the extension of a Soundbar with STB integrated – a single device can now cover the video and audio requirements for all services.
- v. Remote Control – the Smart STB could be deployed without a remote – and just use Voice to navigate all services – from Video/TV to IoT services.

Smart Assistant to be built into the predominant UI

Mobile App for IoT applications

- Set-up is still done using an App
- Smart Assistant to run scenes



Remote control for a STB

- Smart Assistant to trigger simple tasks
- Feedback provided on display. Further navigation to be done on remote control or Smart Assistant



Figure 13 – Its not just the phone app – it's the remote and TV too

As can be seen from Figure 2 below – there are probably 4 Smart STB architectures (2 depicted below)

- i. Addition of push to talk remote with near field microphone using RF4CE or BLE – cheapest way to overlay smart assistant functions
- ii. Addition of 4 far field microphones to small form factor STB – to turn it Smart
- iii. Addition of 4 far field microphones + 2-4W speaker to the STB – mass market device
- iv. Addition of Smart STB to a soundbar – Higher End device that appeals to about 20% of consumers or more at lower prices.

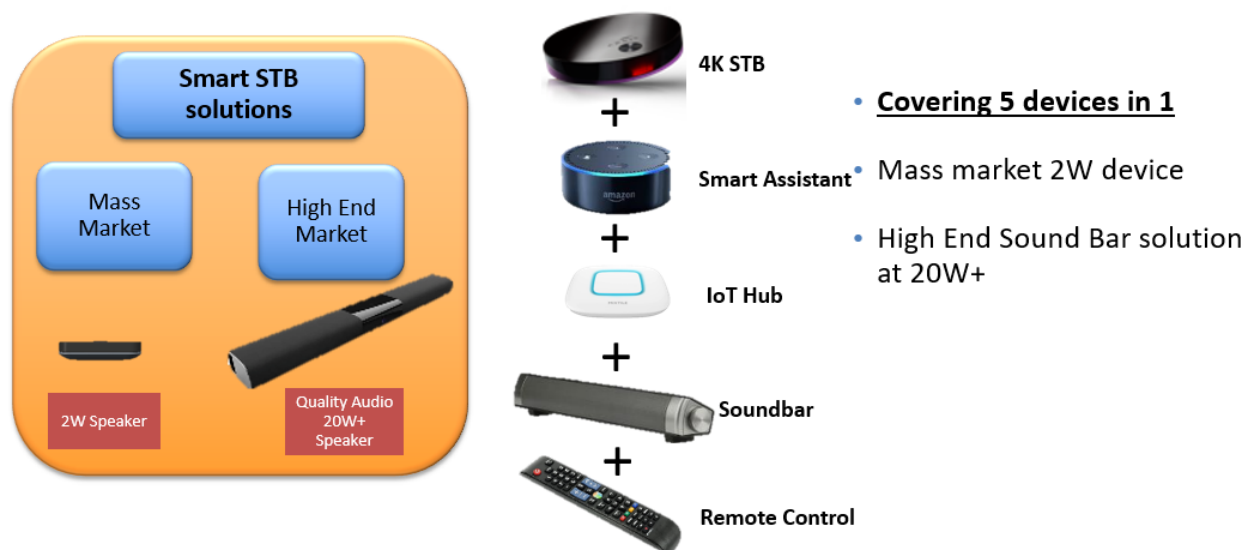


Figure 14 – The 5 for 1 opportunity with the Smart Media Device

2. Which IoT protocol and which Assistant or can there be more than one?

Service Providers can do one thing better than all the major IoT and Smart Assistant players can on their own. They can aggregate the services and the devices into one cohesive experience for the consumer.

From the IoT perspective the following is the high level take on the Radio protocols

- i. 802.15.4 – ZigBee
 - a. ZigBee has a very mature and strict data model for its IoT ecosystem. This makes it straightforward for a Service Provider Hub to control third party ZigBee devices and customer owned and managed (COAM) devices.
- ii. 802.15.4 – Thread
 - a. Thread also runs on 802.15.4 and has a well enough defined data model to easily aggregate Thread based devices.
 - b. ZigBee and Thread coexist and the Dotdot standard supports ZigBee over IP
- iii. BLE
 - a. Bluetooth is more difficult to take ownership of COAM devices. The BT data model is not strict enough that its often required to develop specific code on a device per device basis to guarantee its interoperability
- iv. Z-Wave
 - a. Strict data model that allows it to support Service Provider aggregation.

Additionally, the Cable industry is also supporting the Open Connectivity Foundation and the Opensource standard IoTivity. The OCF absorbed both the UPnP and Alljoyn assets in the last 3 years and has a S/W architecture that works very well with the aspirations and future directions of the Service Provider. Using the IoTivity Server and Plug-in support it has for Protocols like ZigBee, Z-wave, Thread and BLE – it makes a powerful interface layer to be able to host an IoT protocol engine at the same time supporting RESTful cloud interfaces and protocols like COAP, MQTT and even the RDK-B WebPA protocol.

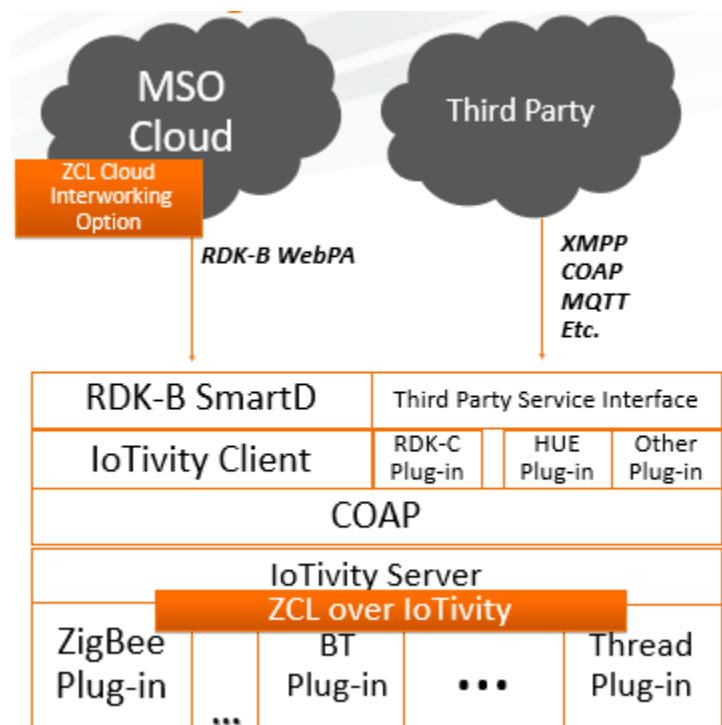


Figure 15 – The Simple Cable IOT Stack

The use of a Smart Assistant ASR and NLU and Skills Framework/Store is the more interesting discussion.

Options exist to

- i. Develop your own ASR and NLU and support your own specific Wake work in your own specific Smart Devices
 - a. This is the most expensive option and requires a large development team with specialized skills in Voice processing and definition of Natural Language understanding and related skills and actions.
Very few if any Service Providers will go this route – because of expertise deficit, expense to implement or contract out. Complexity of Language support does sometimes drive towards developing specialized regional ASR solutions.
The benefit of doing this – is that you own all the data mined with Smart Assistant interactions.
The negative is that you must implement everything including Web searches and try and bring up a comprehensive skills solution.
- ii. License ASR from independent ASR solution providers, license NLU for Web queries (“What is the weather today”) and develop or contract your own NLU for specific Service Provider Skills (“Whats my Wi-Fi password, Buy more broadband”). These skills require integration into the Service Providers own backend.
 - a. This is a path followed today by several Service Providers today – especially for TV navigation services and basic web queries with voice.
The benefit of doing this is that you broaden the Voice Assistant skills (license costs are expensive) but keep your customers analytics.
- iii. Decide to leverage one of Google, Amazon, Cortana, Watson, Bixby, others – solutions.

- a. Develop skills for these ecosystems that drive Service Provider services but through the presence of the third-party device
 - i. For example, develop Alexa Skills to navigate TV or other SP services. Add this skill to the Alexa Skills database.
- b. Add the ASR or NLU to a Service Provider device – standalone Smart Assistant or integrated into existing device like AP or STB

Both these options leverage the very broad and feature rich solutions of each ecosystem – but offer the ASR/NLU provider the analytics of what the consumer is doing. The debate then is to whether allowing one of these companies to engage with ‘your’ (Service Provider) customer is a threat to your own service directions. While the ASR/NLU providers claim they maintain privacy and don’t really use the data – this is debatable as there is at least trend and frequency patterns that also offer consumer insights.

- iv. Possibly the best option (and potentially unique to Service Providers) is one where there are multiple Smart Assistants in the Service Provider device.
 1. The scope exists to be able to add multiple wake words to the smart assistant device
 2. Based on the wake word selected
 - a. Alexa – invoke Amazon ASR/NLU – could add Service Provider Skills but don’t have to – “Alexa Tell <Service Provider> to..”
 - b. Ok Google – invoke Google Assistant ASR/NLU and tasks – could add Service Provider Skills but don’t have to “Ok Google Tell <Service Provider> to..”
 - c. Ok <Service Provider Name Here> whats my Wi-Fi password.
 3. For the Service Provider specific path – the following elements are in play
 - a. Only implement ASR and NLU for specific value add SP skills
 - b. Keeps the analytics and data from the Amazon and Google clouds

See the Figure 16 that illustrates this below

There can be more than one ...

- Leverage Alexa (and others) for the ecosystem they support – including generic enquiry skills
- Build specific NLU for your Backoffice, strategy, customer support and services
 - Customer Support and Self Healing
 - Broadband and TV UX integration

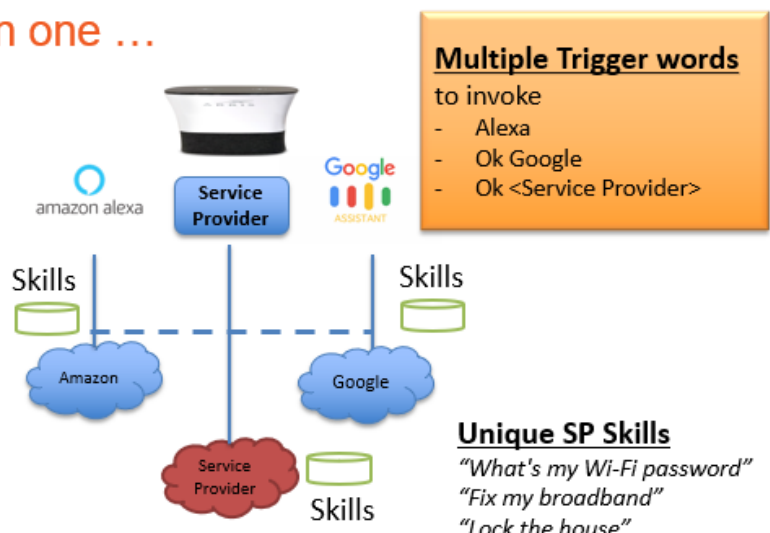


Figure 16 – Multiple Smart Assistants are possible.

Figure 17 below denotes the flow control path from Voice DSP to the various ASR and NLU elements.

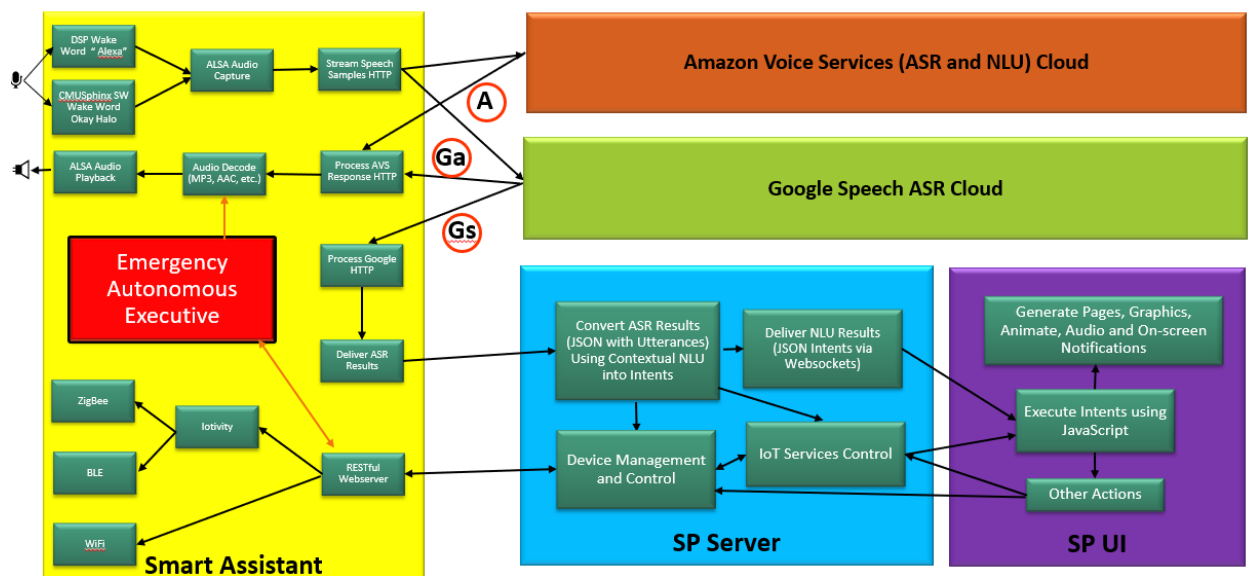


Figure 17 – The flow of voice from Mics to ASR

Examples of the unique set of Service Provider skills include

- *OK Service provider, what is my Wi-Fi password?*
- *..., how fast is my broadband?*
- *..., can I get faster broadband?*
- *..., how much more will I pay for faster broadband?*
- *..., upgrade my broadband*
- *..., open my guest network*
- *..., secure my guest network,*
- *..., what unknown devices are connected to my net?*
- *..., shut down my son's Wi-Fi from 8PM on weekdays*
- *..., what parental controls are active?*
- *..., what devices have poor connection?*
- *..., fix my wife's iPad connection*
- *..., show my front door webcam on den TV*
- *..., block my webcams from the Internet*
- *..., connect me to support center*

And one of the most opportunistic areas to apply Smart Assistant services is to provide customer self healing and interface to Chatbot and future AI services for customer support.

It is not hard to imagine a scenario where the consumer does not use to the phone at all to contact the Customer support desk. Instead it uses smart assistant to ask ‘Ok Service Provider – I’m not happy’ ; ‘Ok Service Provider – my Wi-Fi sucks’ ; ‘Ok Service Provider Fix my iPad’ ; ‘Ok Service Provider – upgrade me to 1Gbps service’

These are powerful new automated features that can reduce customer call center Opex, increase consumer NPS scores and also potentially unlock new pay for services opportunities.

There is still ongoing debate and dialog around

- a) Are Service Providers compromised by integrating with large company ASR/NLU – and are they competing for the same consumer?
- b) Is there opportunity for both Service Provider and ASR/NLU provider to profit together? Certainly, the aforementioned ASR/NLU companies are offering SDK and Cloud API’s to integrate their ASR/NLU Solutions.
- c) What is the risk of sharing the analytics from Consumers with the ASR/NLU providers?
- d) Is there a hidden cost associated with using these -on the surface- free to use – ASR/NLU SDK’s?
- e) Will consumers be confused with different wake words and different Assistant persona’s in the same device?
 - a. Some point out that we already do it today – with Siri on phone and Alexa at home – and we are already distinguishing what each assistant does for us.
 - b. Some will point out that who owns the responsibility if something goes wrong? If you ask the multiple persona smart assistant to “Lock the house” and it is not done. Which one is responsible or is there some accountability ambiguity.

There is also some reluctance of the ASR/NLU to also certify solutions that support Wake words that don’t invoke their Assistant. Amazon at this point – will certify devices that support other assistant’s resident in parallel to their own Alexa.

Google, at this point are not certifying devices with multiple Smart Assistant Personas. They are reviewing the scope to boot to one from several options presented but not all resident together. This is something our industry should lobby to Google to open up this multiple wake words potential and compete on their own merit of consumer wanting to use there services.

One last point to make here – is that several years ago – the Cable industry was very wary of the rise of Netflix and viewed them as a competitor to keep at arm’s length. As more Netflix end points were updated with Netflix App downloaded to Smart TV and OTT STB – it became clear that the best strategy for Service Provider and consumer alike was to integrate Netflix as just another video source – with deep metadata integration being an even better user experience. This has proven to keep the consumer on the Service Providers device even when in Netflix app and always returning to the same HDMI port that the SP’s own services run on.

The same analogy could be applied to Smart Assistant ASR/NLU providers – that integrating them into Service Providers own solutions offers the control still to the SP vs consumers (1 in 3 has one now) adding devices to their homes that the Service Provider is blind on what is being asked from the home.

Conclusion

As the highest growing CPE device – the smart assistant remains a function that the Service Provider must embrace. Couple the growing trend in Smart Home devices – it seems like there is a perfect storm brewing of the interception of IoT and Smart Assistant. This provides the option to put them together – probably even in the same device based on their almost 1:1 relationship and their coverage of voice input and IoT mesh coverage in the typical home. There are certainly lots of decisions to make – like speaker size or whether to integrate into STB. However, the most important decision to make is on ASR/NLU selection. This paper outlined the potential to have more than one – and to leverage the investment of other companies on NLU/ASR to allow the service provider to implement key , critical Skills that are specific to them and their customers – affording the best leverage of Smart Assistant services for their customers.

Abbreviations

ASR	Automated Speech Recognition
NLU	Natural Language Understanding
BLE	Bluetooth Low Energy
ZCL	ZigBee Cluster Library
RDK	Reference Design Kit
SDK	Software Development Kit
AP	Access Point
STB	Set Top Box
OCF	Open Connectivity Foundation
SMD	Smart Media Device

Bibliography & References

Source : “Voice-First Experience” – Parks Associates

http://www.parksassociates.com/bento/shop/whitepapers/files/Parks%20Assoc%20Enabling%20Voice%20in%20the%20Smart%20Home_WP.pdf

Source : “Wordwide Quarterly Smart Home Device Tracker, March 2018” – IDC

https://www.idc.com/getdoc.jsp?containerId=IDC_P37480

Source : Voicebot.ai <https://voicebot.ai/>

Source : “Impact of Voice on Connected Consumer Markets” – Parks Associates

<https://www.parksassociates.com/whitepapers/voice-may2017>

It's ALIVE! Getting to Successful R-PHY Deployment: Do's And Don'ts

A Technical Paper prepared for SCTE•ISBE by

Tal Laufer

Director, Product Line Management
ARRIS
3871 Lakefield Drive, Suwanee GA 30030
+1 470 326 8077
tal.laufer@arris.com

Jeroen Putzeys

VP, Sales EMEA
ARRIS
3871 Lakefield Drive, Suwanee GA 30030
+32 478 662065
Jeroen.Putzeys@arris.com

Uffe Callesen

Technology Architect
Stofa
Slet Parkvej 5-7, Denmark
+45 51328435
ufca@stofa.dk

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
The Story Unfolds.....	4
1. Market	4
1. Cable Market Driver – the "Need for Speed"	4
1.1. Distributed Access Architectures (DAA)	6
1.2. DOCSIS 3.1 Introduction.....	8
2. Company (Stofa) overview	9
2.1. Company overview.....	9
2.2. Network architecture at the starting point	10
2.3. Motivation for network upgrade.....	10
3. Operators Drivers for R-PHY, and Stofa's Specific Drivers	10
3.1. Drivers and benefits of Remote PHY	10
3.2. Stofa's Drivers for R-PHY.....	11
4. Planning	12
4.1. Network architecture	12
4.2. Planning the fiber and interconnect network.....	13
4.3. RF Plant Upgrade and Updating Network Inventory System.....	15
4.4. CIN planning.....	16
4.5. IEEE1588 Timing	17
4.6. Product Qualification Phases	18
4.6.1. Initial Testing and Solution Evaluation	18
4.6.2. Integration into the Network	19
4.7. Automation	21
5. Deployment and Operational Results	23
5.1. Field Deployment	23
6. Performance Improvement.....	24
6.1. RF performance	24
6.2. Space & Power in the Headend	28
6.3. Operational Simplification	30
Conclusion.....	31
Key Benefits	31
Key Challenges	32
Major Takeaways	33
Abbreviations	34
Bibliography & References.....	35

List of Figures

Title	Page Number
Figure 1 - Nielsen's Law of Internet Bandwidth (Growth Rate = 50%/YEAR for high end subs).....	5
Figure 2 - CCAP/Remote Architectures Market Revenue Forecast	6
Figure 3 - Centralized Access Architecture Diagram.....	7
Figure 4 - Distributed Access Architecture – Remote PHY Diagram.....	7
Figure 5 - R-PHY Internal Components	8
Figure 6 - Remote MAC PHY Architecture Diagram.....	8
Figure 7 - DOCSIS 3.1 Deployment Forecast (Source: ABI Research)	9
Figure 8 - Stofa Selected Reference Area	13
Figure 9 - Reference Area with DWDM Solution	14
Figure 10 - CIN Network Design	16
Figure 11 -PTP Latency Calculation	17
Figure 12 - Timing Distribution Network.....	18
Figure 13 - Stofa Initial Lab Setup	19
Figure 14 - PNM Topology View	20
Figure 15 -Software Tool for US RF Network Alignment and Monitoring	21
Figure 16 - Back office Management Architecture.....	22
Figure 17 - RPD Manager Onboarding	23
Figure 18 - US and DS channel SNR distribution	27

List of Tables

Title	Page Number
Table 1 - DOCSIS Evolution increases HFC network capacity (Source: CableLabs)	9
Table 2 - Upstream and Downstream SNR Measurements	24
Table 3 - Space and Power required for M-CMTS and R-PHY Devices in the Headend.....	29
Table 4 - Summary of Space and Power Required for Legacy and R-PHY Architectures	29
Table 5 - Number of Managed Devices in Different Architectures.....	30

Introduction

This paper is an operational overview of one of the first Remote PHY deployments in the world, at a Danish operator – Stofa. Being an early adopter, the entire process of selecting the technology, planning the new network design, and deploying the products was an uncharted territory. Stofa and ARRIS, the selected vendor for the project, have learned many lessons from the experience, which are all detailed in this paper.

In addition, we have analyzed the benefits that the Remote PHY network upgrade has provided to the operators, including measured improvements in signal to noise ratio in the field and space and power saving in the headend.

The results we present demonstrate the great benefits transitioning to Remote PHY can achieve. First, there is a clear improvement in SNR values in the upstream (US) direction, which will allow Stofa to go to higher modulation orders using DOCSIS 3.1 US. In the downstream (DS) direction, results are less conclusive. The transition to R-PHY was accompanied with a move of DOCSIS D3.0 channels to higher part of the spectrum, which may have countered the improvement. In addition, some elements in the network like the drop cable were not replaced, and may affect the measured signal-to-noise ratio (SNR) at the cable modem (CM) side.

Second, we analyzed the space and power saving in the headend that resulted from the migration from the previous network architecture (Modular-CMTS) to Remote PHY. The analysis showed significant space savings – going from about 70 RU to 18 RU in the headend for the MAC core and network support, and less than half the power required in the headend for supporting these functions. In addition, the operational complexity is dramatically reduced, due to the major reduction in devices that need to be managed, and the automated management that is introduced. As part of the project, Stofa's ability to do proactive network management actions is greatly improved, as well as their ability to effectively grow their network and customer offering without any truck rolls or field work.

All these benefits will be reviewed in the paper, as well as the steps Stofa took in order to prepare for the deployment, and successfully deploy the new technology in the field, serving thousands of subscribers.

The Story Unfolds...

1. Market

1. Cable Market Driver – the "Need for Speed"

It cannot be denied that the "need for speed" is dominating our cable operators market at a growing rate. Operators are very mindful of the highly competitive landscape they operate within, and about the need to increase some of the subscribers' bandwidth by 50% annually. We can see the projected bandwidth demand growth according to Nielsen's law in Figure 1 below:




Figure 1 - Nielsen's Law of Internet Bandwidth (Growth Rate = 50%/YEAR for high end subs)

In order to deal with that rapid growth in the demand for bandwidth, operators are searching for new technologies and architectures that can help them supply those speeds and new services, using an efficient scalable design, with predictable and controllable cost of ownership.

Operators may consider making changes to their headend network, their nodes and amplifiers, their Service Groups sizes, the modulation profile used on different areas of the plant and more. The challenge for an operator is to choose the right mix of adjustments that can help them optimize their network and supply their subscribers' demand.

In this paper we will review the case of Stofa, a Danish operator that has chosen to introduce DOCSIS 3.1 and Remote PHY architecture into their network, to handle those mentioned challenges.

Many operators are seeing similar challenges, and therefore we are seeing a significant change in the cable access market, while operators are considering their next technological upgrade and their future network evolution.

In Figure 2 we can see the S&P Global (Kagan) forecast for market transition:




Figure 2 - CCAP/Remote Architectures Market Revenue Forecast

This chart shows clearly the expected transition to Distributed Access Architectures (DAA), which include Remote PHY and Remote MACPHY. Starting gradually from this year with early trials and limited deployments, we can see a projected growth in the penetration of the new technologies starting significantly from 2019 and onwards.

We can attribute this timing to multiple factors:

- **Product Readiness** - Equipment providers have started to offer Remote PHY products already early this year, and they will get some experience with initial early deployments [Silbey]. However, some operators are choosing to wait for the technology to mature, or for a specific flavor of the products to become available before they begin deployments.
- **Network Planning** - Architectural upgrades such as the transition to DAA require many planning phases and a lot of thought. As will be discussed later in the paper, operators have to go through the considerations of which access architecture is the best fit for them, and they must also go through detailed planning activities in order to be prepared for the transition from a network readiness perspective, as well as from an operational perspective.
- **Existing Unrealized Capacity** – Many operators are currently using dense integrated CCAPs, deployed in the last few years. Many of these systems still have unused capacity that can allow for services expansion with licenses addition only. So, for those operators the need to upgrade their network may be less pressing. If they do upgrade, then they may opt to utilize architectures that permit them to re-use their recently-deployed equipment.

1.1. Distributed Access Architectures (DAA)

Distributed Access Architectures represent an evolution in the cable access network structure and operations. The drivers for DAA (and specifically Remote PHY) are detailed in section 1.c. They include the desire to bring fiber closer to the home (in order to be better prepared for FTTx architecture), the need to reduce space and power in the headend, and the drive to provide higher bandwidth to the subscriber.

In order to do that, the cable industry designed a variety of Distributed Access Architectures, where parts of the traditional integrated CMTS or CCAP are moved to the node structure, closer to the subscriber. The

different architectures vary in the amount of functionality that is being moved to the node, and the changes the operator will be making in their headend design, and specifically the video portion of that.

If we start by looking at the **Centralized Access Architecture** in Figure 3, we can see that all CMTS and CCAP functionality is centralized in the headend or hub, supporting high speed data, voice and video services. The processed radio frequency (RF) signals are transmitted to the optical node via analog carriers over fiber, where they are converted to analog signals over coax.




Figure 3 - Centralized Access Architecture Diagram

With **Remote PHY** in Figure 4, the PHY part of the processing is moved to the node, meaning the QAM modulation, FEC and DAC/ADC. The headend equipment is responsible for the MAC processing, and is transmitting the MAC-processed signals to the node via digital optics the Converged Interconnect Network (CIN).




Figure 4 - Distributed Access Architecture – Remote PHY Diagram

The internal components of a Remote PHY system include (per [R-PHY-spec]):

- Ethernet Interface
- Clock Circuitry
- Remote PHY Pseudo-Wire Interface
- Common Layer 1 PHY Circuitry

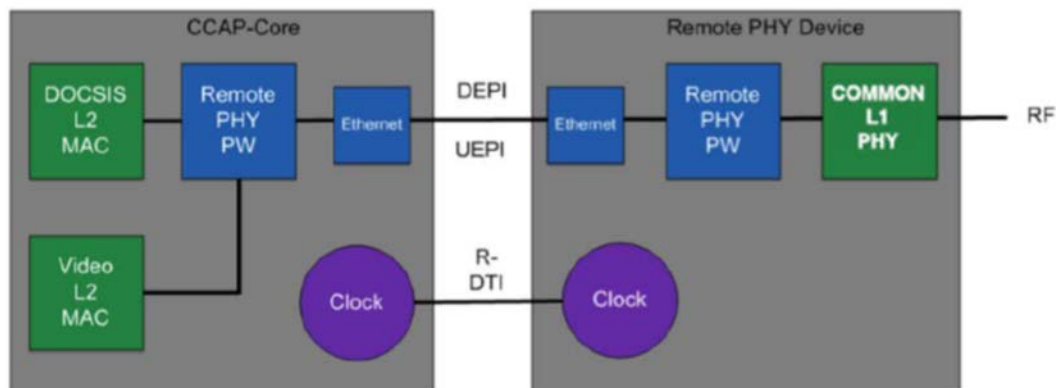


Figure 5 - R-PHY Internal Components

As can be seen in Figure 5, there is a clock component at every device, and those clocks need to be synced, to allow for proper operation of the Remote PHY solution. The synchronization requirements are detailed in the DTI spec that is part of the Remote PHY set of specifications [R-DTI Spec].

Another type of DAA, still being defined under a specification process, is "Flexible MAC Architecture" (aka Remote MAC and PHY, also sometimes referred to as Remote CCAP). In this flavor of DAA, the traditional functions of a CMTS (both MAC and PHY) are moved to the node. The data is transmitted to the node from the north bound router over IP network.




Figure 6 - Remote MAC PHY Architecture Diagram

We will not discuss the transition to Remote MAC PHY, depicted in Figure 6, within this paper.

1.2. DOCSIS 3.1 Introduction

DOCSIS 3.1 is the relatively new broadband data specification, designed to increase the DOCSIS services capacity on the existing HFC networks. DOCSIS 3.1 relies on:

- OFDM
- LDPC
- Energy Management
- Hierarchical QoS
- Active Queue Management
- Advanced Timing Support.

The capacity D3.1 is expected to offer is summarized in the Table 1.

Table 1 - DOCSIS Evolution increases HFC network capacity (Source: CableLabs)

	DOCSIS 3.0	DOCSIS 3.1
Upstream	0.1 Gbps	1-2 Gbps
Downstream	1 Gbps	10 Gbps

Since the DOCSIS 3.1 spec completion, in 2013, we are seeing gradual deployment of D3.1 services. Some reasons for the slow adoption are: Headend equipment readiness, CPE device availability, and compatibility issues between the two given it is a new spec. Some operators are also delaying the DOCSIS 3.1 introduction due to the large investment required for CPE device upgrades.

Figure 7 shows the forecast for D3.1 services deployment.




Figure 7 - DOCSIS 3.1 Deployment Forecast (Source: ABI Research)

2. Company (Stofa) overview

2.1. Company overview

Stofa is located in Denmark and is delivering broadband services to about 500K cable households. Stofa is Denmark's second largest provider of television, broadband, and telephony.

Stofa's customers are mostly antenna associations – organizations responsible for providing infrastructure services to a community of usually a few thousand subscribers. The services they are responsible to provide include high speed data, voice, and video (legacy QAM video) in high quality to their subscribers.

2.2. Network architecture at the starting point

About 60% of Stofa's footprint is rural, but there is a strong concentration of MDUs in Denmark's major cities. As is typical in most of Europe, cable runs are buried rather than aerial (as is largely the case in the US outside urban areas).

Prior to the network upgrade project discussed in this paper, Stofa had been deploying a Modular CMTS (M-CMTS) architecture. Stofa had deployed the CMTS in their headend, and the associated EQAMs were usually deployed in the antenna facilities location.

Stofa has been deploying services using an N+5 plant configuration, meaning 5 levels of amplifiers after the fiber node. Before the upgrade project, for their high-speed data service groups, they had been deployed with a 1 downstream: 2 upstream node combining ratio.

Stofa has fiber links deployed between their headends and regional hubs. Their topology is such that the typical maximum distance between the headend and the node location is 10 km.

2.3. Motivation for network upgrade

In 2016, Stofa decided to launch a project for planning of a network upgrade, and started searching for the best fit in terms of network architecture and products. The Stofa drivers for the network upgrade included:

- A. **Market competition** – growing competition in the Danish market was driving Stofa to look for advanced solutions to increase the bandwidth offering to the subscriber. Stofa chose to upgrade their network to DOCSIS 3.1, and go from a 860 MHz plant to a 1.2 GHz plant for DS, and up to 204MHz in the US. These upgrades were targeted to increase the capacity Stofa planned to offer their customers on both the US and the DS directions: first, by transitioning to D3.1 and reaching a higher modulation order; and second, by increasing the amount of spectrum available for the different services.
- B. **Transition to digital fiber** – The year-on-year growth in both physical fiber count and optical wavelength division multiplexing (WDM) filters to support the growing number of optical nodes was becoming a real issue for Stofa in regards to physical space at the major headends. Either much better utilization of existing fibers or new larger headends was needed. Digital fiber solves this issue by providing the benefits of not only multiplexing many nodes onto the same fiber (not completely unlike WDM) but also by providing the benefits of statistical multiplexing (Customer A not using bandwidth at the “exact” same time as Customer B). The use of the IP protocol on top of Digital fiber also enables dynamic sharing of traffic across different fibers - something that was unthinkable in the past.
- C. **Consolidating hub sites** – Stofa was challenged by the space and power in the hub and headend locations. The operational cost of those facilities were a burden on the Stofa budget, and they were looking to reduce the space required for network infrastructures, and longer term they planned to consolidate many of their headend's into a centralized data center.

3. Operators Drivers for R-PHY, and Stofa's Specific Drivers

3.1. Drivers and benefits of Remote PHY

Potential drivers for DAA architecture include (also according to [Cloonan]):

- A. **Transition to digital optics** – with DAA, the transmission from the headend to the nodes is done using digital signals on the fiber network, as opposed to the centralized architecture, where it is

done over analog signals on the fiber network. The digital transmission allows for more lambdas to be populated over one fiber. With digital fiber, an operator can use up to 80 lambdas on one fiber, whereas with analog optics the maximum is about 32. This allows the operator to better utilize their existing fiber network and drive more bandwidth using it. Digital fibers also allow for greater distance between the headend or hub and the node, allowing the operator to centralize their MAC core functionality in data centers feeding remote nodes. In addition, operation of digital fibers is perceived easier than analog fibers, that may require frequent tuning and adjustments. The transition to IP brought in many operational benefits, and significantly increased the fibers' capacity.

- B. **Headend space and power reduction** – DAA's main principle is about moving functionality out of the headend and hub, and down to the plant. The reduction of functionality in the HE drives lower space and power required for the processing functions, allowing for operational saving on the facility maintenance, and potentially hub consolidation. See previous study conducted on space and power requirements for the different access architectures [HFC-Green-ULM].
- C. **Better SNR at the "end of the line"** – DAA moves RF signal processing from the headend where it is transmitted to the node with AM fiber, and puts the RF processing in the node eliminating the potential signal degradations from AM fiber, and can improve SNR resulting in the use of high orders of modulation with more data transported over the same bandwidth
- D. **Facility consolidation** – As discussed, Remote PHY can help with facilities consolidation, due to space and power saving. In addition, R-PHY allows for the headend/hub to be more remote from the node and subscriber, since the fibers used to carry the data are carrying digital signals, which have less distance limitations. So with R-PHY, the facilities can potentially be consolidated to reduce facilities and maintenance costs for the network.

Additional benefits may apply specifically to Remote PHY:

- E. **Ability to reassign MAC Processing** – when using R-PHY, the MAC processing capabilities are centralized in the MAC core. The Remote PHY Devices (RPDs) can dynamically be moved from one core to another per need, potentially for load balancing purposes. This allows the operator better resource efficiency and flexibility.
- F. **Ability to select best-in-breed from CCAP Cores and Nodes** – Remote PHY spec by Cable Labs [R-PHY Spec] is designed such that cores and nodes from different vendors can interoperate, when they are both complying to the spec.
- G. **Little change in the provisioning, configuration and management systems** – Remote PHY architecture defines that provisioning and configuration of the nodes be done from the MAC Core. Having one centralized point of configuration is similar to the way DOCSIS networks operate today, so minimal changes have to be done in back office systems and operational models.
- H. **Better path towards virtualization** – centralizing the MAC Processing in one location creates a better path towards virtualization, since the MAC processing function is easier to virtualize. In the future, physical appliances such as the MAC Core can be virtualized and moved to off-the-shelf servers, which will further reduce the cost, space and power, and enhance the flexibility of resources assignment.

3.2. Stofa's Drivers for R-PHY

Of the potential benefits of Remote PHY, a few were more impactful in Stofa's decision to migrate their network to remote PHY.

The headend space and power saving, along with potential facilities consolidation were very important to Stofa. The cost associated with maintaining their headends and hubs is a considerable part of their OPEX, and Stofa were looking to both downsize in the headends, and also remove processing equipment from the local antenna associations / communities.

Better signal quality or SNR at the subscriber location was another driver for the project. Stofa's market is very competitive, and Stofa had to improve their existing capabilities to support the market's demand. Better SNR will allow Stofa to migrate to D3.1 and materialize the benefits it will provide with higher modulation orders, which will create more capacity to the subscriber without changing the "last mile" wiring.

A strong motivator for Stofa to make the large resources investment in the project is the future expansion of the network as forecasted from their traffic engineering numbers. In the coming years, Stofa will have to continue to expand their services, growing the capacity provided per subscriber, in order to remain competitive. The current investment in the new technology (with a lot of room to grow) will allow Stofa to add only licenses in the next network upgrade cycles. They will not need to send technicians to install anything, but rather will only have to remotely change the configuration on the nodes already in the field.

In addition, Stofa likes the R-PHY architecture that allows them to keep the expensive and more complicated equipment (hence MAC Core) in the headends, which are owned by Stofa. This allows them better control over the installation environment, and a future path toward virtualizing these functions and centralizing them in data centers.

The centralization of the MAC core processing also has other benefits for Stofa.

First, the single provisioning interface (for a MAC core) allows Stofa to do minimal changes in their provisioning system, and back office tools. Some automation will be required to support the R-PHY nodes, but other changes are minimal which reduces the cost and complexity of this upgrade project. In addition, it also simplifies the training the Stofa personnel will have to take, in order to manage all configuration items on "one box". That applies for both the engineering teams and the operational teams debugging issues on the system.

Lastly, the Remote PHY solution Stofa has chosen allows them to converge all customer services on one platform: DOCSIS, VOD, and broadcast services. All these are processed by the MAC Core and being transmitted over digital fiber to the R-PHY node. The convergence allows for operational simplicity (since data path is unified for all services), and Stofa is also benefiting from sending the video signals efficiently over digital fiber. Other architectures for video transmission exist in the market, but Stofa decided to go with the converged one because of the operational simplicity it offers.

4. Planning

4.1. Network architecture

After it was decided to deploy a converged video and data Remote PHY in the Stofa network, the focus shifted towards crafting the right design that will provide the benefits of R-PHY, and match the Stofa network and requirements.

R-PHY comes with new requirements to the IP network. As the remote-PHY architecture is literally separating two physical elements that were previously on the same circuit board, the network has to provide the same stability and fixed low latency for the solution to work over the existing fiber network.

Replacing the classic analog optical equipment means that a lot of focus was put on the transceivers, and not just for the fibers going to the RPDs. As classic headends turn into IP hubs instead, the need for 100/200G backhaul links needs to be addressed. Stofa was also mindful of temperature limits for the transceivers going into the RPD's.

For the actual design approach – there is no "one size fits all", as optical network design will vary from operator to operator. The following briefly describes the design variants Stofa evaluated.

As previously described, Stofa has many headends in rural areas, typically with a very limited number of fibers connecting them to the central network. In order to select the optimal network architecture, Stofa chose one topology representing a geographic area that includes typical headend sizes and distances between them, as described in Figure 8.

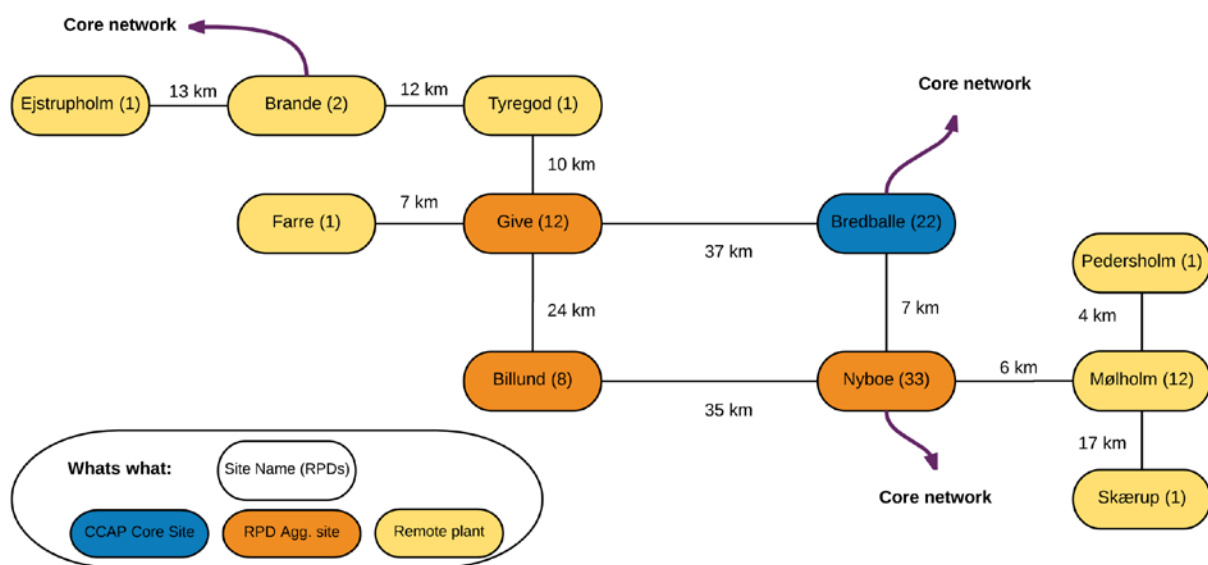


Figure 8 - Stofa Selected Reference Area

4.2. Planning the fiber and interconnect network

Stofa reviewed various options for interconnecting the elements of their network design, as detailed in the section below.

Stofa considered upgrading the existing Point-to-Point fiber network to a DWDM network in order to connect the RPDs to aggregation sites where the MAC cores reside, as described in Figure 9.

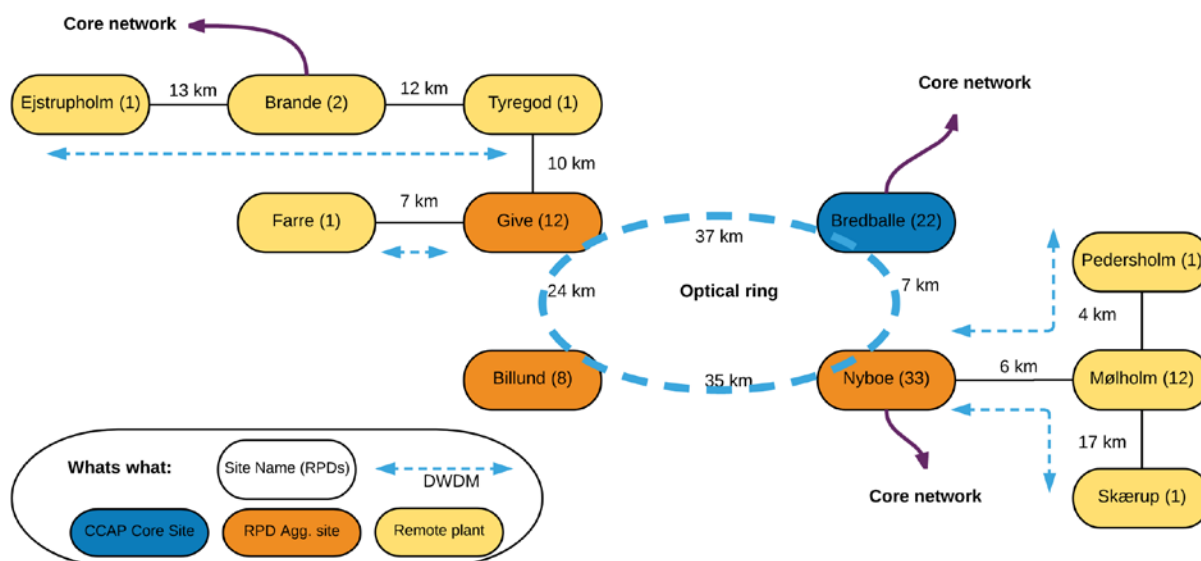


Figure 9 - Reference Area with DWDM Solution

A clear benefit of this approach is that active equipment (including everything belonging to the IP network) could be consolidated to fewer headends. This consolidation also means less investment in costly 100G/200G backhaul transceivers as each 10G RPD link uses a dedicated 10G wavelength over DWDM using 10G DWDM transceivers in both ends. One last benefit is that using DWDM won't drastically change the day to day operations of fiber plant and nodes.

However, many challenges were identified as well. As Stofa has limited fiber between the existing headends, a 50 GHz / 80 Channel DWDM system would be needed in most areas, and in some cases even 80 channels would not suffice (where more than 80 RPDs each using one DWDM Channel, needed transport across a single set of fibers). The range of the DWDM transceivers would also be problematic in some rural areas, requiring optical amplification at some sites that were initially designed as passive sites. On top of all that, the complexities of ensuring redundancy in the DWDM network and the much higher cost of DWDM transceivers compared to 10G Long Reach (LR) optics were additional concerns for Stofa.

Hence, the DWDM solution was not an ideal fit for Stofa. The next logical approach would have been to design a single-service network to support just the RPDs. In a layer 2 (switched) design, guaranteeing low latency and symmetrical traffic flows in the network is relatively simple. Operators have to deal with the known limitations of layer 2 networks, like the number of supported VLANs, spanning-tree for redundancy etc. Nevertheless, this approach is feasible.

Despite its benefits, a layer 2 (L2) network was not considered as there are some serious downsides. Stofa, like many other operators, is going to be using the converged interconnect network (CIN) for other services as well, making it impossible or overly-complicated to use L2 in practice. For L2, an operator may even need to build this as a completely isolated network – not benefiting the rest of the business.

Another strong argument against the use of an L2 network in the Stofa network was the following. With DAA, an operator will be aggregating IP traffic where their fiber infrastructure is aggregated. This leads to physical sites with massive backhaul capacity requirements and that is probably the biggest downside of deploying DAA over L2 networks – an operator will likely need a lot of expensive 100G+ transceivers that will be poorly utilized in an L2 network since effective load balancing of IP traffic requires IP routing which in turn requires a L3 based network.

For Stofa the preferred design option was a converged IP (L3) network that supports all existing services and fulfills the requirements for R-PHY. Being an early adopter, this did not leave Stofa with many options other than to study roadmaps and choose the best future solution that complies with the network requirements (Further details of CIN design in a later section).

4.3. RF Plant Upgrade and Updating Network Inventory System

Stofa decided, as part of the DAA transformation to also migrate the HFC network to 1.2 GHz Downstream/ 204 MHz Upstream operation– a necessary move to be able to offer symmetrical services and to prolong the lifetime of the initial R-PHY rollout. On the pre-existing HFC network, Stofa had limited topology data linking cable modems, taps, amplifiers, and nodes. A key goal for the HFC rebuild was to improve topology data for the RF network.

This was to allow better integration with monitoring and PNM solutions in order to allow Stofa to perform targeted proactive network maintenance, and to be able to automatically identify affected customers during outages.

It was decided to do the HFC rebuild in 3 phases on a per service group (SG) basis:

- A. During the first phase a SG is redesigned, resulting in updated and accurate network documentation. No site surveys are done yet (unless known issues are present). Some, but not all, inconsistencies in the network documentation are found and fixed at this point. The updated design data is automatically imported into the PNM solution to enable use of the PNM tools to troubleshoot issues after phase 3.
- B. During the second phase a SG has its passives and amplifiers replaced, still keeping the original US/DS split. Quality Assurance during this phase means that service groups are deployed with identical components, guidelines for mounting, cabling etc. Any inconsistencies between the new network design and the real world are captured during this phase and scheduled for field fix. In addition, cabinets are checked for defects and improvements needed for better airflow. As the new Node/RPD generates more heat than the old optical node the cabinet is given a new lid and door – both with ventilation holes. A small improvement that yields a 10-15C temperature decrease in the Node/RPD. After phase 2 the HFC network is completely rebuilt (but using the old nodes and frequency plans).
- C. The third and final phase involves deploying the R-PHY device/Node and swapping diplex filters and upstream modules in the amplifiers, as well as changing the split to 1DS-1US. A new frequency plan is effectively in operation as soon as the RPD is online. RF filters on each tap are swapped to reflect the new frequency plan (Broadcast TV subscriptions are managed using bandpass filters).

4.4. CIN planning

DAA brings several specific requirements to the CIN, the Converged IP Network. The solution chosen by Stofa (depicted in Figure 10) is entirely IPv6 based and requires all devices in the CIN to be IPv6 capable (don't assume this is the case). Also, the ability to either support or effectively transport IEEE1588 timing information needed for RPD operation is crucial. Being an early adopter of DAA, there was a very limited choice of possible switches and routers that comply to these requirements. This was made even more challenging because of the requirements for highly scalable backhaul links (nx100G per router). It would have been possible to deploy a cheaper platform based on 40G backhaul links, but the lifetime of such a solution would not be in line with the wish to roll out a platform with a very long lifetime in the field and network.

Stofa decided to build a multi-service IP/MPLS network to serve as the CIN. A key driver for this was the wish for full flexibility on the association of RPDs to MAC cores, and removing any geographical constraints. Ideally, the ability to simply move RPDs (through provisioning) to a new core would also mean that transitioning to a virtual core at a later stage would be seamless.

The wish to implement (and provide device support for) 802.11X and MAC SEC should be carefully considered. For 802.11X specifically, there are challenges if RPDs are deployed in Daisy Chain or Ring Topologies; we have noticed 802.11X vendor implementations may be pre-mature. It is recommended that the CIN edge devices used to support the implementation be compatible with those recommended and tested by the RPD vendor.

Another observation: for the near future – we expect the backhaul links to be a major cost driver—especially if deploying 100G/200G. The Cable Labs Coherent Optics specification [P2PCO-SP] could help drive costs down though.

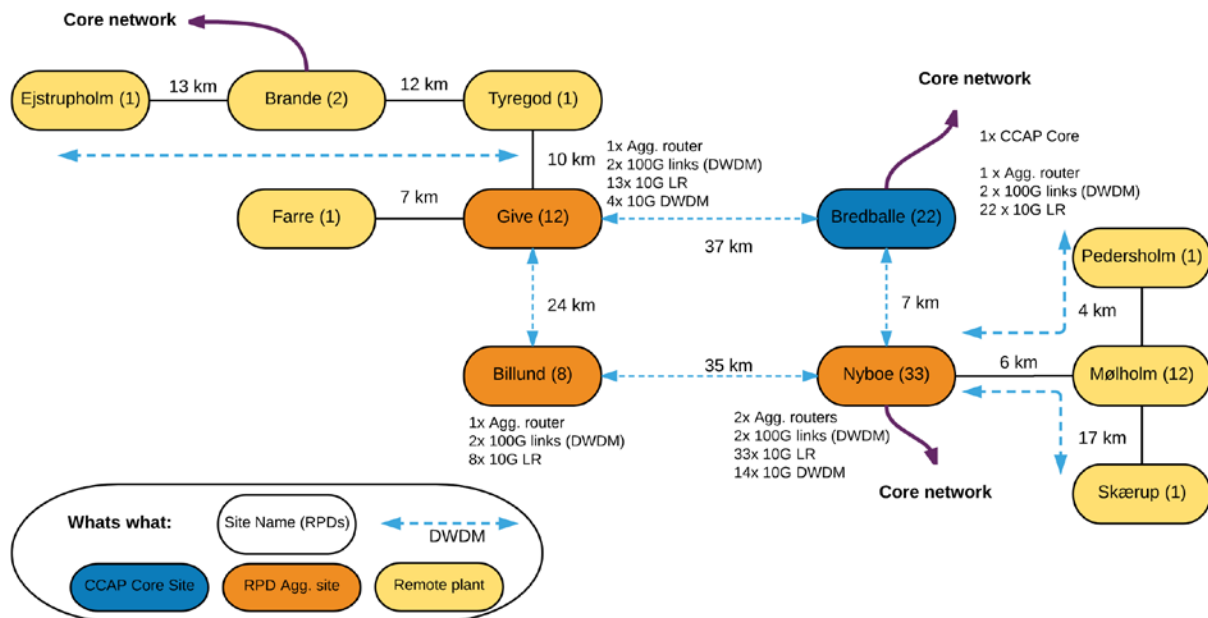


Figure 10 - CIN Network Design

4.5. IEEE1588 Timing

The implementation of Precision Time Protocol (PTP) Grandmasters (GM) and supporting the flow of PTP traffic in the network was originally a major concern for Stofa as there was no in-house knowledge of IEEE1588. Stofa limited its lab tests to GM devices from two vendors and the experience from using both was that the basic PTP functionality needed for DAA is almost trivial. If selected CIN devices can provide consistent low latency forwarding ideally using QoS based on DSCP classification, the operator only needs to consider possible asymmetric routing for the PTP traffic – the reason for this is illustrated in Figure 11:

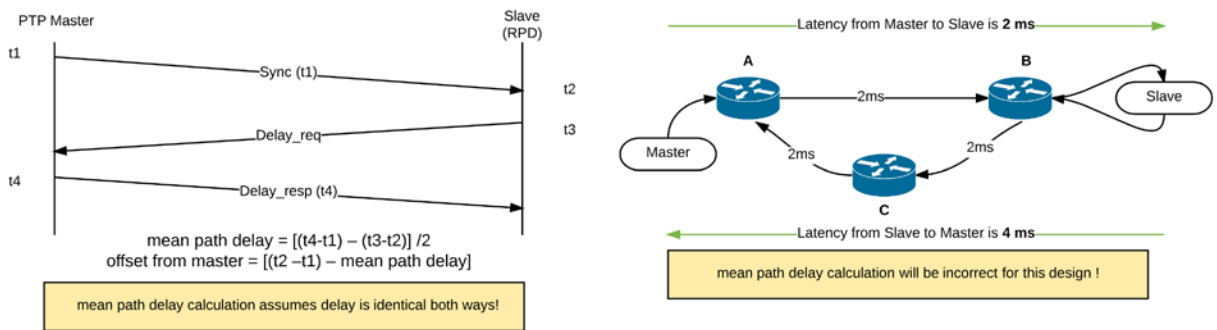


Figure 11 -PTP Latency Calculation

PTP used with Remote-Phy allows for a packet delay variation (PDV) of 2 ms – which means that R-PHY should actually work with up to 2 ms of difference in latency in an asymmetric routing situation. But with more focus being put on Latency in DOCSIS networks Stofa has decided to aim for symmetrical routing for PTP – to be able to implement any latency improvements done to DOCSIS without having to rework the CIN.

Failover of PTP in case of GM failure or network issues is still very loosely defined in the R-PHY specs and recovery implementations vary from vendor to vendor. Currently with the Stofa implementation (depicted in Figure 12), the workaround is to have two active GMs in the network sourcing the PTP traffic from the same IPv6 address but using different prefix lengths. In effect this means that the IP routing protocols makes the primary GM reachable on the network until the link to it fails, in which case the routing protocol switches to the secondary GM. This does not account for cases where only PTP communications with the GM fails – In these cases manual intervention is needed but with the holdover time for the MAC core and RPDs of at least 2 hours, this is tolerable.

The Stofa CIN network is actually license-based upgradeable to PTP Boundary clock support. It remains to be seen if this step will be needed in the future. So far Stofa have not needed PTP test equipment to validate accuracy in the field, because of the low IP hop count in the network, and simplified routes to the RPDs. However it is important to understand that the operator will need test equipment synchronized to the same source as their GMs (typically GNSS) to measure and troubleshoot timing accuracy, for more complicated networks, and for better analysis of the impact on PTP accuracy as it traverses the network.

As more experience is gained Stofa expects to increase both distance and hop count between CCAP cores and RPDs – a handheld PTP reference might prove very useful at that stage.




Figure 12 - Timing Distribution Network

4.6. Product Qualification Phases

4.6.1. Initial Testing and Solution Evaluation

As mentioned, Stofa is an early adopter, and was one of the first operators in the world to qualify and deploy a Remote PHY architecture with live customers. Correspondingly, the first MAC core and nodes that were used in the initial testing in the Stofa lab were of the first few beta units of the product worldwide. Those components had to be integrated with a new DHCPv6 and timing servers, and new routers chosen for the CIN of the R-PHY network. Given all that, Stofa wanted to be best prepared for the testing, and started as early as the equipment vendor could allow them to test.

In order to get a head-start on the new R-PHY technology, Stofa received a very early pre-production unit into their lab many months before it was ready for production.

As a preparation phase, Stofa decided to test the integrated CCAP solution in the lab, although the solution they are planning to deploy is Remote PHY and not an I-CCAP solution. The rationale for that was that the hardware used for the MAC core operation was the same as the I-CCAP hardware, and the CLI commands and general flow of operation are the same between the I-CCAP and the MAC core solutions. Stofa tested the I-CCAP in the lab, and they also launched a limited field trial with the I-CCAP, in order to learn more about the field behavior of the platform and ability to converge DOCSIS and VOD video services. The lab and field trials gave the engineering and operation teams some experience with services configuration, monitoring, and debug of the platform. After this trial Stofa gained valuable experience and confidence on the Remote PHY architecture.

The RPD and node were staged in the Stofa lab and connected in a very basic network scheme to make it operational, as depicted in Figure 13.




Figure 13 - Stofa Initial Lab Setup

The test plan that was run on the set up including the following tests of functionality:

- A. Node registration in the MAC core – GCP tunnel established, timing server connectivity operational
- B. DOCSIS – cable modem registered, high throughput achieved to CM
- C. Voice over DOCSIS
- D. Video on Demand (VoD) services
- E. Broadcast Video
- F. DOCSIS 3.1 DS – OFDM
- G. Multiple R-PHY nodes on single core
- H. Routing – IPv6
- I. Routing – OSPF and OSPFv3 operation
- J. Terminal Access Controller Access-Control System (TACACS)
- K. Lawful intercept using SNMPv3
- L. Load testing, making sure QoS is maintained, US and DS
- M. DOCSIS Load Balancing – making sure load balancing rules are triggers are respected
- N. Integrated Upstream Agility (Dynamic Modulation Profile switching)
- O. Sparing – Routing Switch Module, Downstream card, Upstream card
- P. Link redundancy – to the Remote PHY module
- Q. Timing server failure – connectivity lost to the timing server, testing the holdover operation

Additional tests were run in later phases that included more than one RPD per core, and additional error cases coverage.

4.6.2. Integration into the Network

The new R-PHY product must fit into the existing operator's network and be able to integrate with existing components that are part of the normal network operation and management. The integration with these components had to be tested in advance, and guiding documents generated in order to instruct the operational team during deployment.

The main systems that required integration are:

- **DHCP servers** – Following internal analysis, and also aligning with the vendor recommendation, Stofa decided to use separate DHCP servers for the RPD nodes in the network. It made sense separating the DHCP instance from the one used for subscriber devices and maintain separation of services for the sake of security and ease of operation. The DHCP servers selected had to be IPv6 supporting, to match the core and node routing capabilities. The new DHCP servers were set up in the lab with the R-PHY devices and tested with the rest of the architecture for the first time.
- **TACACS Server** – Stofa manages their users and permission levels using TACACS servers, the core had to be tested along with the existing servers.
- **Proactive Network Management Solution** – in order to improve network visibility and allow fast action to handle potential customer affecting issues, Stofa decided to adopt a new network management system, which made integration testing much easier, given the R-PHY and PNM products came from the same vendor. The PNM solution chosen collects information from the plant, alerts the operator of a proactive alarm notifications which allows them to fix it before it is causing any significant service degradation and or outage to the customer. More details and case studies can be found in [Cunha PNM].

Below in Figure 14 is a visualization of the topological information provided by the tool, allowing for geographical location of potential outage causing issues.

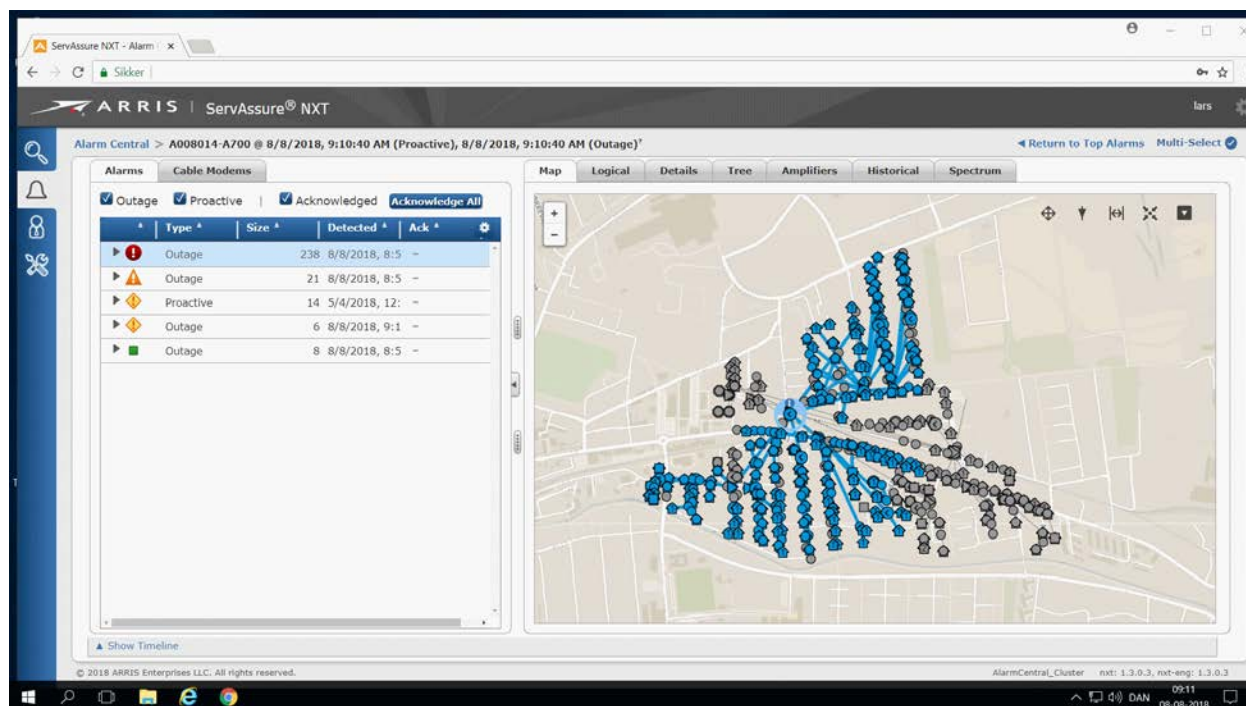


Figure 14 - PNM Topology View

- **US RF network alignment and RF Quality Monitoring Solution** – With the analog links becoming digital links between the node and hub, a replacement solution was required for the legacy analog US RF alignment and RF monitoring systems including some of the field tools. Stofa decided to use a tool that connects remotely to the RPD and provides real time US RF

spectrum measurements through which the US path can be aligned, and which can be used for periodic or reactive RF network monitoring and maintenance. This tool is measuring the upstream RF spectrum:

- Utilizing Fast Fourier Transform (FFT) data from the RF burst receiver, spectrum from a specific DOCSIS upstream is displayed (Power Density vs. Frequency)
- Uses the FFT data to measure the noise level in the upstream channel

Example screens are shown in Figure 15.

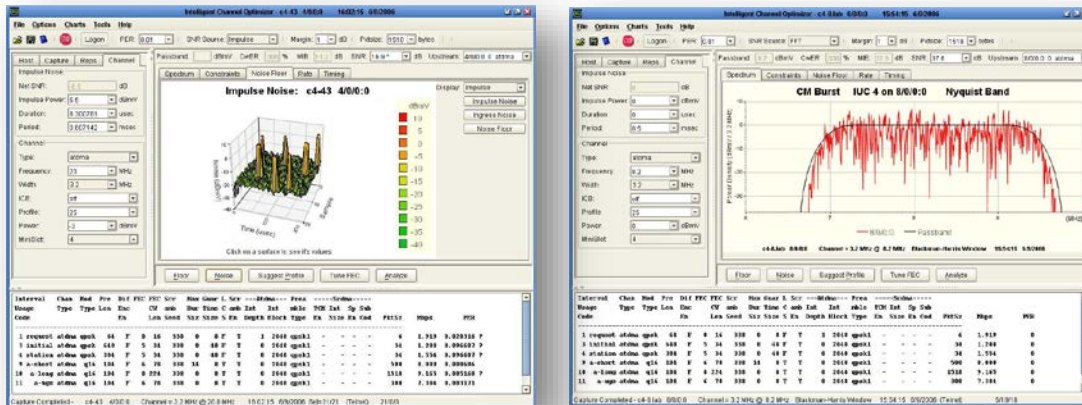


Figure 15 -Software Tool for US RF Network Alignment and Monitoring

- **Implementation of an Optimized R-PHY node RF specification enabling a seamless integration of the R-PHY node with a N+5 RF cascade** – With support from Stofa, the vendor defined the D3.0 and D3.1 RF performance and level design for the R-PHY node with RF cascade to achieve a most optimized EOL (End of Line) MER supporting 4K QAM Downstream and 2K QAM Upstream operation.

4.7. Automation

The introduction of Remote PHY architecture requires a significant upgrade of the network operational management. Distributed Access Architectures drive more functionality closer to the subscriber and into the field. This causes an explosion of managed devices out in the field. In locations where there were only analog components, now with DAA there are a few hundred more digital and managed devices. The Remote PHY devices are IP devices that are part of the IP network, require software upgrades, and may suffer logical and physical outages. All this strongly drives the need for more automation in the network, for onboarding, provisioning, and monitoring those RPDs.

If we summarize the different functions in the back office that need to be managed and integrated to support the operation of the R-PHY devices, it may look like this, in Figure 16:

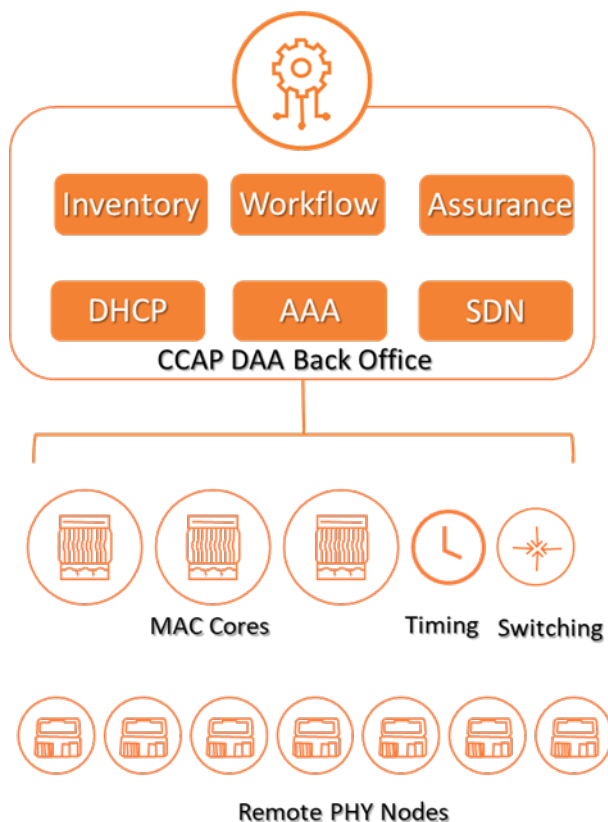


Figure 16 - Back office Management Architecture

Stofa considered different options for managing the network, including building their own automation tools, but decided it will be more beneficial to invest in a long-term management system for the network that can expand to additional functions and support future migration to virtualization.

Specifically, Stofa decided to deploy an RPD manager, which supports automation of the onboarding of a new node. It performs this function from the box into the network, with configuration required from the operator's side. This is very critical to operators like Stofa that are planning to deploy dozens of RPDs a week (or more) as part of a fast roll out plan or are needing to manage a large RPD deployment.

The RPD manager allows the field technician to scan the barcode on the RPD using their mobile device at the point of deployment, and loading that information into the centralized database, to create the network inventory. This information will include the RPD serial number, the node geo-location, time of day, and more parameters relevant for the node installation.

After the node is identified, it is matched to a MAC core (according to the operator's policy), and downloaded software version and initial configuration.




Figure 17 - RPD Manager Onboarding

In order to manage the matching of RPDs to cores, the RPD manager interfaces to the DHCP server for the Remote PHY devices. It is also able to download configuration templates to the MAC cores, to eliminate time-consuming manual CLI commands.

Automating the RPD management allows the operator to do bulk software upgrade or configuration changes, thus reducing the cycles and manual work required to manage the field devices. Its operation diagram is included in Figure 17.

In the future, the same network management system is planned to also support the provisioning of the CIN, allowing for end to end provisioning and monitoring of the Remote PHY network, from one single system.

5. Deployment and Operational Results

5.1. Field Deployment

With the objective to start the first field trial, the lab setup was replicated into a field location with 4 RPDs connected to the MAC core. The first trial phase included launching only one RPD into the live network, supporting DOCSIS, VoD and Broadcast services to live customers. The first node was, naturally, the one that took the most time and effort. The dominant challenges encountered were mostly found in the deployment of the "supporting network", and were mostly related to routing configuration. The routes to DHCP and TACACS servers had to be established and verified. The new CIN network had to be configured to match the future services expansion, which also took some time and effort. Eventually, the cut over was declared a success, with no issues on the network.

The single node was left running for a few weeks, to make sure all services had good stability. After few weeks, 3 more nodes were cut over, with no issues at the time of the migration or after it.

The field trial was considered a huge success with practically no issues encountered for more than a month despite using pre-production software– and eventually the full production deployment was started as the vendor released field mature software.

6. Performance Improvement

6.1. RF performance

RF Signal data was captured before and after the 3 deployment phases described above. In this table, we summarized the difference in average US and DS signal to noise levels.

DS noise was measured and reported by the cable modem, US noise was collected from the R-PHY node, as described in Table 2:

Table 2 - Upstream and Downstream SNR Measurements

	Upstream average dB				Downstream average dB		
	Before	After	Change		Before	After	Change
Min	28,0	31,1	-0,8		37,6	35,3	-3,5
Avg	33,6	36,2	2,5		40,4	40,3	-0,2
Max	37,4	37,7	8,2		42,0	41,8	2,1

The US SNR shows an average 2.5dB improvement, with up to 8.2dB improvement in some areas. This will allow Stofa to go to higher D3.1 modulations in the future, and provide higher US capacity to their subscriber.

We note that mild improvement to downstream signal quality is seen, on the average.



There are few possible causes for this:

- A. A decision was taken to move the DOCSIS spectrum to 722 – 906 MHz, from 218 – 338 MHz. This was done to keep DVB-C video multiplexes clear of possibly impacting LTE ingress noise. In the long run this spectrum will be used for OFDM which is more resilient to LTE interference. The fact that similar SNR was maintained on a much higher frequency range can be considered as an improvement in the network operation, compared to previous state.
- B. It should be noted that no changes to the RF drop cable were made except for F-connector replacements (if deemed necessary following visual screening). In many cases the RF drop cables have been out in the network for over 20 to 30 years whilst exposed to varying environmental conditions (low and high temperatures, rain, wind and sun) and possibly mechanical stress (installation conditions). All are factors impacting the RF attenuation loss and RF shielding quality between the tap and home, especially in the higher downstream frequency ranges (750 to 1.2GHz).
- C. Deployment of new 1.2GHz R-PHY nodes, RF amplifiers and RF passives follow a strict and highly optimized RF network design and implementation process. RF network (R-PHY node to tap RF outlet) design characteristics assume low drop cable attenuation and cable quality conditions. In cases where non-typical drop cable RF losses had to be compensated, Stofa had historically implemented one-off work around solutions (such as higher RF amp output,

installation of a lower tap value, etc...). Some of these one-off compensation fixes may be negated through the RF network upgrade process.

- D. Impact of poor in home RF network quality on End-of-Line downstream performance. Any network changes/RF performance improvements within the home were deemed out of scope for the R-PHY and D3.1 network upgrade project.

The graphs below in Figure 18 show the channel SNR distribution, before and after the network upgrade.







Figure 18 - US and DS channel SNR distribution

In these diagrams we can see in more detail the SNR levels reported for every channel (SC-QAM) that is in use by Cable Modems in the examined part of the network.

Here we can see the clear improvement in US SNR. The DS SNR remains distributed, and is assumed to be as such from the reasons mentioned in the section above.

6.2. Space & Power in the Headend

One of the main drivers for the network upgrade project has been to reduce the space and power needed in the headends for network operation.

Before the network upgrade, the following list of devices were required to provide the full set of services, DOCSIS, VoD and Broadcast, in the headend:

- A. CMTS – supporting all DOCSIS services
- B. Edge QAM (since using an M-CMTS architecture) – supporting DOCSIS PHY and VoD
- C. Broadcast Edge QAM
- D. DTI server

Additional headend equipment required is the RF combining network – including RF combiners and splitters to combine the different RF outputs from the EQAMs for unicast and broadcast traffic. Operator will also need optical transceivers to transmit the analog RF signals over the fiber to the nodes. In some operators' network, there is also an RF Switch attached to the CMTS, it was not used in this case.

In some limited cases, Stofa placed the EQAM devices in the hub, closer to the subscriber. We will focus on the more typical deployment scenario where all equipment is deployed in the same headend.

The devices that will reside in the headend in the newly established R-PHY architecture are:

- A. MAC Core
- B. Timing servers
- C. Aggregation Router(s)

A note about the aggregation routers: the switches or routers that are required to transmit the R-PHY signals from the MAC core to the nodes can reside in the headend, data center, or hub. In many cases, some of the routing devices were already pre-existing in the headend before or are used for other network services. In our analysis, we have chosen to include one additional "dedicated" aggregation router for the sake of comparison to the previous state of network devices in the headend.

For Remote PHY, given the decentralized nature of the architecture, PHY functions are distributed in the field, inside cabinets (in the Stofa case), or field nodes (on a pole or in storage locations). Those devices take space, naturally, but that space is assumed to already "be there" given the fiber node that is enclosing the Remote PHY device was there for the legacy services, to provide analog fiber transmission. In some cases, there may be a need to upgrade the fiber nodes or their power supplies to meet the new R-PHY requirements. These upgrades will incur cost and resources that are not calculated in this section, given they are outside the headend, and may not be applicable to all operators.

Summarizing the space and power consumption for all the devices that combine the two configurations, assuming power consumption at 25⁰ C and calculating the amount of space and power required to make 96 SGs operational, the numbers are given in Table 3:

Table 3 - Space and Power required for M-CMTS and R-PHY Devices in the Headend

Device	Count	Space per Unit [RU]	Estimated Total Power [W]
CMTS – UBR 10K	2	18	4470
EQAM – DOCSIS and VoD	2	13	3000
EQAM – Broadcast	1	1	400
DTI server	1	1	80
E6000 MAC Core	1	16	3125
Timing Server	2	1	40
Aggregation Router	1	1	320

It is important to mention that the R-PHY solution can double and more its capacity without any additional hardware required. We have chosen to run the comparison on 96 SGs, since we feel this is representative enough, and in addition, the lower density on the legacy devices makes the numbers for higher capacity very difficult to support within one hub.

If we combine those numbers for the case of 96 SGs support, we get the totals in Table 4:

Table 4 - Summary of Space and Power Required for Legacy and R-PHY Architectures

	Space [RU]	Power [W]
Legacy Architecture	64	7950
Remote PHY	19	3485
Improvement	Over 3x	over 2x

The difference in space and power consumption between the "before" and "after" states is quite outstanding.

However, we should consider a few points that may narrow the gap a bit, but would still keep the obvious benefit of the transition to R-PHY:

- The legacy architecture Stofa started with was in service for many years, and included older types of devices, that may have not taken advantage of more current technology advancements such as newer chipsets, and routing technologies. It is safe to say that there are newer I-CCAP solutions in the market that will be better than the M-CMTS architecture in terms of space and power. However, even those alternatives come short when compared to Remote PHY, simply because R-PHY moves the PHY processing out of the headend and thus removes space and power required for it.

- Stofa has been using M-CMTS architecture that required a few more devices such as the timing server, and a separate EQAM. If the starting point has been I-CMTS or I-CCAP, it is safe to assume that the initial space and power consumption would have been lower. In addition, in order to support redundancy, external RF switches have been in use. Again, there were alternatives in the market that supported that sparing protection without the need for external devices, and with lower space and power consumption.
- The legacy architecture included external Broadcast QAM processing – there are solutions in the market that already allow for that processing to be done inside the CCAP chassis, thus reducing the need for that EQAM.

As mentioned, even considering all the points that may narrow the gap between the legacy architecture and the Remote PHY architecture, the new design provides significant space and power savings in the headend.

The benefit becomes even more evident when we consider that the existing devices used in the headend for the MAC core processing can double or more their capacity, supporting few times more the number of service groups estimated in this analysis, without any additional hardware component needed.

In the future, the MAC core function can be virtualized, meaning transferred to reside on a virtual machine, using an off the shelf server. Depending on the performance of future off the shelf servers versus future appliance-based MAC Cores, it is possible that this may provide even greater space and power savings, as well as increase the network resource assignment flexibility.

6.3. Operational Simplification

The new network architecture introduces better operational simplification, compared to the previous network Stofa was managing.

If we look at the number of headend devices that need to be managed to support 200 service groups, we can see the significant difference in Table 5:

Table 5 - Number of Managed Devices in Different Architectures

Before Network Upgrade	After R-PHY Introduction
4 CMTS	1 MAC Core
3 Universal Edge QAM	2 timing server
1 DTI servers	1 Aggregation Router*
8 devices	4 devices

Note - the number of additional aggregation routers depends on the operator's architecture

Less devices that need to be managed means operational simplicity: fewer physical devices mean less network cables and interfaces to manage them. It also means fewer management tools (software UI or command line) need to be used and trained for (which will also lower the support cost). It also means less

interoperability issues between the devices, and quicker debugging process (since there are less devices to check for the source of issue). It also means less IP addresses need to be assigned and less inventory the operator will have to keep in order to deal with outages.

In the remote PHY architecture, all the nodes are managed and provisioned through the MAC Core, per the R-PHY Spec [R-PHY Spec] (there is no direct configuration interface into the remote PHY devices). So although R-PHY introduces many new devices into the field, they are all provisioned and managed from one central point, which makes it easier from an operational perspective, and also for problem diagnosis.

According to technicians' feedback from Stofa, understanding and operating the new architecture was significantly easier compared to the previous one, due to the centralized management. It required less training, and shortened the debug cycles for issues in the lab and in the field, compared to the previous network architecture.

The additional level of automation provided by the RPD manager tool helped a lot to facilitate the quick and painless installation process. The RPD manager allowed to shorten the new node deployment period from roughly 15 minutes on the legacy platform to just 5 minutes on the DAA platform (including scripted CIN, DNS and IPAM provisioning currently done outside of RPD manager) – While this might not sound like a huge improvement - it is a critical optimization when deploying dozens of RPD's each day.

Conclusion

Key Benefits

Stofa gained many benefits out of the network upgrade project and the transition to R-PHY.

These are the ones that were most important to Stofa, covered in this paper:

- **Better Signal to Noise Ratio (SNR)** – as measured, the transition to Remote PHY provided significant increase in US SNR (average of 2.5dB). The improvement of SNR on the DS direction is not manifested in the measurements. We do believe there is some slight improvement that is countered by the transition to the higher frequency range. In addition, there is a potential DS SNR improvement that is limited by other elements in the existing network that have not been upgraded such as the drop cable and in-home end-of-line network.
- **Path to efficient use of DOCSIS 3.1**– the network upgrade project included all devices upgrade to 1.2 GHz spectrum support in the downstream. This enlarges the available downstream spectrum for transmission, using DOCSIS 3.1 above the 1 GHz previous top bar. It also allows upstream high split, of 204 MHz. In addition, the SNR increase in both US and DS directions allow for higher modulation orders to be used for signal transmission, hence allowing more bits per Hertz on the existing infrastructure. These changes significantly increase the plant capacity, and allow Stofa to increase their bandwidth service tiers offered to their customers.
- **Space and power saving in the headend** – the space and power needed to support the R-PHY architecture in the headend was shown to be significantly lower than the previously used architecture – 3x saving in the space required, and more than 2x saving in power, for the same amount of service groups. The saving is very impactful in terms of operational cost, and will allow Stofa to consolidate headends and hubs.

- **Future growth made easier** – the Remote PHY solution deployed, as well as the supporting CIN network, were chosen out of consideration of the future Stofa needs. The design was created focusing on future growth, such that additional capacity and capability added to the network will not require field technicians and truck rolls. Future upgrades to the network will be done by remote software configuration, licenses enablement, and capacity increases in software. Stofa fully focused on guaranteeing that, as part of their network planning process, and have invested in resources and tools today, for the network needs of tomorrow.
- **Path to virtualization** – In the selected distributed access architecture, Remote PHY, the MAC core functionality remains in the centralized headend, allowing for better control and access to it. This design will enable migrating that functionality in the future to virtualized servers, performing the same functions, using a virtualized MAC core on off-the-shelf servers. These virtualized MAC cores, or vCores, will be designed according to Network Function Virtualization (NFV) standards, and managed according to Software Defined Networks (SDN) principles, paving the path to a fully virtualized headend design.
- **Improved operational simplicity** – as shown, the network upgrade provided a greater operational simplicity, for engineering teams, support technicians, and customer services. The simplicity is gained along with the decrease on the number of devices to be managed in the network, the centralized management of all the remote PHY devices in the field, and the additional level of automation obtained using advanced PNM and orchestration tools.
- **R-PHY project drove an upgrade to the “DOCSIS supporting network”, more investment made in CIN network upgrade** – Along with the access network upgrade, Stofa decided to invest in upgrading their CIN network to be a L3 supporting network. In this design, given everything is routed, Stofa can load share traffic dynamically across all the expensive backhaul links, and not have links dedicated to one service or another. Secondly, with RPDs connected directly to the IP network (using low-cost 10G LR transceivers), they are not locked to a specific MAC Core site, which gives network configuration flexibility as well.

Key Challenges

Such a complex migration project cannot come without challenges, along the way. Some stemmed from the still young technology being used, and some from the significant change this upgrade has made to the network and the operations:

- **Remote PHY Spec stability** – One of the significant challenges in the R-PHY project was derived from Stofa being a very early adopter of the new architecture. During the product development phases, the R-PHY specifications were still changing, which caused delay in product delivery, and some recurring changes in the network design for deployment. Stofa's strategy has been to keep everything standards based so quite a lot of time has been spent discussing how grey areas of the specifications could/should be implemented. Close cooperation with the equipment vendor is and has been key for Stofa to be able to start the rollout at such an early stage.
- **Network planning** – In order to create the best design for the Stofa network, the operator had to investigate new areas for them, such as the CIN routing upgrade, and the timing distribution architecture. Stofa upgraded their routers to support future growth of the services, and thus purchased and integrated new spine and leaf routers, creating a resilient new backbone network. This network also had to support the timing distribution requirements needed for Remote PHY operation, guaranteeing low delay and symmetric forward and return paths for the timing

synchronization traffic. This was a new domain for Stofa, in which they needed to ramp up fast, which posed some challenge for the network architects during the project.

- **Operational processes update** – as with every major network upgrade project, the transition to a new architecture required massive updates to the operational processes defined for the Stofa engineers and technicians. Deep level trainings and educational sessions were held to bring the team up to speed on the new architecture and create new processes that will optimize the installation process and network management for all types of services.
- **Change in field measurement tools** – the change of architecture was also accompanied by the need to change some of the measurement and debug tools that were previously used with the M-CMTS network design. Specifically, the RF alignment procedures used in the past were difficult to use with R-PHY, as they required proprietary support in the nodes and headend devices. To resolve that, Stofa have transitioned to standard based methods of data collection from the plant, using DOCSIS standard features (like FFT MIBs) to obtain the SNR and noise values from the plant.

Major Takeaways

For operators considering an upgrade to R-PHY, or getting ready for the deployment, we believe the Stofa story as told in this paper will provide a lot of value, by preparing the operators for the main decision points and activities you will undertake as part of the project. Here are the main points for consideration that we have drawn from our experience:

1. **Start testing early** – getting yourself familiarize with the product and architecture is going to pay off many times over later. Start testing as early as you can, even with similar flavors of the architecture. Getting your lab ready for the new architecture testing is a non-negligible effort, and you will save yourself time if you start that early, even before the product is available.
2. **Planning is key** – one cannot underestimate the importance of pre-planning in a complex project like an architecture upgrade. Planning the lab and field qualification in advance will allow the technicians and network engineers time to figure out their needs to support the effort. In addition, planning the deployment phases in advance will uncover migration challenges, and operational issues that can get resolved in advance, as to not slow down the roll out.
3. **CIN network and Timing architecture are of great importance** – Deploying DAA is massively different than I-CMTS or M-CMTS, and their supporting networks are more complex, specifically the CIN and Timing architectures. Operators must dedicate resources to plan these aspects of the network, with the right experience and skill sets. Specifically, the CIN network has to be designed, as the operator must make sure they understand the requirements specific to DAA - IEEE1588(PTP), non-symmetric routing etc. make sure you dedicate the right resources to plan for that
4. **Gradual migration** - Do not underestimate the importance of starting out slowly – it is important to have enough time to capture both process glitches and technical issues and to implement the required improvements. The rollout will be unlike anything your HFC techs have tried before. Consider all tools and measurements that will be impacted by the architecture change and plan ahead. It is recommended to start field deployment with one node, and gradually increase scale for full deployment.
5. **Plant Data Quality** – Poor HFC network documentation will seriously slow down the deployment rate and results in miscommunications to customers about planned outages that may get delayed due to operational reasons. It is worth considering a site survey, in cases where there is doubt.

6. **Importance of Automation** – the significant change in the network is an opportunity to automate more processes, making it easier for the deployment team and longer term also easier for the maintenance and support teams. Specifically, automating the RPDs management is beneficial, to allow for unified interface to the devices, bulk configuration and management, inventory tracking, and debuggability of the units in the field.
7. **Partnership** – Along the entire project, the Stofa and equipment vendor's teams have kept very open communication line, and full transparency on the status of their development and readiness for the project. The close communication enabled us to quickly react or hurdles in the plan, and form work arounds together in a way that will not delay the project and push it forward.
We recommend weekly meetings between the different integration teams, and additional ad-hoc meetings (quickly scheduled) to discuss specific issues and challenges.

In conclusion, the authors consider the Remote PHY deployment and network upgrade project to be highly successful. It provided the benefits Stofa was looking to get from transitioning to the new technology, and guarantees the Stofa network can grow and expand in the future, providing excellent service to all its subscribers.

Abbreviations

ADC	Analog to Digital Converter
CCAP	Converged Cable Access Platform
CIN	Converged Interconnect Network
CM	Cable Modem
CMTS	Cable Modem Termination System
D3.0	Data Over Cable Service Interface Specification version 3.0
D3.1	Data Over Cable Service Interface Specification version 3.1
DAA	Decentralized Access Architectures
DAC	Digital to Analog Converter
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DWDM	Dense wavelength division multiplexing
FEC	Forward Error Correction
FFT	Fast Fourier Transform
Gbps	Gigabits Per Second
GHz	Gigahertz
GM	Grand Master
HFC	Hybrid Fiber Coax
HSD	High Speed Data
MAC	Media Access Control interface
MACPHY	DAA instantiation that places both MAC & PHY in the Node
MPEG	Moving Picture Experts Group
MIB	Management Information Base
MPLS	Multiprotocol Label Switching
MSO	Multiple System Operator
OFDM	Orthogonal Frequency Division Multiplexing
NFV	Network Function Virtualization
PHY	Physical interface

PNM	Proactive Network Maintenance
PTP	Precision Time Protocol
QAM	Quadrature Amplitude Modulation
R-MACPHY	Remote MAC-PHY
RPD	Remote PHY Device
R-PHY	Remote PHY
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SDN	Software Defined Network
SG	Service Group
SNR	Signal to Noise Ratio
TACACS	Terminal Access Controller Access-Control System
VLAN	Virtual Local Area Network
vCore	Virtual Core
VoD	Video on Demand

Bibliography & References

[Cloonan] T. J. Cloonan et. al., “Looming Challenges and Potential Solutions for Future Distributed CCAP Architectures,” in Proc. 2015 SCTE Cable Tec Expo)

[Silbey] Mary Silbey, Light Reading, "Cable DAA Debuts Worldwide"
<https://www.lightreading.com/cable/ccap-next-gen-nets/cable-daa-debuts-worldwide-/d/d-id/741925>

[R-PHY spec] - Data-Over-Cable Service Interface Specifications - DCA - MHA v2, Remote PHY Specification, CM-SP-R-PHY-I10-180509

[R-DTI Spec] - Data-Over-Cable Service Interface Specifications DCA - MHA v2, Remote DOCSIS Timing Interface, CM-SP-R-DTI-I05-170524

[HFC-Green-ULM] J. Ulm and Z. Maricevic, "Giving HFC a Green Thumb, A Case Study on Access Network and Headend Energy & Space Considerations for Today & Future Architectures", SCTE Cable Tech Expo 2016

[Cunha PNM] G. Cunha, "We Have Arrived. Our Light Bulbs Finally Have IP Addresses! Approaches for Proactively Managing Customer Experience and Reducing OPEX in a Cable Operations Environment, SCTE Cable Tech Expo 2017

[P2PCO-SP] – "P2P Coherent Optics Architecture Specification", P2PCO-SP-ARCH-I01-180629

Learnings From a DAA/DOCSIS 3.1 Early Adopter

Launching and Maintaining a Next-Gen HFC Plant

An Operational Practice prepared for SCTE•ISBE by

Jim Walsh

Solutions Marketing Manager
VIAVI Solutions

David Hering

Senior Product Line Manager
VIAVI Solutions

David Judge

Cable Broadband Engineer
Vodafone New Zealand

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Distributed Access Architecture Drivers.....	3
The Solution: Distributed Access Architectures	3
1. Distributed Access Architecture Variants	4
2. Other DAA Benefits	5
Expected Challenges for Operators	5
1. Removal of RF from hubs	6
2. Ethernet Timing Concerns Due To Separation of MAC and PHY	7
3. Expanded Use of Previously-Specialized Technologies.....	9
4. Increased Complexity Amid Architecture/Vendor Proliferation	9
One Early-Adopter Company's DAA Voyage & How They Overcame Challenges	10
1. Architecture Details	10
2. The Deployment Process:.....	11
3. The Results	11
4. Lessons Learned.....	12
General Conclusions.....	12
1. Headend/Hub Construction.....	13
2. Fiber Construction	13
3. R-PHY Installation and Cutover	13
4. Maintenance.....	13
Conclusion.....	13
Abbreviations	14

List of Figures

Title	Page Number
Figure 1 - Distributed Access Architecture Variants	4
Figure 2 - Benefits of Deeper Ethernet	5
Figure 3 - Challenges Created By Removal of RF From Hubs.....	6
Figure 4 - Virtualized Upstream Ingress and Sweep System	7
Figure 5 - PTP Use in Remote PHY	8
Figure 6 - Example of How PTP Error Can Manifest	8
Figure 7 - Complexity Introduced By DAA Variant/Vendor Proliferation.....	9
Figure 8 - Interim DAA Maintenance Architecture	10
Figure 9 - Framework for DAA Deployment and Maintenance	12

Introduction

The Hybrid Fiber Coaxial (HFC) plant is rapidly evolving to economically address subscriber bandwidth and service quality demands, and technologies like DOCSIS 3.1 and distributed access architectures (DAA) are key to this evolution. While these breakthrough technologies will enable HFC to remain the dominant broadband service medium for years to come, they also present challenges to the folks who must maintain the plant. This paper will touch on the drivers behind adoption of these technologies and follow the early planning, preparation, and live-plant rollout of DAA and DOCSIS 3.1 from an early-adopter MSO and test vendor including early successes and lessons learned. Examples will be provided demonstrating how industry-wide collaboration and standardization have been and will remain keys to success.

Distributed Access Architecture Drivers

As an industry, we have grown accustomed to our subscribers demanding more bandwidth each year, usually at the same price as last year's service. Each time we think that bandwidth demand models must be wrong, and that there is no way subscribers can possibly consume data at levels predicted for just a couple of years out, a new killer app emerges to prove us wrong. Over the top video (OTT) is a prime example. Even just a few years ago not many could have imagined our customers' binge-watching habits or that three kids could be simultaneously talking with their friends via FaceTime. In the past cable operators have used the same playbook to address the need for increased bandwidth:

- Add more carriers
- Get more bits/Hz of spectrum
- If all else fails – node splits

Unfortunately, the traditional tools in MSO toolboxes alone are not enough anymore

- **Add More Carriers** – Analog reclamation has freed up space, and optional 1.2GHz/204MHz DOCSIS 3.1 extensions help but are not always practical/affordable to implement
- **Get more bits/Hz of spectrum** – DOCSIS 3.1 helps here with the potential to move to modulations as high as 4096QAM, but moving above 1024QAM often requires SNR improvements that are challenging with current system limitations
- **Node Splits** – Splitting nodes enables smaller service groups, but increases the total number of nodes. Each added node brings with it more gear in the hubs that require CAPEX, rack space, and power/cooling. Rack space and power/cooling budgets alone take this option off the table in many cases.

The Solution: Distributed Access Architectures

Fortunately, the best and brightest minds in cable predicted this crunch and developed/specified a solution to enable the required node split rate while also enabling more bits to be squeezed from each Hz of plant. Hub constraints can be eased by distributing network functions out into the field – functions that previously resided in equipment that occupied space and consumed power in the hubs. The fidgety analog optical link can be converted to a robust, commodity 10G Ethernet optical link with higher SNR's, enabling higher modulation orders (and therefore more bits/Hz).

1. Distributed Access Architecture Variants

Presently the two dominant DAA variants are Remote PHY (R-PHY) and Remote CCAP (R-CCAP), although R-MACPHY solutions under development are gaining mindshare among operators. Within R-PHY there are several sub-variants. We will not cover the technical differences or pros/cons of the variants in detail, as there have been many great papers on this topic in SCTE Cable Tech Expo proceedings from previous years. We will cover just enough to enable an understanding of their impacts on test and measurement.

Distributed Access Architecture Variants

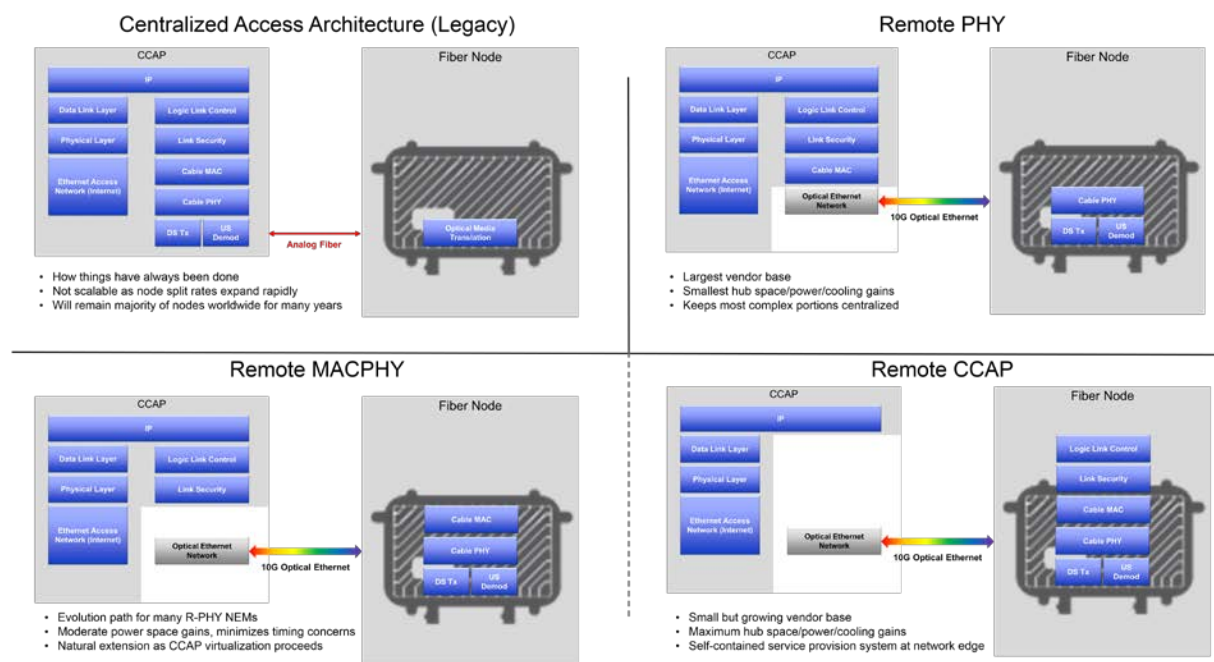


Figure 1 - Distributed Access Architecture Variants

As its name implies, R-PHY moves just the PHY layer out into the node while the MAC layer remains centralized. This MAC/PHY separation creates challenges related to Ethernet timing. This will be discussed in detail in later sections. R-PHY has been the most popular option from established CCAP vendors although they have also been following up with R-MACPHY offerings as well. R-PHY tends to deliver the smallest rack space/power/cooling savings but is viewed by some as the simplest from an overall implementation and management standpoint as the MAC layer management remains centralized.

At the opposite end of the spectrum is R-CCAP, moving nearly all CCAP functionality out into the field. This option took off, as one vendor was very early to market with a solution enabling simultaneous transition to DAA and DOCSIS 3.1. Maximum rack space/power/cooling benefits are generally obtained with this option but concerns about single-sourcing and decentralized management have limited adoption among major operators.

Commonalities for all DAA variants:

- Redistribution of hub/headend-based capabilities into the outside plant

- Replacement of analog optical link with commodity Ethernet optical link
- Removal of RF combining network and related test points from hub/headend

The first point above is perhaps key to most of the challenges that operators will face with DAA deployments. By splitting up previously co-located layers, a brand new optical media translation interface has been created. This is disruptive to traditional demarcation points for existing MSO groups and processes. It is not clear nor obvious where the split between headend or field groups lies with this messy separation. This creates an environment rife with opportunities for responsibilities to fall through the cracks and finger-pointing during troubleshooting. Figuring out and clearly defining ownership of all aspects of this new interface is critical to long-term DAA success.

2. Other DAA Benefits

The obvious benefit of DAA is the enablement of rapidly-accelerating node split rates without expanding current hub/headend real estate and power-cooling footprints, but there are several other compelling benefits.

- Higher SNR digital optical link enables higher modulation order attainment in DOCSIS 3.1 (More bits/Hz!)
- More robust optical link – higher plant reliability and reduced maintenance costs due to replacement of fidgety analog optical link
- Increased service offering flexibility is enabled as Ethernet pushes deeper into the plant. Serve high-usage customers or FTTH clusters via EPON, provide Ethernet for business services, or 5G backhaul from the mux deployed deep into the network instead of running dedicated fiber for each.




Figure 2 - Benefits of Deeper Ethernet

Expected Challenges for Operators

While the benefits of DAA are significant, they come with challenges. In this section we will focus on the challenges faced by those who must turn-up, monitor, and maintain the HFC plant. The challenges generally fall into several general categories:

- Visibility & injection limitations via removal of RF from hubs
- Ethernet timing due to MAC-PHY separation
- Expanded use of previously-specialized technologies
- Increased field complexity amid architectures/vendor proliferation

1. Removal of RF from hubs

Upstream ingress has been the bane of cable operators since the earliest days of 2-way HFC communications, and there are no indications that this will change even in fiber deep architectures. Historically 75%+ of upstream ingress enters the plant through sources originating as home or drop issues. This is not expected to change as plants transition to fiber deep. Customers will continue messing with their home wiring and drops will continue to degrade. Hub-based return path monitoring systems have been a critical tool used in conjunction with field meters to efficiently address upstream noise, but these hub-based systems rely on a return RF feed to operate. All DAA variants have the common trait of removing the RF combining network from hubs disallowing the use of traditional rack-mounted gear and forcing a change in methodologies. Return sweep capabilities face a similar predicament, but with the added challenge related to the inability to combine telemetry signaling into downstream RF in the hub. On the downstream side, forward sweep and leakage taggers also rely on RF combining access which will no longer be present in hubs in a DAA environment.

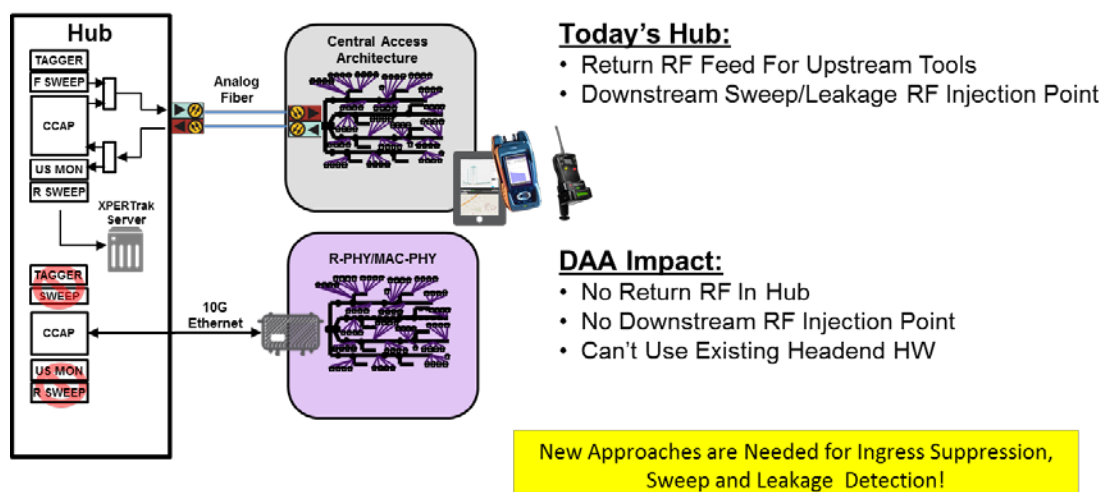


Figure 3 - Challenges Created By Removal of RF From Hubs

One solution to the upstream ingress and return sweep visibility challenge is to virtualize the upstream spectrum analysis and return sweep receiver functionalities into the Remote PHY Device (RPD). This virtualization, combined with a software agent (See RCI in figure below) orchestrating communications between the field meter, central server, CCAP, and RPD enables technicians to use the same process and field meter for sweeping the upstream of both legacy centralized access architecture (CAA) nodes and DAA nodes. The RCI would typically run as an instance on a virtual machine co-located with CCAP or could run on standalone commercial off the shelf (COTS) x86 server.

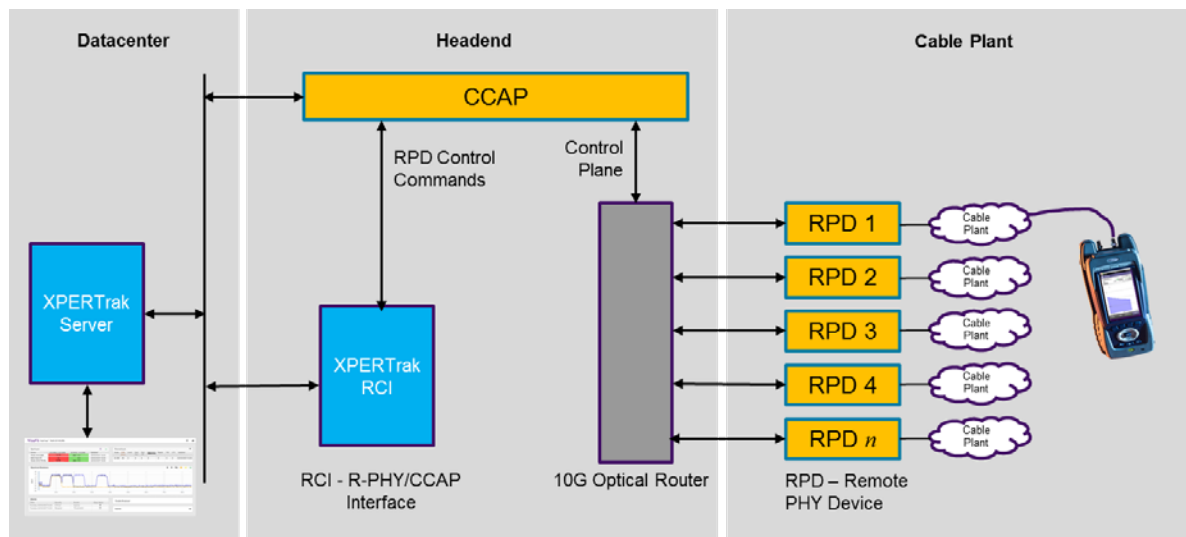


Figure 4 - Virtualized Upstream Ingress and Sweep System

Sweepless sweep has generally been accepted as a substitute for meter-based sweep for the downstream, and leakage tagger capabilities are currently implemented in several DAA vendors systems, with more in development. Through these virtualized solutions, critical maintenance capabilities can be retained throughout the DAA transition while minimizing the complexity faced by technicians in the field.

2. Ethernet Timing Concerns Due To Separation of MAC and PHY

DOCSIS upstream communications (DOCSIS 3.0 and 3.1) at their very core are reliant upon the maintenance of precision timing within MAC domain groups. Modems are assigned precise time slots in which to transmit packets across a shared upstream path, and if a modem transmits out of their assigned slot the potential exists for collisions with packets from other modems resulting in BER or dropped packets. To achieve this synchronization, time information is communicated to the CCAP and RPD's by the Precision Timing Protocol (PTP) Master, and from there each RPD communicates it to CPE via DOCSIS Timing Protocol (DTP). The illustration below shows the typical PTP process.

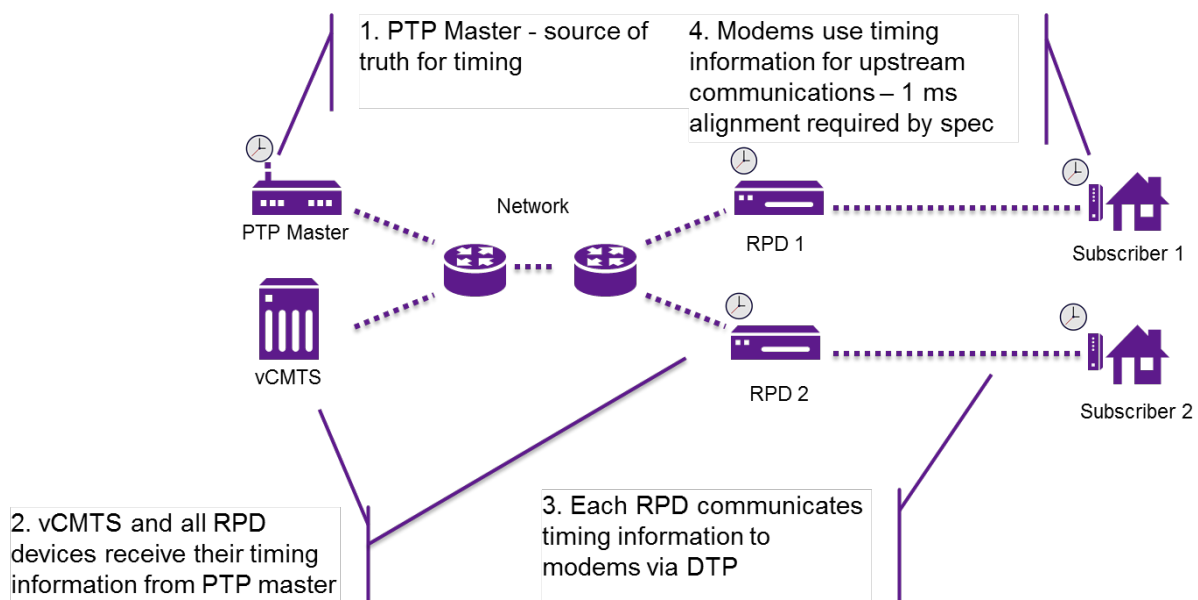


Figure 5 - PTP Use in Remote PHY

The inherent problem with PTP when the MAC and PHY are split is the assumption of symmetry in delay from RPD and PTP Master. The PTP Master communicates total bi-directional delay time for the round trip from itself to each RPD and back. The RPD's assume that this delay is symmetrical and simply divides the delay by two to calculate the delay in each direction. If something impacts only one leg of this journey creating an asymmetrical delay, an error in the delay reported vs actual for the delayed leg of one half of the asymmetry will occur. See below for a numerical example of how delayed PTP messages could result in improper clock synchronization.

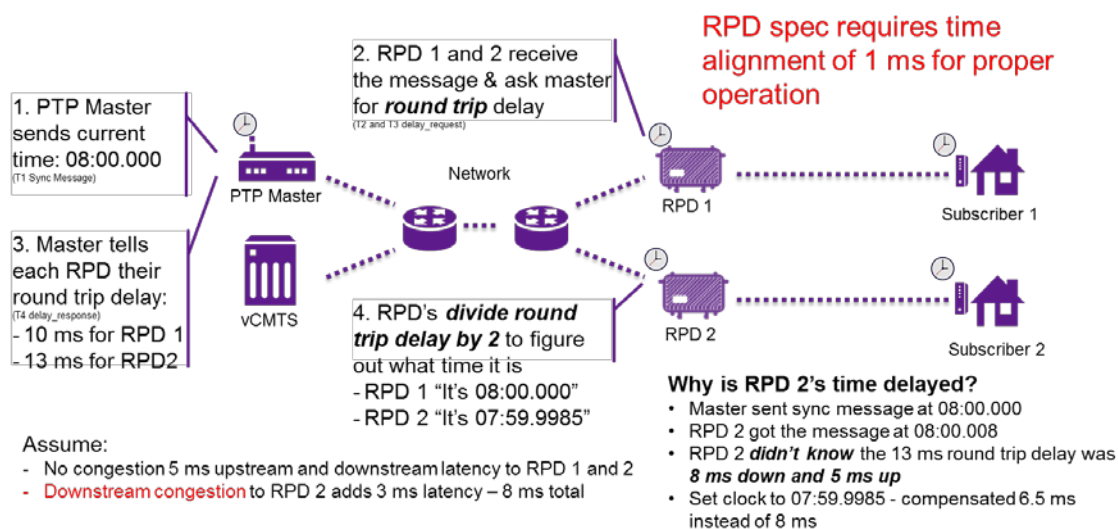


Figure 6 - Example of How PTP Error Can Manifest

3. Expanded Use of Previously-Specialized Technologies

While fiber and Ethernet technologies and testing are not new to the HFC, historically they have been delegated to a specially trained and equipped subset of technicians. The rollout of DAA will push both technologies deeper into the network and force operators to train and equip a much larger portion of their workforce to properly install, maintain, and troubleshoot issues with these technologies. The criticality of clean fiber connections cannot be understated with up to 80% of issues in fiber networks resulting from dirty or damaged connections. Technicians will need training on proper fiber handling and inspection/cleaning and access to the tools to do this efficiently. Dense Wavelength Division Multiplexing (DWDM) has found limited applications in the HFC to date but will play a much larger role in future DAA-enabled networks. The same goes for Ethernet testing and troubleshooting. This can no longer be the exclusive realm of business service techs. Many more maintenance techs will now need to test Ethernet services as part of RPD turn-up and troubleshooting.

4. Increased Complexity Amid Architecture/Vendor Proliferation

The transition to DAA will be a multi-year process for most operators, not a one-time event. It is expected that the majority of most operators' nodes will still be CAA for at least the next 5-10 years or more. It is also expected that large operators (at least) will implement multiple DAA architectures as specific situations dictate, resulting in a multi-year period where any given region may consist of a mix of legacy CAA nodes and DAA nodes using various architectures from multiple vendors.

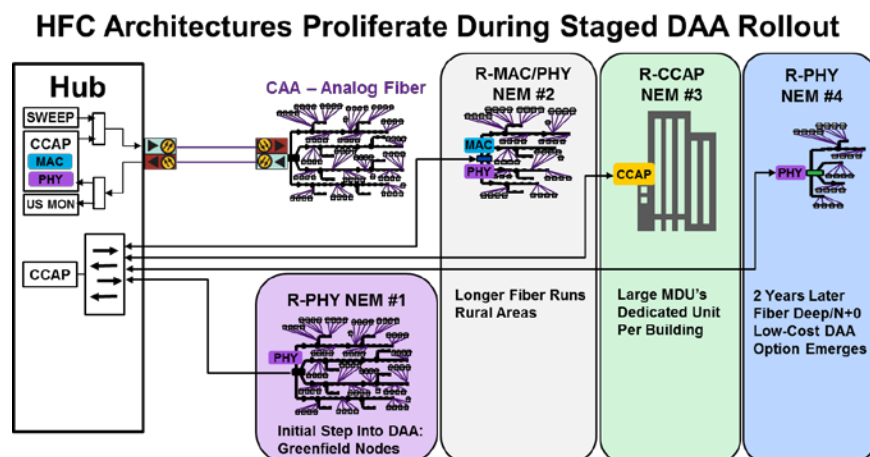


Figure 7 - Complexity Introduced By DAA Variant/Vendor Proliferation

Unless standard maintenance processes that apply across the heterogeneous mix of nodes can be implemented, life for technicians will become quite complex. Imagine the confusion if a tech must use one set of processes and tools for a CAA node they are working on in the morning, a different set for the R-PHY node assigned after lunch, and yet another for the R-CCAP node serving the MDU they are assigned to troubleshoot in the afternoon. Not exactly a recipe for success or operational efficiency! Care must be taken to ensure that maintenance needs are considered during network design and tool selection to ensure that saving pennies on the front end doesn't result in extra dollars of cost to monitor and maintain them.

One Early-Adopter Company's DAA Voyage & How They Overcame Challenges

In this section we will chronicle the voyage undertaken by Vodafone New Zealand (NZ), an early adopter for both DOCSIS 3.1 and DAA technologies. Vodafone is a leading broadband provider in New Zealand and has a sizeable video and broadband customer base in their HFC footprint. Vodafone NZ is historically an early adopter and has not shied away from being on the bleeding edge of new technologies to stay ahead of the competition and provide best-in-class services to their subscribers. Like many cable operators globally, the entry of FTTH competitors was a major driver behind initiatives enabling faster broadband offerings and improved service quality.

1. Architecture Details

Both Christchurch and Wellington employed traditional HFC architecture. The fiber infrastructure was in good shape and could therefore be reused. The new architecture incorporated WDM for the forward path, combining the GPoN traffic, digital TV and telemetry onto a single fiber. The other fiber is used purely to transport return RF back to the hub for use by existing return path monitoring and sweep gear. This solution required no new fiber therefore existing test and measurement gear could be reused.

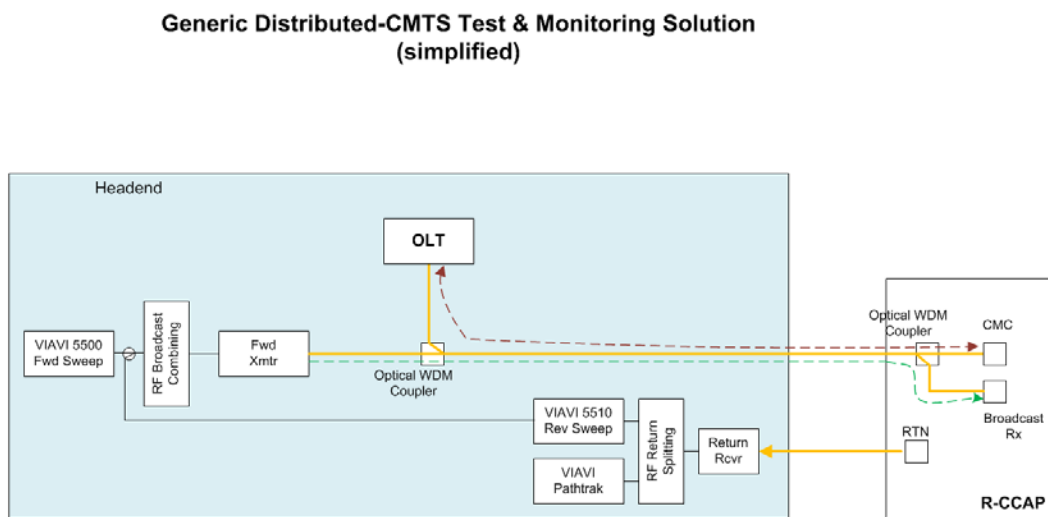


Figure 8 - Interim DAA Maintenance Architecture

Nodes with lower subscriber numbers use a single DAA CMTS to provide desired capacity. For larger nodes, multiple DAA CMTS's are co-located at the optical node. VFNZ has up to two DAA CMTS co-located in an optical node. The forward fiber is split across the two DAA CMTS, resulting in a 3-fiber solution (one shared forward, one each reverse). This solution required both existing fibers to operate, so only one new fiber was required to provide back-feed of RF to the hub for monitoring and sweep use. This was recognized as an interim solution until the D-CMTS vendor implemented the CableLabs specified Narrowband Digital Forward and Return (NDF/NDR) capabilities to support a virtualized test solution.

2. The Deployment Process:

The deployment process itself occurred very quickly and went better than most could have reasonably expected. The entire footprint was cut over to R-CCAP and DOCSIS 3.1 in 6 months, and customer disruption was minimal with outages generally limited to ~60 minutes for single D-CMTS nodes and only during maintenance windows. Sounds easy – right? The secret to the success achieved by Vodafone NZ was careful planning.

Discussions with NEMS and test partners began as far back as 2015 to plot the course for a successful deployment. Attention to detail in system architecture/design assured that maximum capacity could be achieved at minimum cost while still retaining test capabilities critical to maintaining the plant once turned up. After architecture details were nailed down focus shifted to deployment planning.

Configuration and test of all gear before sending to the field for deployment paid great dividends. Centralizing the configuration function helped ensure consistency of configurations across the footprint, enabled use of lower-skilled personnel for field deployment, and minimized opportunities for error.

All units were tested before being sent to the field for deployment. These tests included basic unit functionality, as well as service testing, using racks of CPE in the lab. Pre-deployment testing minimized time spent troubleshooting issues in the field, enabling further standardization of the deployment process. It is much quicker to diagnose defective gear in a controlled environment with skilled personnel available vs in the rain at 2AM with minimal support available. Swapping-in a properly configured replacement unit is also much easier as part of a controlled work procedure vs configuring a replacement in the field.

The field deployment process was also carefully planned and refined over time before large-scale rollout occurred. Contractors performed much of the field work although some Vodafone NZ personnel were available to support as-needed. Typically, two people worked in parallel:

- One tech would splice in new fiber tails and wire in the WDM
- The second tech would swap out the OTN

Once both techs were finished, they would use field meters to perform basic checks (Level/MER), and multiple homes off each node would be polled every 5 min by Vodafone NZ systems to ensure satisfactory performance.

3. The Results

The result of this careful planning and disciplined execution is that everyone won!

Vodafone NZ Subscribers

- Gained access to gigabit service options
- Receive even better quality of service

Vodafone NZ

- Achieved a 300% gain in plant capacity reusing much of their infrastructure
- Implemented future-proofed architecture enabling scalability for future growth
- Retained key maintenance capabilities
- Gained technical knowledge/expertise by being early adopter

4. Lessons Learned

Many things went right in the Vodafone NZ R-CCAP/DOCSIS 3.1 rollout, but there were still some lessons learned along the way. The risks of being an early adopter were known up front, so it was no surprise when these risks materialized into challenges. Being one of the first large-scale DOCSIS 3.1 deployments globally, we surfaced issues related to the immaturity of support for the new protocol and firmware on both the CCAP and CPE. Interoperability issues frequently emerged requiring close cooperation with vendors to address. While these troubleshooting fire drills are disruptive and time-consuming, they are also invaluable learning opportunities for the engineers involved. Working side-by-side with the vendors' subject matter experts allowed Vodafone NZ personnel to gain a deeper understanding of the service provision equipment and DOCSIS protocols than would have otherwise been possible through routine system operation. Anyone who has been through the early adopter experience and the invaluable learning that it provides knows the truth and wisdom behind the quote "that which doesn't kill you only makes you stronger."

Outside of this specific vendor example there are other common themes encountered in working with operators globally on DAA planning and deployments. One is that planned/specified test capabilities often lag behind the launch of new network gear. DOCSIS 3.1 has been deployed in the field for quite some time, yet we are just now starting to see NEMS implement some of the exciting PNM capabilities documented in Section 9 of the DOCSIS 3.1 PHY Layer Specification. Understandably, functionality to ensure the successful passing of packets receives top priority, but as a result, test capabilities often don't release until well after CCAPs and CPE are fielded. The same has held true for the NDF/NDR capabilities documented in the R-PHY spec – few vendors have released support to date although most have a committed roadmap to do so.

General Conclusions

Through work with cable operators globally in developing best practices the following framework has been established for recommended testing for Remote PHY rollouts. This list is generally applicable to all DAA rollouts with the exception of PTP timing tests that apply specifically to Remote PHY.

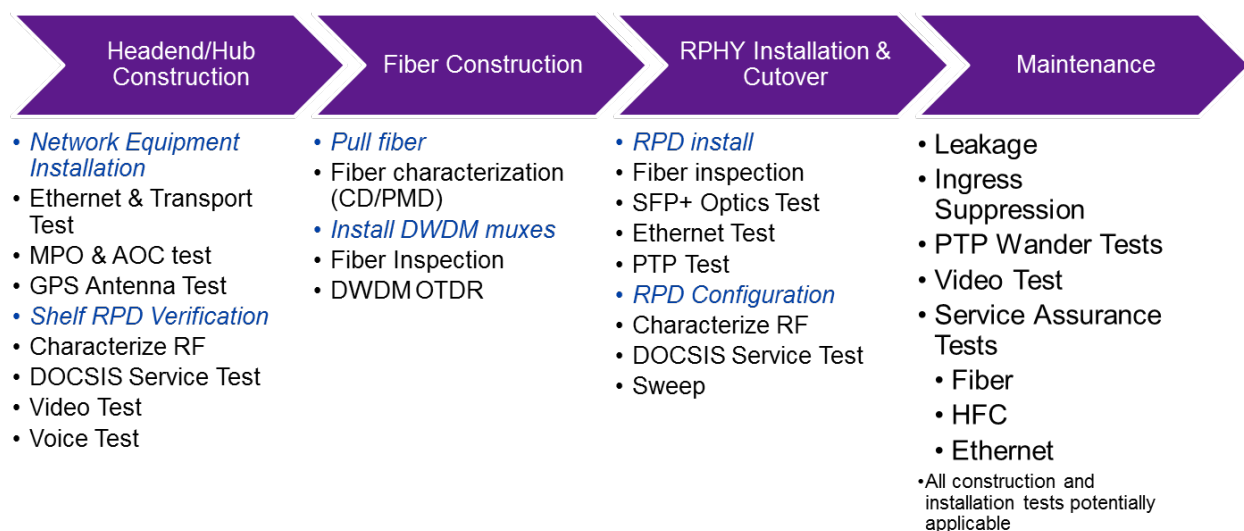


Figure 9 - Framework for DAA Deployment and Maintenance

Again – detailed discussion of each point in this framework is beyond the scope of this document but a brief description is provided below with general test content and rationale.

1. Headend/Hub Construction

The goal of this section is to ensure that the infrastructure needed to feed an R-PHY deployment in the field is functional and correct before any work is done to validate the field aspects. Ethernet and transport gear can be tested including optical cabling interconnections (MPO & AOC test) as well as GPS antennas required for PTP functionality. Shelf RPD verification can be used to pre-test the hub/headend components and services delivered before field R-PHY gear is deployed.

2. Fiber Construction

This section focuses on validating that fiber has been successfully deployed without any damage or excessive bending. Like most aspects of new technologies, the secret to success often lies in disciplined execution of the fundamentals. Techs need to remember “inspect before you connect” – dirty or damaged fiber interfaces are still a leading cause of fiber network performance issues. The DWDM muxes will often be a new ingredient for many HFC plants and will require specialized equipment to properly validate.

3. R-PHY Installation and Cutover

This is the most exciting part of the process, where end-to-end performance can finally be tested for the first time. As part of the RPD install process the SFP and related optics must be tested after ensuring that the fiber is clean and undamaged. Once operational, Ethernet testing is required, as well as verifying adequate throughput and that timing is working properly via PTP testing. As a matter of good practice this is an appropriate time for sweep to characterize the RF performance and to ensure DOCSIS services are operational and working at full performance as well.

4. Maintenance

Most of the plant issues that techs spent their time addressing in legacy plant will still be present in DAA nodes (minus the dreaded analog optical link setup!). If the R-PHY planning process, including system/tool selection has been thoughtfully carried out, the tools and processes that techs use will be very similar, other than the changes noted earlier in this paper. Upstream ingress will still be the #1 issue and having capable tools to address this is still critical. Leakage is still a valuable maintenance capability, even in shorter cascades that often accompany DAA implementations, for localizing potential upstream ingress sources along with many other plant issues. Most RPD vendors have integrated leakage tagging capabilities to enable continued use of deployed field gear for drive-out and walk-out use cases. One new test for techs working on R-PHY nodes will be the PTP wander test as a validation and troubleshooting tool. Service assurance is also still important. Be sure to consider how this will be handled for HFC, fiber and Ethernet aspects of the network. It’s also important to note that all construction and installation tests are also potentially applicable in the maintenance category.

Conclusion

Hopefully after reading this paper you have a better understanding of how DAA is enabling the explosive node split rate required to meet subscribers’ ever-increasing broadband appetite. Redistributing certain network functions from overcrowded/overheated hubs out into the field eliminates many barriers to

gigabit service delivery but creates challenges. At the heart of most challenges is the new optical media translation interface that is created by the redistribution. Understanding the technical and organization challenges will determine whether a project succeeds or fails. Specific challenges including the elimination of RF in the hub for test gear has been addressed by building virtualization capabilities into the RPD's themselves and orchestration software to tie it all together. All this theoretical guidance is great, but the Vodafone New Zealand example demonstrates that DAA can indeed resolve critical business problems for operators and continue to be maintained by the same workforce with mostly same field tools and processes. Keys to success include careful planning of the architecture, and the deployment/validation process, and plant maintenance during what is for most a multi-year transition period. Finally, a framework was provided to guide thought regarding test considerations for the entire lifecycle of a typical Remote PHY deployment.

Abbreviations

AOC	active optical cable
CAA	centralized access architecture
CCAP	converged cable access platform
CMTS	cable modem termination system
COTS	commercial off the shelf
CPE	customer Premises Equipment
CWDM	coarse Wavelength Division Multiplexing
DAA	distributed Access Architecture
DTP	DOCSIS timing protocol
DWDM	dense wavelength division multiplexing
HERD	headend rearchitected as datacenter
HFC	hybrid fiber-coax
MPO	multi-fiber push on connector
NDF	narrowband digital Forward
NDR	narrowband digital Return
OTT	over the top (video)
PTP	precision Timing Protocol
R-CCAP	remote ccap
R-PHY	remote phy
RCI	remote phy / ccap interface
RPD	remote phy device
SCTE	Society of Cable Telecommunications Engineers
SFP	small form-factor pluggable
SNR	signal to noise ratio

LoRa, The LPWA Choice For Cable Operator's Entry Into The Smart Home Market

A Technical Paper prepared for SCTE•ISBE by

Bill Beesley

Principal Solution Architect
Fujitsu Network Communications
2018 Telecom Parkway, Richardson, TX 75082
972.479.2098
bill.beesley@us.fujitsu.com

Ladan Pickering

Software Strategic Planner
Fujitsu Network Communications
2801 Telecom Parkway, Richardson, TX 75082
972.479.2371
ladan.pickering@us.fujitsu.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
Conclusion.....	6
Abbreviations	7
Bibliography & References.....	7

List of Figures

Title	Page Number
Figure 1 - IoT Demands a New Approach to Connectivity.....	4

Introduction

According to IHS estimates, the Smart Home market is expected to grow at CAGR of 26% to reach 75.4 billion devices by 2025. LPWA IoT solutions are low cost, simple, and effective approach to solving automation and security in homes. This paper will describe how Low Power Wide Area Networking using LoRa gateways and endpoints can allow operators to support reliable connectivity for IoT solutions while minimizing capital and operational support costs. The paper will also outline unique use cases for LoRa and IoT that will provide innovative new revenue opportunities for cable operators. A survey of LPWA technologies in this paper will show that LoRa is the leading choice. Since Cable operators are already in the homes and have the relationship with the customers, they are in a unique position to offer automation and security solutions and win a large portion of the Smart Home market.

Content

Unless one has been living under a rock for the last five years, they have noticed an explosion growth of connected devices, often referred to as the Internet of Things, or IoT. According to a recent Forbes article, the analysis firm Gartner has predicted that the number of connected devices will grow from 8.4 billion in 2017 to 20 billion by the year 2020. These will include the familiar smart home devices such as thermostats, garage door openers, doorbells, lightbulbs and personal assistants. However, to reach the device numbers predicted by the analysts we are going to have to think beyond connecting the devices we use in our homes. In short, IoT is going to have to go outside and play and as it does they are going to have to deal with challenges in both networking and battery life that will present opportunities for service providers. Over the next few years we will see IoT bring “smarts” to our cars, air quality monitoring in cities, parking spaces, trash cans and cows. Yes that’s right, we are going to have an “Internet of Cows”.

The Internet of Cows Problem

In 2013, Fujitsu Limited rolled out its GYUHO software as a service solution targeted to support the food and agricultural industry. One of the interesting use cases of the technology was to detect when cows were going into estrous. As it turns out, cows start to move around more as they begin to enter their prime mating window and by attaching pedometers with wireless radios, ranchers can monitor their herd and receive alerts when a cow is about to go into estrous. In a test group of 40,000 cattle, the software allowed the ranchers to improve their fertilization rates from 1 in 2 to 2.5k attempts to 1 in 1.58 mating attempts. In addition to optimizing mating times to achieve better impregnation rates, they are also able to determine the sex of the calf with a 70% accuracy. Improving both of these metrics allow the ranchers to achieve more profit while lowering their feed and other operational costs.

Our cow use case has both of the requirements that current radio technologies such as WiFi or 4G/5G are not able to support well. The first challenge is one of proximity to WiFi hotspots or cell towers. Cows generally do not live in metropolitan areas where these facilities are often in such abundance as to be taken for granted. To support our Internet of Cows economically both from a capital and operating expense, we need to be able to reach multi-kilometer distances with low cost devices. The data throughput requirements are minimal since we do not anticipate our cows will be streaming Netflix, but we do expect them to be able to easily wander over a large area. The second challenge is one of battery life. It is just not feasible to expect our cows to drop in for a recharge every few hours, and the intent is to make the rancher’s lives simpler and more profitable so changing out batteries on a regular basis is an expense and effort we want to avoid. The technology used to support our Internet of Cows would ideally be able to run for months or years before having to have batteries replaced or recharged.

It is not just our Internet of Cows that need support for low data rate, low power and long distance. There are multitudes of commercial IoT use cases that have similar requirements.

In summary, we are looking for a technology that can support relatively low data rates over a large distance and with low power demands to support long battery life.

One of these use cases that is getting a lot of press lately is the idea of “Smart Cities”. Kansas City, Kansas has been working to create a 51 block Smart City in their city core as part of an initial roll out of advanced services citywide. Within this area citizens not only have access to free WiFi, but the city has also connected its safety infrastructure such as street lights and emergency vehicles to the network to provide better access to information, and to support the ability to more closely and efficiently manage this infrastructure. In addition, they have deployed trashcans that have sensors to analyze temperature, humidity and fill level. This allows the city maintenance staff to determine exactly where and when public trashcans need emptying and route their resources more efficiently. This has allowed them to greatly reduce the amount of staff and fuel previously used visiting every single trashcan in the city, regardless of whether it was in need of service. Now, while the trashcans are located well within the range of traditional WiFi, the trashcans do not have need of WiFi data rates, nor would it be ideal to burden replacing batteries during each trashcan visit as this would offset much of the operational savings. Solar could be one option, but not all trashcans are in direct sunlight so ideally, this solution would be supported by a radio infrastructure capable of supporting low power and therefore, extremely long battery life.

	Optimized for	Direction	Latency	Cost	Complexity	Spectrum
Cellular & WiFi	Voice and Data	Large Download, Small upload	Low	High (Setting up a network is very expensive).	High (Connection based)	Licensed
IoT	Cost, Battery life	Large Upload, small download	NA (Wide range of use-cases)	Low. Can build private network	Low (Connectionless)	Unlicensed




Figure 1 - IoT Demands a New Approach to Connectivity

As we can see from Figure 1, the primary requirements for commercial IoT deployments are focused around low cost and long battery life instead of optimizing for voice and data services as traditional wireless networks were designed to support. If we are going to optimize a network to support the needs of billions of small, chatty devices that need to be left out in the wild to live for years without human intervention, then new radio technologies must be considered.

There are several technology standards that are emerging to support IoT and they all have their strengths and shortcomings. In this paper, we will compare four of the most popular, specifically comparing them

against LoRa. These are Long Term Evolution Category M1 (LTE-M), Narrow Band IoT (NB-IoT), WiFi, and of course, LoRa. SigFox was specifically not considered because this is a closed network run by a private company. Random Phase Multiple Access (RPMA), an Ingenu product and nWave's Weightless protocol primarily used in smart meters were also not considered, not because of any technical shortcomings but because they lack any significant adoption beyond the companies who started their development.

LTE-M was developed as part of the 3rd Generation Partnership Project (3GPP R13, 2016) standards work. It supports data rates from tens of kbps to 1Mbps. It has an advantage for existing wireless providers in that they could leverage their existing network. It has disadvantages in that it is not yet commercially available, although there are pilots. Additionally, the modules are expensive relative to other IoT options and when available it is expected that license fees will be charged per device, and that there will be fees for both data connectivity and usage.

NB-IoT is also a part of the 3GPP R13 2016 standards development. It was designed primarily for low data rates of less than tens of kilobytes per second and as with LTE-M should allow providers to leverage their existing infrastructure. But as with LTE-M, devices are still in the trial phase and it is expected that fees similar to existing wireless infrastructure will be incurred on the device.

802.11 WiFi is another technology heavily leveraged by the IoT industry, primarily because it is a well-understood, global standard known for its ubiquity and ease of use. It supports high bandwidth data rates but at the cost of high power demands that impact battery consumption. It is also a relatively low range technology with poor building penetration.

LoRa was established as a standard by the LoRa Alliance (<https://lora-alliance.org/>) in 2015. The standard utilizes unlicensed spectrum in the 915Mhz band in North America to support very low power transmission over distances greater than 10 kilometers in rural areas. It has advantages of a built-in security model, low battery usage, low cost and very long range with the ability to also penetrate buildings. Data rates are very low ranging from .3 to 50 kbps, but for many commercial use cases this rate is acceptable. Currently end devices are in the \$20 range with targeted prices of \$2-\$5 as volumes increase and gateway costs of around \$500-\$700 with a gateway being capable of supporting more than 62 thousand devices depending on how often they are transmitting. In addition, the LoRa network is open with no connectivity or data fees, or licenses on a per device and or gateway basis making it ideal for high volume, low cost IoT deployments.

There are currently world wide deployments of LoRa gateways by both public and private consortiums. More than 83 network operators such as The Things Network (<https://www.thethingsnetwork.org>) and Senet (<https://www.senetco.com>) are in operation, and they are working aggressively to promote and expand deployments. Companies or even individuals can join groups like the Things Network simply by purchasing and connecting a LoRa gateway. There are also a multitude of commercial IoT devices available for purchase as well as open hardware and software development kits designed to seed innovation.

What is in it for the Cable Operator?

All of the top tier cable operators in North America have invested heavily in the deployments of public WiFi networks and they should continue to do so as these networks will be in high demand for IoT and end user usage that demands low latency and high bandwidth. The consideration of deploying LoRa technology will require consideration of new approaches to monetizing the network, as LoRa does not provide a mechanism for charging the end user or device for access or consumption as operators would traditionally have done. IoT will instead offer new business models whereby the suppliers of hardware,

software and most importantly, services will profit from the technology investments. In our Internet of Cows example, the rancher was provided the entire solution as a service which included charges for the equipment and recurring charges for the monitoring services. The Kansas City trash can example was also delivered as a turnkey service to the City. There will inevitably be open consortiums like The Things Network, but these are intended to encourage innovation, not provide sustainable commercial solutions. There will also be new entrants like Senet who will attempt to build out new networks, but they either will largely be leveraging other providers' infrastructure or starting from ground zero, neither of which is an economic nor a time to market advantage. Cable Operators are uniquely positioned to begin offering turnkey IoT Solutions to markets and end customers. The existing HFC network infrastructure is currently servicing around 120 million households in the US, providing cable operators with the fundamental infrastructure necessary to begin deploying IoT infrastructure. Having not only last mile access infrastructure but also infrastructure that reaches into the house and business is a great advantage in deploying new services. Security services for homes and small business is a past example of where Cable Companies have been able to successfully develop a new service offering. As such, Cable operators are in a unique position to create new revenue streams supporting end-to-end IoT solutions, especially for emerging smart cities initiatives such as the Kansas City example. Not only does this have the potential to add new revenue to the bottom line but by offering this as a service to the franchise authority to create Smart Cities, they can further strengthen their partnership in the markets they are serving.

Conclusion

IoT deployments will continue to accelerate and offer an increasingly diverse set of new services to support the needs of consumers, businesses and municipalities. In order to support low cost, long range and long battery life, new transport technologies will need to be deployed in order to create a ubiquitous and always available network. Cable Operators are uniquely positioned to begin building these networks to support a variety of new services capabilities that will allow them to continue to serve their customers with the most advanced technologies and solutions.

Abbreviations

3GPP	3 rd Generation Partnership Project
IoT	Internet of Things
Kbps	Kilobits per second
LPWA	Low Power Wide Area
LoRa	Long Range
LTE-M	Long Term Evolution Category M1
Mbps	Megabits per second
NB-IoT	Narrow Band IoT
RPMA	Random Phase Multiple Access
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

Developing The Connected World Of 2018 And Beyond , Lars Knoll, Forbes March 16, 2018
<https://www.forbes.com/sites/forbestechcouncil/2018/03/16/developing-the-connected-world-of-2018-and-beyond/#328ea3971e51>

Fujitsu to Rollout Global Sales of "GYUHO" SaaS, Fujitsu Press Release, October 15, 2013
<http://www.fujitsu.com/global/about/resources/news/press-releases/2013/1015-01.html>

Low Power Wide Area Technologies for IoT Use Cases

Technology Assessment for MSOs

A Technical Paper prepared for SCTE•ISBE by

Joe Walsh

Managing Director

inCode Consulting

+1 703.835.4386

jwalsh@incodeconsulting.com

Table of Contents

Title	Page Number
Table of Contents	2
Executive Summary	3
Content	4
1. LPWA Market Overview	4
1.1. LPWAN Set to Overshadow Traditional Cellular IoT Connections 2X.....	4
1.2. Licensed LPWA to attain a rapid growth at 49% CAGR; Unlicensed growth will prevail, but slower	4
1.3. Unlicensed will sustain a niche for Private networks, however Licensed Public deployments will far outpace.....	5
1.4. All Operator roads lead towards NB-IoT, though approaches may vary	6
2. Technology Selection by Use Cases	7
2.1. Framework for LPWAN Technology Analysis	7
2.2. Competitive boundaries between licensed and unlicensed set to settle at 60-40 split of applications	9
2.3. Case Study: Smart Bike	10
3. Roadmap and Deployment Options	11
3.1. LoRaWAN will have an opportunity to gain a foothold in the US however global momentum for 3GPP will tilt ecosystem development.....	11
3.2. 5G aligned developments will fortify NB-IoT's future and further bolster the ecosystem development.....	11
3.3. Module costs & batteries will continue to decline, spurring innovation and opening LPWA to new use cases	12
3.4. NB-IoT can complement LoRaWAN and enable MSO to win mega deals	13
Abbreviations	14
Bibliography & References.....	14

List of Figures

Title	Page Number
Figure 1 - U.S. IoT Connections Breakdown By Technology	4
Figure 2 - U.S. LPWAN IoT Connections Breakdown By License Type.....	5
Figure 3 - U.S LPWAN IoT Connections Breakdown By Network Type	6
Figure 4 - U.S. MNOs IoT Evolution	7
Figure 5 - Factors for Comparing LPWAN Technologies	7
Figure 6 - Technology Suitability for Top 20 LPWAN Applications	9
Figure 7 - Technology Suitability for Top 20 LPWAN Applications	10
Figure 8 - 3GPP vs LoRaWAN Development	11
Figure 9 - 3GPP IoT Standards Development	12
Figure 10 - Module Cost Forecast	13
Figure 11 - Potential LPWA Market Options and Consequences	13

Executive Summary

IoT is set to expand at a very rapid pace over next few years. While the majority of IoT connections in the United States of America are based on Wi-Fi and other short-range technologies, Cellular (2/3/4G) and LPWA currently account for ~20% of the IoT connections and are expected to grow at a CAGR of 37% between 2018 to 2023. The LPWA segment is set to grow at an astounding 58% CAGR and touch 512 million connections by 2023, while cellular IoT connections will grow at 16% CAGR and rise to 206 million – accounting for less than half of LPWA connections. As operators race to compete with new models to serve LPWA connections, the connectivity revenue per connection will also continue to decline at over 13% CAGR, thanks to the range of options on the table.

Within LPWA category, the licensed technologies are set to grow faster than unlicensed technologies. Licensed LPWA market is forecasted to grow at 49% CAGR ('18 – '23) reaching \$1,836 Million (based on 306 Million connections) in 2023, while unlicensed LPWA market will grow at a slower pace - 37% CAGR – to capture \$718 Million (based on 206 Million connections). Customers will look to deploy LPWA solutions as either Public or Private network depending on use case requirements, addressable market, ecosystem maturity, security, and regulatory considerations. While the unlicensed Private networks will remain favorable for certain enterprise verticals, nearly 80% of the growth will come from licensed public connections once LPWA scales.

To select the right LPWA technology for a given use case, MSOs should evaluate technologies across three key areas: Technology Fit, Operational Impact and Business Relevance.

- Within technology fit, the major evaluation factors include throughput (or payload size), battery, mobility, and tracking requirements.
- Operational Impact attributes to be considered are Quality of Service, deployment type (private vs public networks), coverage, and topology (urban vs. rural).
- Finally, the business relevance criteria encompass time-to-market and operator incumbency.

Based on the above framework, critical applications requiring high QoS and accurate tracking deplete battery life and will need licensed technologies like NB-IoT. Non-critical and delay-tolerant applications will likely prefer non-licensed technologies like LoRaWAN. The middle market applications will accommodate both technologies; however, NB-IoT is considered a favorite. Segments with low operator presence, like agriculture, have initial propensity for unlicensed technologies. Thus, LoRaWAN can carve out a market niche in rural, smart-city, and utility sectors, while the price sensitive middle-market and the premium critical apps market will likely be NB-IoT's home turf.

Ongoing 3GPP standardization of licensed LPWA technologies – LTE-M, NB-IoT, EC-GSM – will continue to increase the applicability and robustness of these technologies to a wide variety of use cases. Engineering advances will continue to reduce module prices. With Operator subsidies, module prices for licensed LPWA technologies will likely be at par with unlicensed LPWA technology modules in the next 3 to 5 years. Similarly, improvements in battery technology and form factors will not only drive adoption of existing use cases but also enable new use cases.

Cable operators (MSOs) deploying LoRaWAN have a window of two to three years to monetize it before wireless operators (MNOs) deploy nationwide NB-IoT and/or LTE-M networks with enabled devices/modules. MSOs should consider building a regional NB-IoT network as well as evaluate means to provide nationwide IoT coverage via roaming and / or wholesale agreements.

Content

1. LPWA Market Overview

1.1. LPWAN Set to Overshadow Traditional Cellular IoT Connections 2X

The Low Power Wide Area (LPWA) market is a fast-growing market with some wide array technologies available to address the multitude of IoT use cases. Majority of the IoT connections are Wi-Fi and short-range technologies like Zigbee, Bluetooth, etc. Cellular and LPWA together make up less than 20% of the IoT connections. In 2017, the total number of Cellular and LPWA IoT connections in the United State of America were about 108 million. These connections are expected to be 722 million by 2023, growing at a CAGR of 37%. The LPWA subset will proliferate at an 58% CAGR to reach 512 million connections by 2023, while cellular IoT connections will grow at 16% CAGR and rise to 206 million – accounting to less than half the number of LPWA connections. Factors like reduction in module costs, improvement of battery form factor, increased clarity on technology applicability by use case, and demonstration of Return on Investment (proof points) will likely be the key drivers of this growth.




Figure 1 - U.S. IoT Connections Breakdown By Technology

1.2. Licensed LPWA to attain a rapid growth at 49% CAGR; Unlicensed growth will prevail, but slower

The LPWA market can be segmented by the type of spectrum a technology uses, resulting in two segments; i.e., licensed and unlicensed technology markets. Technologies like LoRaWAN, Sigfox, RPMA, and Telensa operate in the unlicensed spectrum, whereas LTE-M, NB-IoT, and future 5G IoT technologies will leverage the licensed spectrum.

On the other hand, the use cases can also be segmented based on expected level of service for a given application, resulting in non-critical, middle ground, and critical applications. Applications like Smart Meters are generally classified as non-critical as the mobility and latency requirements are not stringent, while applications like people or pet tracking are considered critical applications as they require a high

degree of QoS. An example of a middle ground application is commercial automation which has moderate QoS requirements and delay tolerance.

We expect the Unlicensed LPWAN market to grow at 37% CAGR from 2018 to 2023. In year 2023, the unlicensed technologies are expected to connect about 206 million connections and generate revenues of \$718 million. In comparison, the licensed technologies will grow at 49% CAGR during the same period and reach 306 million connections and \$1,835 million in revenues by 2023.




Figure 2 - U.S. LPWAN IoT Connections Breakdown By License Type

1.3. Unlicensed will sustain a niche for Private networks, however Licensed Public deployments will far outpace

Enterprises will choose between private and public LPWA networks based on addressable market, ecosystem maturity, use case requirements, security, and regulatory considerations. Based on the current state of the ecosystem and the hurdles in the realization of “Private LTE” networks, it is expected that Unlicensed LPWA will be more favorable for these enterprise private networks.

Some enterprises will choose to leverage Private networks driven by:

1. **Lowest cost of service:** Private networks will continue to have lower Total Cost of Ownership (TCO) at least for next few years, especially for sectors with built-in long-term business case needs like Utilities. It is critical that these networks are easy to deploy and operate.
2. **Security / Regulatory:** Increased focus on security and need for regulatory adherence for private network deployment.

The majority of enterprises will prefer Public networks, primarily for two key reasons:

1. **Use Case Flexibility:** Public networks can support a wide variety of use cases.
2. **Industry Momentum & Ecosystem:** Globally, the majority of the networks deployed are forecasted to be public as the number of vendors, devices, and systems proliferate.

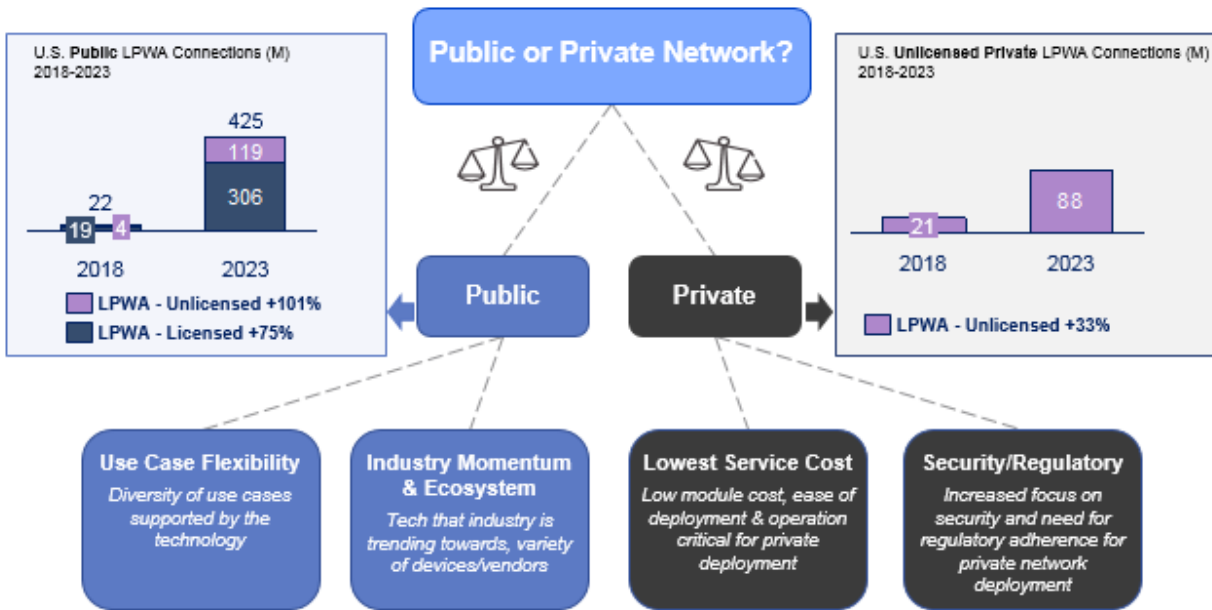


Figure 3 - U.S LPWAN IoT Connections Breakdown By Network Type

1.4. All Operator roads lead towards NB-IoT, though approaches may vary

Cable MSOs have entered the IoT market but have mainly focused on unlicensed technologies. Comcast launched LoRaWAN-based machineQ in late 2016 and has acquired spectrum to launch any cellular IoT service, if it wishes to do so. Cox has launched an asset management solution under the Cox2M brand. Cable operators face a critical decision on technology choice between the unlicensed and licensed technologies, as well as between NB-IoT and LTE-M, if pursuing licensed technologies for their IoT network.

Wireless carriers (MNOs) like AT&T and Verizon initially launched LTE-M service. Given the global industry momentum of NB-IoT, MNOs have now announced NB-IoT launches, while T-Mobile already has launched NB-IoT and is planning for an LTE-M network too. MNOs are also launching their connectivity management and application enablement platforms; e.g., ThingSpace by Verizon and M2X Flow by AT&T, as well as partnering with other ecosystem players for analytics and services. Similarly, Sprint is building an “operating system” in partnership with its sister-company and chip manufacturer, Arm, to manage device connectivity across cellular, Wi-Fi, and LoRaWAN networks.

Overall, all mobile and cable operators are leaning towards NB-IoT as it enables them to uniquely and efficiently address several of the IoT use cases, has a broad ecosystem, strong support for roadmap development, robust security, and module costs will almost be at par with unlicensed technologies in next few years (at least enabled through bundled pricing models).

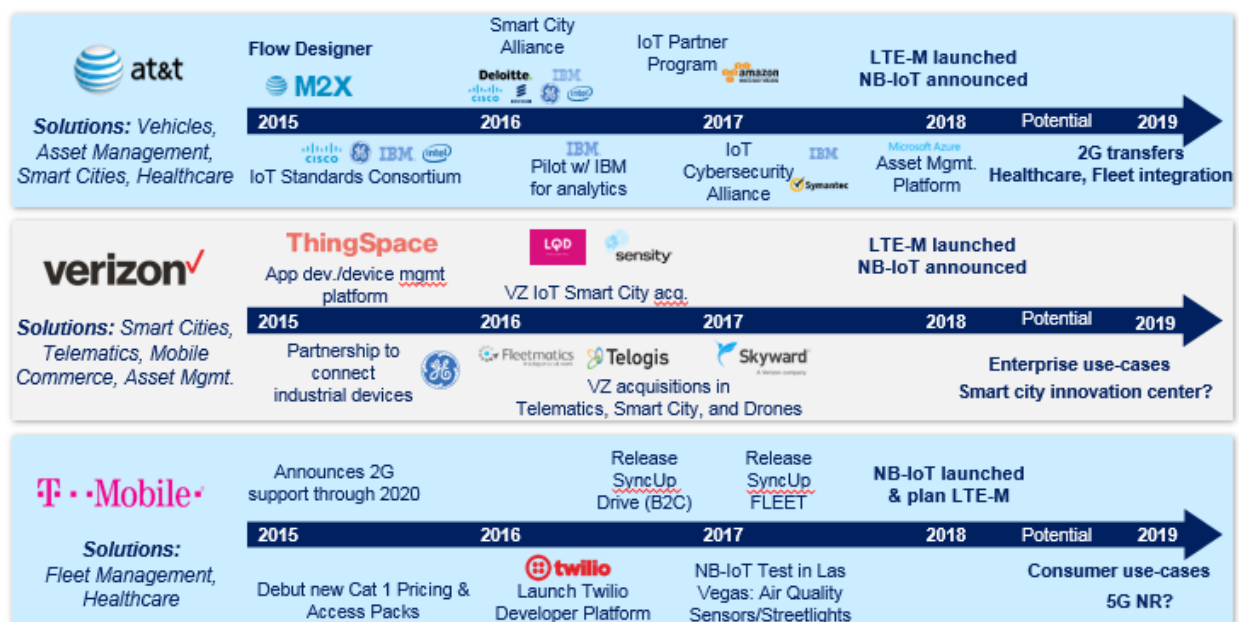


Figure 4 - U.S. MNOs IoT Evolution

2. Technology Selection by Use Cases

2.1. Framework for LPWAN Technology Analysis

To better understand which technologies are best suitable for a target use case, one needs to analyze the use cases across three key criteria – Technology Fit, Operational Impact, and Business Relevance.

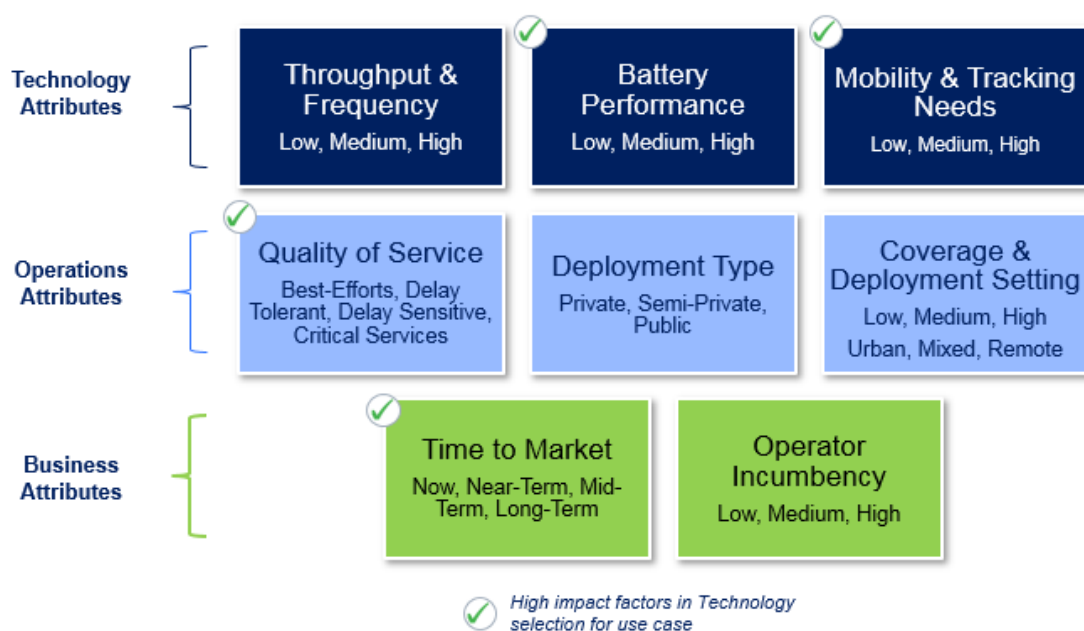


Figure 5 - Factors for Comparing LPWAN Technologies

- (1) **Technology** requirements like *throughput (or payload size), frequency of transmission, spectrum used, battery performance, and mobility*. Applications like smart meters and agriculture sensors have small payload and infrequent message transmission requirements which can be served by LoRa. Whereas NB-IoT is better suited for public safety or ride sharing apps which require frequent broadcast of sizeable payloads and mobility support. Applications with acute needs to maximize battery performance are most serviceable by LoRaWAN.
- (2) **Operations** requirements like *Quality of Service, deployment type, and coverage*. Given that unlicensed technologies are best-effort, they can't guarantee a QoS level. So, delay-sensitive / critical applications characterized by real-time low latency needs are better served by licensed, cellular technologies like NB-IoT or LTE-M. In terms of deployment type, private networks leverage unlicensed technologies. Utilities and Industrials are likely to prefer private networks. Additionally, coverage and distribution impacts use case addressability and level of service. Cellular technologies generally have good urban coverage and have coverage holes in rural, sparsely populated areas, indoors or underground locations. Such areas are sweet spots for LoRaWAN (and other unlicensed technologies). Utilities, agriculture, industrials settings are where LoRa can maximize their advantages, closely followed by home security and assisted living apps.
- (3) **Business** requirements like *Time to Market and operator incumbency*. MSOs looking to deploy LPWA networks and go-to-market quickly initially chose LoRaWAN or other unlicensed technologies as these had a head start over licensed technologies. As 3GPP standards are completed and ecosystem matures, the time-to-market advantage of unlicensed technologies is degrading quickly. Finally, existing client relationships will affect MSO's offerings and technology selection. Applications like Fleet Management have a strong wireless operator incumbency and MSOs may find it difficult to penetrate. On the other hand, agriculture, industrials, and smart cities are some of areas offering great potential for MSOs to launch IoT offerings.

A simplified summary of technology suitability by LPWA's likely target applications is presented below:

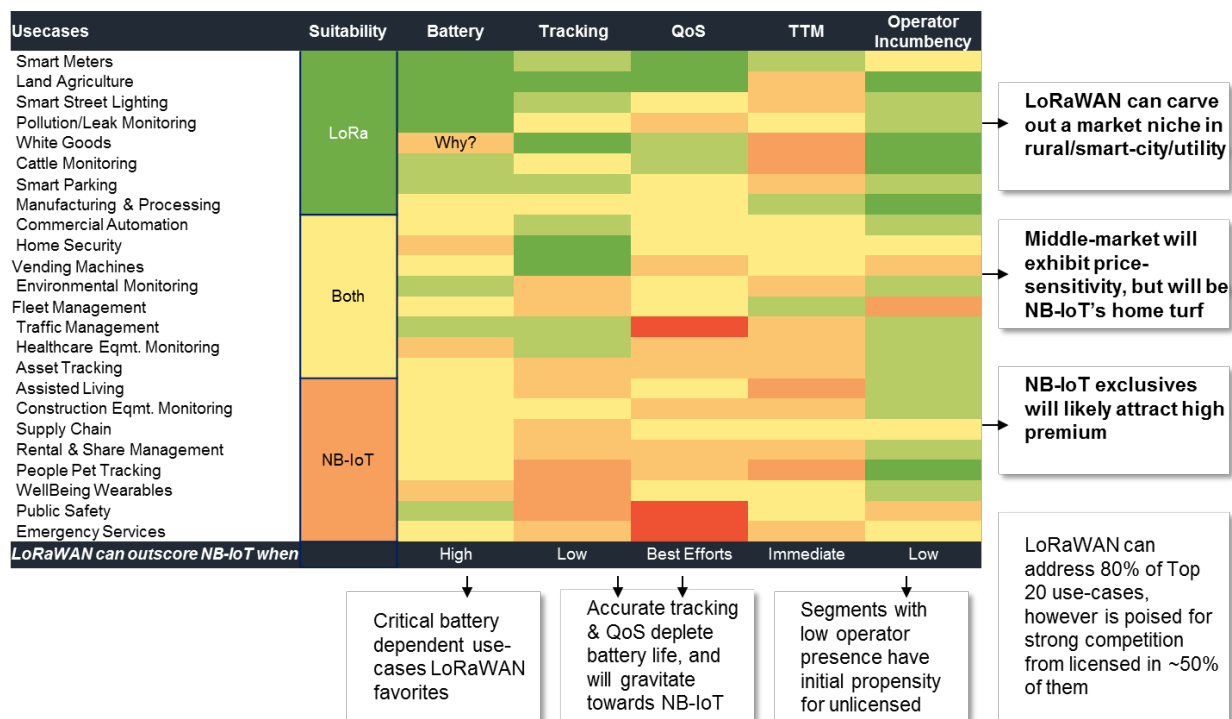


Figure 6 - Technology Suitability for Top 20 LPWAN Applications

As seen, while LoRaWAN can potentially carve out a market niche, it will likely fall behind NB-IoT in broader market addressability.

2.2. Competitive boundaries between licensed and unlicensed set to settle at 60-40 split of applications

By 2023, licensed LPWAN connection will account for 60% of LPWAN connections, while all unlicensed LPWAN technologies will account for 40%. A simplified view of top twelve applications and corresponding forecast of technology connections is show below:




Figure 7 - Technology Suitability for Top 20 LPWAN Applications

2.3. Case Study: Smart Bike



Service providers globally are leveraging a combination of LPWA technologies to address unique customer painpoints while setting new benchmarks for innovation. For example, Ofo, a leading Chinese bike sharing company, offers what is termed docking-free pushbikes for rent, meaning bikes can be collected and deposited from any legal parking spot. Users can locate bikes using their smartphone and unlock it by scanning a barcode. Ofo developed an IoT smart lock based on NB-IoT technology that lowers power consumption, enables wide-area coverage, and slashes system resource delays at low cost. NB-IoT lets Ofo ensure it has bikes located at key locations when commuter demand is highest. Meanwhile, bikes can be unlocked in less than 5 seconds. Both improvements have greatly boosted user satisfaction and enabled Ofo's transformation from a new startup to global bike-rental company with over 10 million bikes and a valuation of over \$2 Billion in under 3 years.

The company has even gone a step further to begin equipping its bicycles with LoRa devices and wireless RF technology (LoRaWAN) to complement its licensed connectivity option to achieve full network connectivity even in remote areas and dense buildings. As clearly seen, unlicensed and licensed LPWA technologies can co-exist and even be complementary in enabling the art of the possible. As "sharing economy" heavyweights Lyft and Uber both jump into bike-sharing (with Lyft's purchase of Citi Bike reportedly for over \$250M in July 2018), it's clear that when it comes to LPWA, it's still Day One!

3. Roadmap and Deployment Options

3.1. LoRaWAN will have an opportunity to gain a foothold in the US however global momentum for 3GPP will tilt ecosystem development

LoRaWAN has been growing in the US as it has low module and servicing costs, strong battery performance, and does not require licensed spectrum. Additionally, with an early start in 2015, it has been deployed and has a strong ecosystem, that is likely to thrive for the next several years.

However, when looking long term, it seems that the balance will tilt in favor of licensed, 3GPP technologies. Operators can leverage their existing infrastructure to deploy and expand the licensed LPWA technologies. 3GPP has successfully developed many Mobile technology standards and has structure processes to define, simulate, adapt and launch global standards. These consistent standards spur innovation as well as interoperability with other devices and systems. Security, a critical aspect of IoT connectivity, is greater in 3GPP licensed technologies. Thus, while LoRaWAN will continue to thrive, 3GPP LPWA will grow significantly and overtake LoRaWAN ultimately.

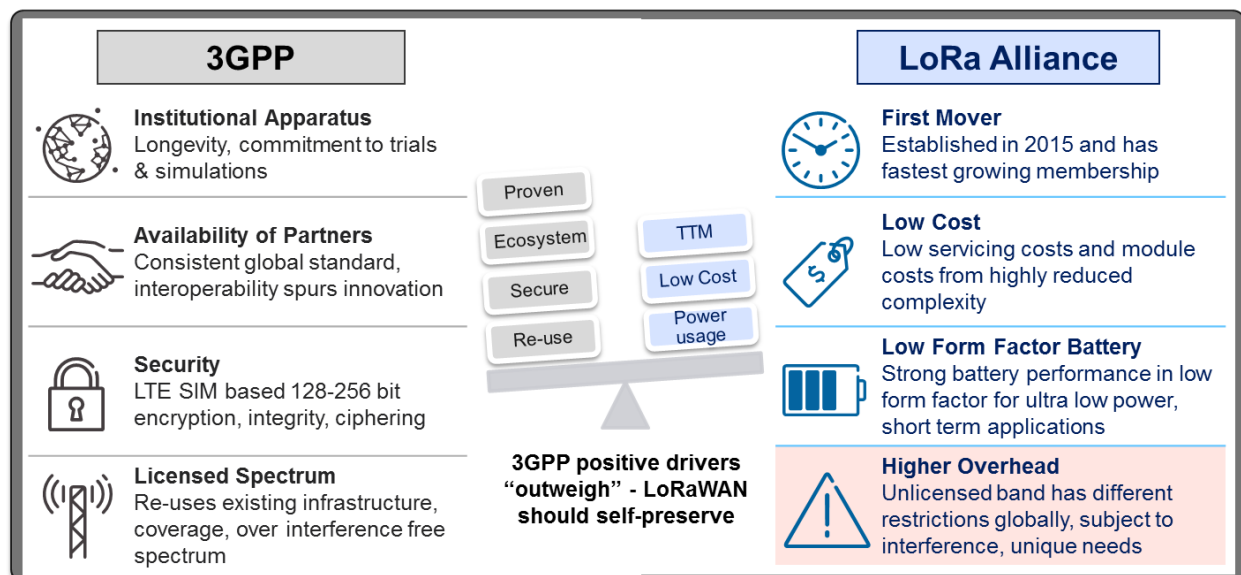


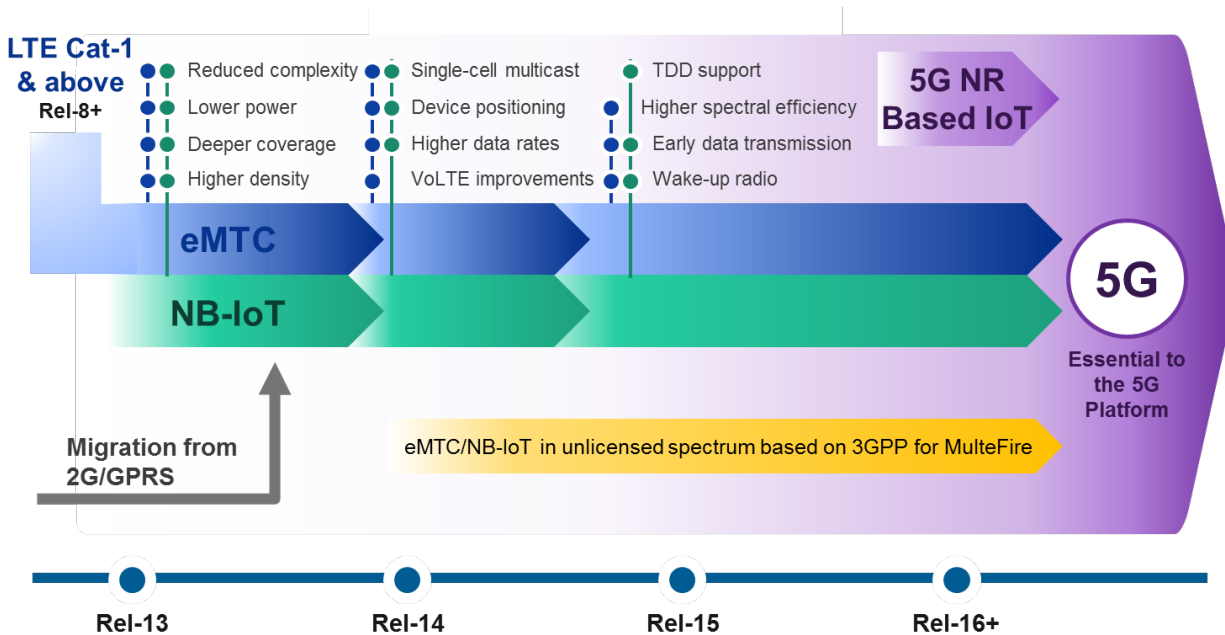
Figure 8 - 3GPP vs LoRaWAN Development

3.2. 5G aligned developments will fortify NB-IoT's future and further bolster the ecosystem development

LTE-M and NB-IoT were introduced in 3GPP Release 13 and provide a complementary cellular solution to enable all IoT use cases by efficiently and cost effectively connecting a wide variety of devices. 3GPP continues to evolve these technologies and bring new capabilities and efficiencies.

Release 14 introduced positioning, broadcast, VoLTE improvements, and two new device categories — Cat-M2 and Cat-NB2. Release 15 will add TDD support for NB-IoT, higher spectral efficiency and power-optimizing features such as wakeup radio and early data transmission. Both, LTE-M and NB-IoT, are agnostic to core networks; i.e., they will work in 5G Non-standalone (EPC) and Standalone (5G Core) modes. There is support for both LTE-M and NB-IoT in-band deployments with 5G NR. Release 16 will further increase device density and network efficiency using non-orthogonal multiple access (NOMA),

enabled by resource spread multiple access (RSMA). Grant-free uplink will allow IoT devices to send sporadic small data bursts to the network without scheduling, thereby reducing overhead signaling.




Source: Qualcomm

Figure 9 - 3GPP IoT Standards Development

3.3. Module costs & batteries will continue to decline, spurring innovation and opening LPWA to new use cases

When deploying large scale IoT devices, module and battery cost become major considerations from a CapEx perspective. However, these are one-time costs and when compared to recurring revenues from connectivity and professional services (e.g. installation, analytics), these constitute a small portion of the customer's TCO (under 5-8% as estimated for a sample asset-tracking use case). Hence, operators or solution providers should consider exploring bundled pricing plans subsidizing device (module, battery, MCU, casing etc.) costs for price sensitive customers who are hesitating to go forward with their IoT plans or looking for cheaper alternatives.

Range of batteries are available for LPWA like LiPo, Alkaline, Lithium-Thionyl-Chloride (LTC), Coin cell (Wearables), and Ultra-thin wireless nanotag. Form factors continue to be reduced and eventually some apps may not even rely on a battery. E.g. Fujitsu Laboratories has developed the world's smallest LPWA sensor device which is powered by a single solar cell.



- NB-IoT will benefit from economies of scale.
- Less complex LoRa modules need cohesive industry effort to retain price advantage

Figure 10 - Module Cost Forecast

3.4. NB-IoT can complement LoRaWAN and enable MSO to win mega deals

Cable operators (MSOs) deploying LoRaWAN have a window of two to three years to monetize it before wireless operators (MNOs) deploy nationwide NB-IoT and/or LTE-M networks with enabled devices/modules. MSOs who wish to address the full span of use cases or meet diverse customer needs, should consider either deploying NB-IoT network in their footprint or negotiate a wholesale deal with another NB-IoT network provider. Additionally, as the LPWAN battleground matures and connectivity price drops, a nationwide footprint will become table-stakes. To remain competitive, MSOs can evaluate means to provide nationwide IoT coverage via roaming and / or wholesale agreements.

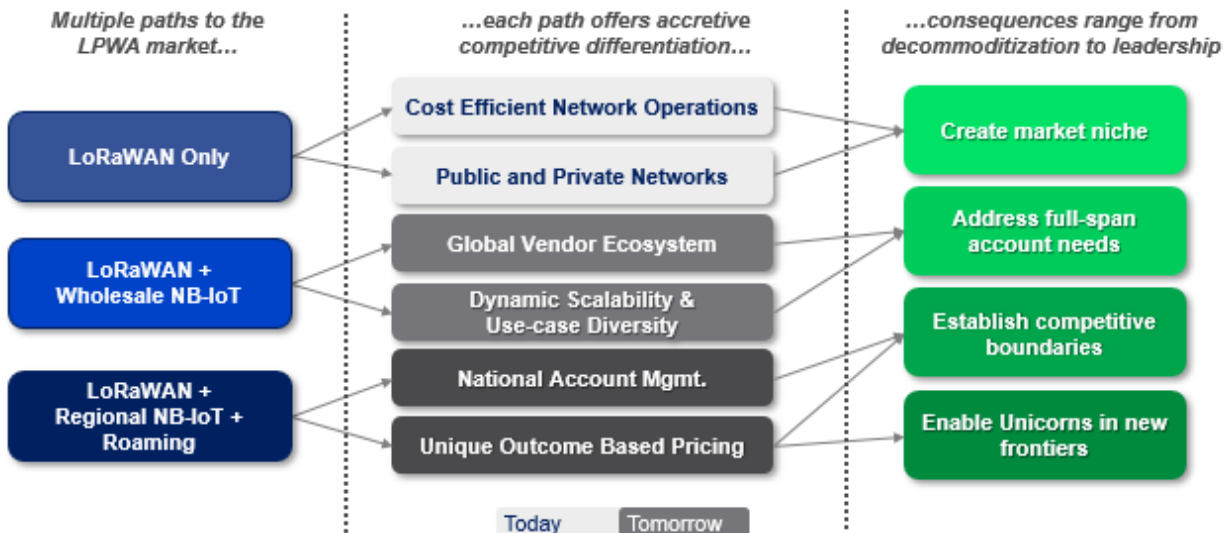


Figure 11 - Potential LPWA Market Options and Consequences

Abbreviations

IoT	Internet of Things
LPWA	Low Power Wide Area
CAGR	Compound Annual Growth Rate
MSO	Multiple Service Operator
QoS	Quality of Service
NB-IoT	Narrow Band – Internet of Things
LoRaWAN	International Society of Broadband Experts
LTE-M	Long Term Evolution For Machines
EC-GSM	Extended Coverage Global System For Mobile
MNO	Mobile Network Operator
RPMA	Random Phase Multiple Access
TCO	Total Cost of Ownership
LoRa	Long Range
3GPP	3 rd Generation Partnership Project
VoLTE	Voice over LTE
TDD	Time Division Duplex
EPC	Evolved Packet Core
NOMA	Non-Orthogonal Multiple Access
RSMA	Resource Spread Multiple Access
MCU	Multipoint Control Unit
LTC	Lithium-Thionyl-Chloride

Bibliography & References

1. ABI Research: <https://www.abiresearch.com/>
2. Analysys Mason: <http://www.analysismason.com/>
3. GSMA: <https://www.gsma.com>
4. Machina Research: <https://machinaresearch.com/>
5. Ofo: <https://www.ofo.com/us/en>
6. 3GPP: <http://www.3gpp.org/>
7. Qualcomm: <https://www.qualcomm.com/>
8. LinkLabs: <https://www.link-labs.com/>
9. Gartner: <https://www.gartner.com/en>

Navigating IoT Technologies, Standards and Frameworks for Managed IoT Service

A Technical Paper prepared for SCTE•ISBE by

Gary Gutknecht
CTO Connected Home
Technicolor
Sugarloaf Pkwy, Lawrenceville, GA
+1-317-809-2417
gary.gutknecht@technicolor.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	4
1. Introduction.....	4
1.1 Smart Home Trend.....	4
1.2 Network Service Provider Opportunity for Managed IoT Service	6
1.3 NSP IoT End-to-End Layers.....	7
1.3.1 Devices.....	8
1.3.2 Connectivity Layer.....	8
1.3.2.1 Messaging Protocols.....	12
1.3.2.2 Message Broker/Gateway Function.....	13
1.3.2.3 IoT Device Management.....	14
1.3.2.4 BSS/OSS Integration	15
1.3.3 Service Layer	15
1.3.3.1 Edge Compute	16
1.3.3.2 Containers for Embedded Implementations	17
1.3.3.3 API Gateway Function	18
1.3.3.4 IOT visualization.....	18
1.4 IoT Architecture Approaches.....	18
1.4.1.1 IoT Gateway Function.....	20
1.4.2 Platform Layer.....	21
1.4.2.1 Messaging Management.....	22
1.4.3 Cloud IoT Platform	23
1.5 Harmonization of Standards.....	25
1.5.1 Connectivity Harmonization	25
1.5.2 Data Model Harmonization.....	26
1.5.3 Service Layer Harmonization	26
Conclusion.....	27
Abbreviations	27
Bibliography & References.....	27

List of Figures

Title	Page Number
Figure 1: NSP End-to-End IoT Service Layers	7
Figure 2 : Super Sensor (1).....	8
Figure 3: Wireless Connectivity for IoT	9
Figure 4: Connectivity Stack Comparison.....	12
Figure 5:Virtualization of IOT services	17

List of Tables

Title	Page Number
Table 1 : Smart Home IoT Evolution.....	5
Table 2 : NSP's Differentiation	6
Table 3: Short Range IoT Connectivity Comparison	10
Table 4 : Messaging Protocols Comparison	13
Table 5: IoT Gateway Stack.....	21
Table 6: IoT platform layers	22
Table 7: Infrastructure-as-a-Service Cloud IoT Frameworks.....	24
Table 8: Cloud IoT Framework Comparison	25

Introduction

Network Service Providers (NSPs) have a major opportunity and advantage in offering Managed IoT Services. They have the organizational and business structure to successfully build all the necessary IoT layers of Connectivity & Networking, Core IoT Platform and Services. However, IoT technologies and ecosystems are complex and fragmented making it challenging for Service Providers to formulate a winning strategy. A significant number of complex heterogeneous IoT options for sensor connectivity, networking and application layers make it challenging to understand the best solution for targeted use cases. Functional overlap is pervasive when considering networking and IoT sensor application layer options such as Thread, Open Connectivity Forum and Dotdot, which makes it difficult to pick the best approach and understand how they will work together. IoT connectivity protocols and standards such as Wi-Fi, Zigbee, Z-wave, BLE, NB-IoT, LoRa and SigFox can be confusing without an understanding of their technical features, optimizations and use cases. In addition to these challenges, IoT solutions connect to their closed service layers using different messaging protocols (CoAP, MQTT, HTTP, AQMP), data-models and proprietary APIs, which make service integration difficult. This paper will provide an overview of e2e IoT network layers and make comparison of different IoT technologies in each layer with an emphasis on use case alignment. In addition to topics above, the paper will also include Service Provider e2e considerations such as security, privacy, reliability and scale. A review of harmonization efforts among standards at each layer, and a brief introduction to IoT data-model normalization efforts in the industry (e.g. Semantic Web of Things) will be covered. The reader will gain a clear understanding of the current e2e IoT technology landscape in a structured taxonomy and have a current and practical view of how to apply this understanding to their IoT decisions.

Content

1. Introduction

By 2023, there will be over 50 Billion devices connected to the Internet and much of the device growth over the next 5 years comes from Internet of Things (IoT) for consumer, enterprise and government applications. Of these use cases, the Smart Home IoT device and services market worldwide is estimated to be \$138 Billion by 2023 according to MarketsandMarkets (July 2017) and grow at a 13.61% CAGR between 2017 and 2020. The current evolution in IoT which focuses on making devices inter-connected and smarter, builds on the technology evolutions and disruptions we have witnessed in the last decennia around cloud computing and Big Data. Like those technology evolutions, the Smart Home IoT evolution will perform the same paradigm shift from closed ecosystems with lots of industry specific and fragmented standards of today, to an open and common framework to interconnect Smart Home IoT devices and services. Until this point, compliance efforts and integration issues for NSP planning to offer IoT managed services will lead to long development and deployment cycles. These means that Technicolor needs to participate in the market evolution today to gain a leadership role within our NSP customer base. Missing the initial innovation cycle could result in revenue opportunity impact in 2019/20.

1.1 Smart Home Trend

Historically, the Smart Home market started with point solutions (security, home automation etc.) that where closed platforms providing some portal or simple mobile application User experience (UX) for the

consumer to manage the service. The first evolution transition came with the introduction of NEST which was a producer of programmable, self-learning, sensor-driven, Wi-Fi-enabled thermostats (2011), smoke detectors (2013), security cameras (acquired Dropcam), and other security systems.

Table 1 : Smart Home IoT Evolution

	Before 2010	2010 - 2015	2016 – 2020	2020+ (what do we believe)
Breadth	<ul style="list-style-type: none"> Point Solution 	<ul style="list-style-type: none"> Solution Sets 	<ul style="list-style-type: none"> Broader and Multifunction 	<ul style="list-style-type: none"> Universal Multivendor Plug-n-play
Sensor Types Introduced	<ul style="list-style-type: none"> Fire/Smoke/CO Motion/Glass Break/Locks Thermostat Lighting 	<ul style="list-style-type: none"> Consumer Video Cams Smart Speakers Smart Lighting 	<ul style="list-style-type: none"> Intelligent Video Cam/Mic Smart Speaker/Voice Assist Smart Display Facial Recognition Air Quality Awareness Sound Recognition Point Function Robots 	<ul style="list-style-type: none"> Super-Sensors Augmented Reality Virtual Reality Multi-function Robots
Openness	<ul style="list-style-type: none"> Closed System 	<ul style="list-style-type: none"> Simple Open API Some multi-vendor Closed hardware 	<ul style="list-style-type: none"> Robust API Multi-vendor Automation which is Cloud-to-Cloud 	<ul style="list-style-type: none"> Open Hardware Virtual Service Orchestration
UI/UX	<ul style="list-style-type: none"> Simple UX 	<ul style="list-style-type: none"> Complex (Techie) UX 	<ul style="list-style-type: none"> Consumer UX, CUI? 	<ul style="list-style-type: none"> AI
Intelligence	<ul style="list-style-type: none"> Simple Sensor Intelligence 	<ul style="list-style-type: none"> Introduction of ML, Analytics and Voice Assist, and Cloud Management 	<ul style="list-style-type: none"> Pervasive Voice Assist Advanced ML and AI Data Analytics 	<ul style="list-style-type: none"> Predictive Cognitive
Market Delivery	<ul style="list-style-type: none"> Fragmented Network of Distributors, Reseller and VARs 	<ul style="list-style-type: none"> Emergence of OTT direct to consumer, Continued Dist./Reseller/VAR 	<ul style="list-style-type: none"> OTT direct to consumer, Emergence of Managed IoT SaaS for Service Provider Managed offerings 	<ul style="list-style-type: none"> Service Provider Managed become significant OTT direct to consumer decrease

Today, the Smart Home solution space consists largely of silo solutions but with open interfaces for multi-vendor solution to be automated. All IoT solutions for Smart Home have some cloud capability to self-provision and manage devices, often via user friendly Mobile application. Although some multi-function solutions are available, these solution vendors tend to focus on a specific solution area (Lighting, Home Security, Smart Speaker etc.). The IoT industry has started to address multi-vendor interoperability and automation of IoT devices and applications. For example, many if not all smart home IoT solutions have open published APIs for cloud-to-cloud interaction, and common connectivity interfaces to allow more integration of vendor solutions. Due to the success of voice assistant technology from Amazon Alexa and Google Assistant many IoT solution vendors are integrating Amazon and Google into their solution offering. Amazon Echo appeared on the market in 2015 and now dominates with 69% of US Smart Speaker market share in 2017 (source: voicebot.ai) and has an estimated installed base of 20+ million echo units shipped. Globally, this is a small installed base and the voice assistant market is in early market adoption phase but the estimated to grow is significant in consumer (home, car, mobile) and enterprise/commercial applications.

Going forward the emergences of Machine Learning (ML) and Artificial Intelligence (AI) has positively impacted scaling and service value in many categories such as Voice and Video recognition. Many more

innovations will emerge in the coming years as these technologies are advanced and implemented in real-world applications.

NSP's have a major opportunity to participate in the Smart Home IoT value-chain because they bring several advantages in a managed service context. Unlike over the top IoT providers, NSP's have an installed base of connected home subscribers to market and manage IoT solutions.

1.2 Network Service Provider Opportunity for Managed IoT Service

Network Service Providers (NSPs) have strategic assets, operational and business models that, if leveraged, can create key differentiation compared to Cloud IoT managed service player like Google Cloud Platform (GCP) and AWS. NSPs are in the best position to curate a cohesive managed service offering for Smart Home with multiple direct and indirect service models that can be utilized to generate revenue and value. This requires an e2e IoT framework that integrates with NSP network infrastructure, device life cycle management and BSS/OSS systems. The table in this section provides a summary of these potential differentiators.

Despite NSPs having broad capabilities, given the enormous opportunity and the incredible number and variety of competing companies, widespread success in the IoT market will not be easily achieved. NSPs face the persistent threat of disintermediation by over the top (OTT) challengers as well as increasing difficulty deploying new technologies among ageing and diverse internal systems. Technicolor faces on-going challenges as well. Popular OTT services and a general focus shift from raw internet access to multiprotocol connectivity and higher-level services has created an environment where in-home routers and gateways must evolve to remain differentiated and resist commoditization. Additionally, and as is reasonable given their size and maturity, Technicolor and NSPs both move less nimbly than many emerging companies in the space which creates time-to-market challenges that are difficult to overcome.

The Smart Home market is not without major competition for NSPs primarily coming from Cloud IoT Players (Google, Apple, Amazon etc.) which have aggressively position both home networking products (OnHub, Echo) with integrated IoT capabilities as well as Cloud IoT platforms for service delivery. In addition to players like Google and AWS which are directly competing for subscriber in the home, these vendors have also enabled new entrants because they have opened their Cloud platforms and infrastructure as a service (GCP, AWS IoT etc.). The first wave of Cloud IoT platform solutions (2016-2018) that have targeted scaling problems have allowed new entrants to build managed services that compete in traditional telco/cable market space. Cloud IoT solution providers have been able to apply intelligence to complement traditional networking protocols resulting in an augmented user experience and operational scale of these new products. Despite this development, the embedded devices in the home still poses a real challenge since they require specialized skills in embedded and real time development due to their constraints in CPU power and memory budget on top of the required domain knowledge necessary for every specific industry. NSPs have the broad set of capabilities and processes to solve this challenge, that Technicolor can help them achieve this goal.

Table 2 : NSP's Differentiation

IoT Advantage	Details
Customer Base	<ul style="list-style-type: none">• NSP have trusted customer base delivering Broadband, Video, Voice and other Home Automation Services.• NSPs have most experience with scaling networks and customers.• Have the marketing and sales channel to help consumers deal with complex choices.

IoT Advantage	Details
Connectivity	<ul style="list-style-type: none"> NSP are most experienced in connecting millions of devices, enforcing quality-of-service guarantees and offering pricing plan granularity (per device, usage based), and ability to run analytics engines in the network to manage data flows more effectively and to accelerate response time.
CPE Life Cycle Management	<ul style="list-style-type: none"> NSP are most experienced in device enablement, authentication, management, maintenance and replacement in a way that ensures network security, device directory maintenance and rights management.
Customer Management	<ul style="list-style-type: none"> NSP have major advantage over non NSP based IoT SPs when it comes to customers care, man power for truck rolls, automation system to measure and optimize customer experience, billing and service pricing infrastructure. Additionally, regulatory requirements at scale can be implemented across products and services.
Vertical Service Creation or Customization	<ul style="list-style-type: none"> NSP have key advantage in leveraging common infrastructure to address different vertical market.
Drive Standards	<ul style="list-style-type: none"> NSPs can drive vendors to implement standards Helping and supporting standardization efforts in areas like semantic web to overcome the current gaps of device centric communication standards. NSP needs to standardize on a framework which cover multiple layers and coordinate with SP community but allow device manufacturers to differentiate on algorithms and services that improve the current products in a model that still allows competition.

1.3 NSP IoT End-to-End Layers

NSP's must determine which layer of the end-to-end service they want to add-value or own. At a high-level there are three distinct layers that can be evaluated by the NSP to add value. Starting at the bottom we have **Devices**, making up the **Connectivity** layer, which deals with devices discovery, network connectivity and control (QoS, Security etc.). Above this layer is the **Platform** layer, which deals with many critical functions such as device management, service discovery, messaging data management and network and service integration into operational and business systems. Next is the **Service** layer which provides the consumer application or services being offered, and all related systems that enable the applications (AI, Machine Learning, Cloud Infrastructure etc.). Below is diagram and a table that summarizes these important layers, key function and examples.

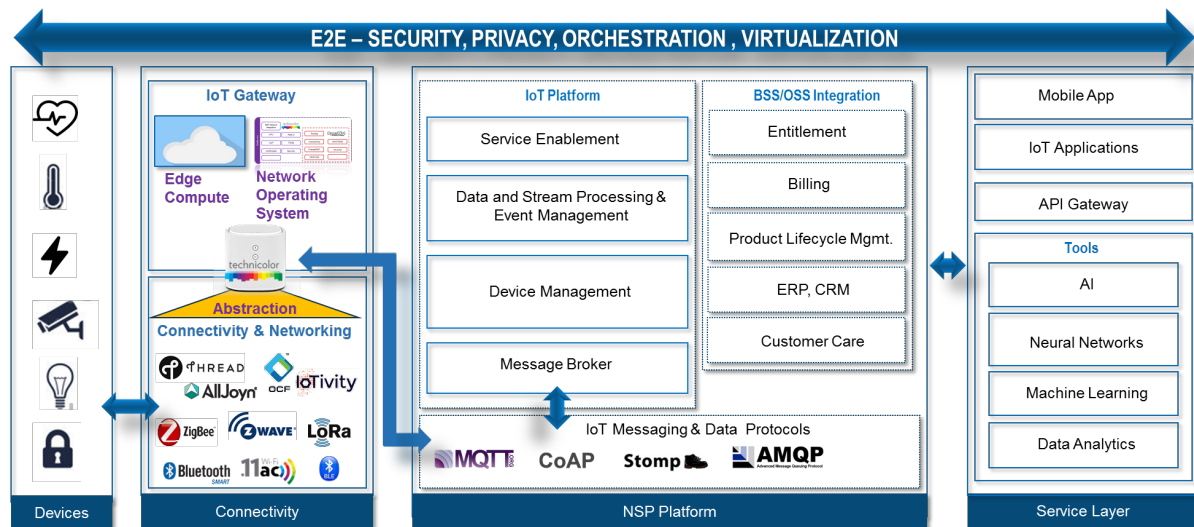


Figure 1: NSP End-to-End IoT Service Layers

1.3.1 Devices

The number and diversity of IoT sensing devices is vast and fragmented and will not be discussed in detail in this paper. However, from an NSP view, it is important to mention the complexity of managing SKUs and kits for curated IoT services will be a major challenge. There is a new opportunity to provide general-purpose sensors which eliminate the need for discrete number of sensors. Many sensors are product specific and limited in one functional area (glass break sensor, motion sensor etc.). Since many smart devices and sensors are silo managed products it makes it extremely difficult for the NSP to kit many sensors, and it is equally challenging for the consumer to manage. Additionally, many homes have devices that are not smart but provide some alerting/alarming capabilities (smoke and CO detector), and to upgrade these to Smart IoT is expensive and time consuming. Furthermore, these discrete sensors may have intelligence or awareness of other sensor devices.

Against this backdrop is a new concept for general-purpose sensors that can integrate many sensor capabilities and are low cost. These sensors can be placed in home in key locations and providing a panopticon of sensor awareness and capabilities “super sensor”. A super sensor can eliminate many SKUs for NSP, provide contextual awareness with multiple sensor inputs and make non IoT device smart (e.g. non-smart Fire Alarm). The super sensor utilizes Wi-Fi connectivity to the IoT cloud application which eliminates the need for multiple IoT radios to interconnect appliances, alarms and sensors.

A recent Carnegie-Mellon research project and paper “Synthetic Sensors: Towards General-Purpose Sensing” demonstrates that a super sensor can eliminate many sensors in the home. (Gierad Laput, 2017) In the project a super sensor integrated many sensors except video to keep the cost low. The super sensor is capable of sensing; sounds, vibrations, ambient temperature, air pressure, humidity, illumination, color, motion, magnetism, EMI and RSSI.




Figure 2 : Super Sensor (Gierad Laput, 2017)

This research showed that a single super sensor could provide a whole room awareness which would require 10s of devices in a single room. Furthermore, combinations of sensor inputs provide more accuracy in detecting events than purpose designed sensors and eliminated need for heterogeneous IoT protocols and frameworks. It is important to note this sensor does not cover video and IoT use cases such as remote door locks, doorbell etc.

1.3.2 Connectivity Layer

Today there exist a very fragmented and heterogeneous world of different competing in-home device connectivity solutions (Zigbee, Z-wave, Thread, etc.) that are all defined with different device constraints,

different wireless protocols and use cases in mind. These industry initiatives are solutions that are focusing on the effort of device manufacturers and standardization committees to make them interoperable in their own walled garden (a.m. Silo). The impact on the market viability for both mass market and NSP trying to package an IoT offering are very challenging. Imagine today a home with 20 smart devices with some joined to third-party IoT hubs and all these being connected to each vendor's cloud service, each with its own proprietary API. These silos are then linked to some other cloud service used by a controller like an Amazon Echo, Google Home, or smart home app. One can imagine the challenge for both consumer and NSP to manage this complexity.

The connectivity layer consists of IoT device connectivity and networking of IoT devices in the Smart Home. The IoT connectivity layer has universally moved to wireless technology. To an end user these are invisible network connections to their IoT devices in the home, and they don't want to be troubled with installing or diagnosing these connections. Many established and emerging wireless technologies are available in the market, each addressing some combination of optimization (power, performance, range, bandwidth, latency and cost). The leading established wireless networking standards for IoT are Wi-Fi, Bluetooth, Zigbee and Z-wave however, newer generation of technologies such as 6lowpan, WeMo and Thread exist. A comparison of IoT wireless connectivity standards should be viewed based on use case requirements such as power consumption, throughput, latency and range. The following diagram shows three distinct groupings based on Data Rate, Range and Power consumption.

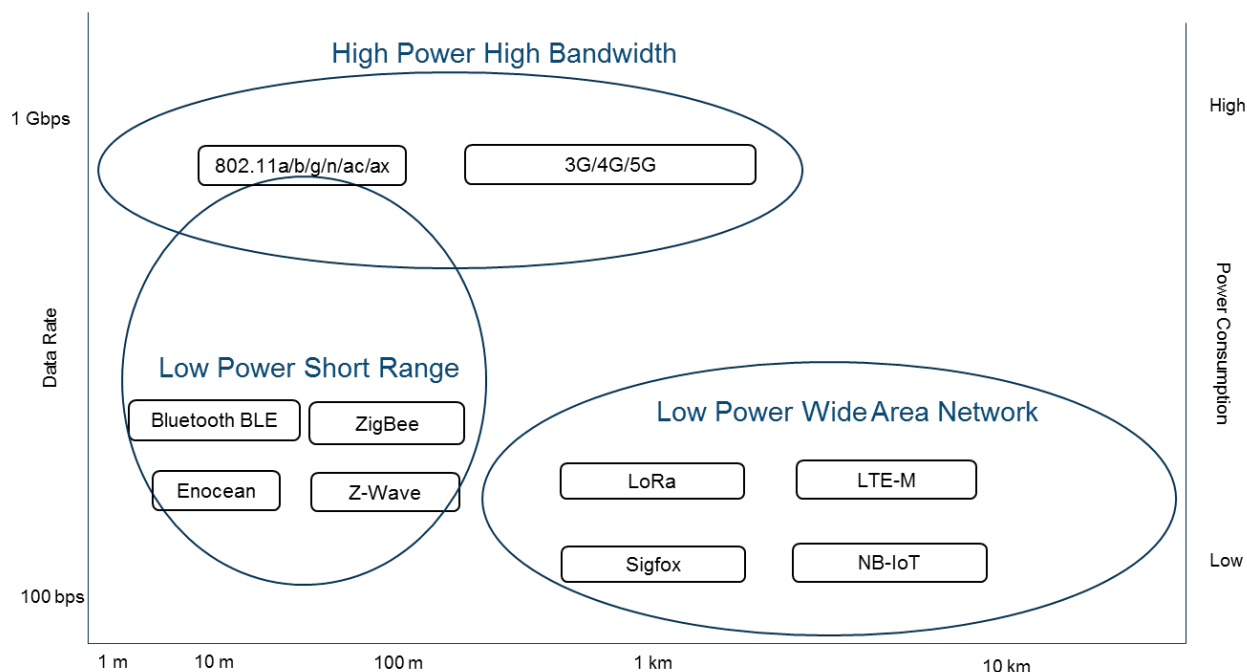


Figure 3: Wireless Connectivity for IoT

Wireless connectivity technologies can be grouped into 3 general segments. There is the low power short range technologies such as Bluetooth, ZigBee and Z-Wave. There is the lower power long-range or wide area network technologies like LoRa, Sigfox and NB-IoT. Lastly there are the high power wireless broadband protocols such as WiFi and 4G/5G, although WiFi is more of a short range wireless broadband solution.

While the new generation of transport protocols are achieving ipv6 using 6LoWPan (Thread, BLE) and are transport agnostic, most of them don't provide a clean separation between connectivity and application

specific functionality and hence need complex transformations between them to make them interoperable. The myriad of wire protocols, data formats and data models have given rise the current complex landscape in IoT connectivity. The complexity is exacerbated by the fact that many vendors create walled gardens with their cloud support functions to operate these solutions.

Smart Home IoT connectivity is mostly concerned with range of 100m or less and depending on the use case low power or high bandwidth. Therefore, the predominate technologies are Short Range. A more detailed comparison of short range connectivity options is provided in the following chart. It is important to note that range of different technologies can depend on indoor, outdoor obstructions and some are subject to interference. For example, technologies using 2.4GHz frequency band could be subject to significant radio performance degradation to do interference. Equally variable is the power consumption comparisons since different IoT devices and or use case may utilize power than others. Therefore, this chart should be view as a general range and power consumptions numbers.

Table 3: Short Range IoT Connectivity Comparison

Feature	Wi-Fi	Bluetooth	ZigBee (Alliance, n.d.)	Z-wave	EnOcean	Thread	6LoWPan
Open	Yes IEEE 802.11	IEEE 802.15.1	Yes IEEE 802.15.4	No, Proprietary based on IEEE 802.15.4	Yes, EnOcean Alliance.	Uses IEEE 802.15.4, IETF 6LoWPan	Yes, IETF RFC 6282 , uses 802.15.4
Range	100-150 feet	v4 = 300 feet v5 = 600 feet	30-100 feet	50-100 feet	100-300 feet	100 feet	100 feet
Frequency	2.4/5 GHz	2.4GHz	2.4GHz	908/915 MHz	315, 868 MHz and 2.4 GHz	2.4 GHz	2.4 GHz
Data Rate	300 – 1300 Mb/s	v4 = 1 Mbps v5 = 2 Mbps	40-250 Kbps	9.6-100 Kbps	125 Kbps	250 kbps	250 kbps
No. Devices	Router dependent	7	65,000	232		250-300	250-300
Topology	Star	P-to-P Mesh	Mesh	Mesh		Mesh	Mesh
Hub Required	No	Yes	Yes	Yes		Yes	No
Security	WPA2	AES-CMAC encryption ECDHE (Elliptic Curve Diffie-Hellman)	AES-128 symmetric encryption	AES-128 symmetric encryption	AES-CBC and variable AES (VAES)		
Power Consumption	High Power	Low Power	Low Power	Low Power	No Power	Low Power	
Cost							
Good For	High bandwidth Mid-range	Short Range PAN and LAN	Lower Power Short Range	Low Power Medium Range			

It is evident in this evolutionary phase of IoT that the primary focus of solutions and technology are more driven by IoT use case and user experience versus universal interoperability. Therefore, several connectivity technologies will exist for the foreseeable future. The NSP will need to make key decisions on which combination of connectivity technologies they need to deploy if they want to add value in this layer. Another factor is cost of technology that impact device costs to consumer or network related costs.

Market penetration of various wireless connectivity technologies among sensor and appliance equipment providers varies significantly. Wi-Fi and Bluetooth dominate home due to their maturity and pervasiveness, others depend on maturity and vendor solution support. It is hard to find good analysis of which wireless technologies dominate the Smart Home beyond Wi-Fi and Bluetooth, but Z-wave appears to have the most significant ecosystem support followed by ZigBee, and then a large gap in terms of adoption exist with new generation technologies such as Thread, NB-IoT and so on. We believe that Wi-Fi will dominate applications that do not require low power or some other constraint. In the low power solution area there is a lot of competition with Z-wave having a largest ecosystem. It is worth noting that Z-wave which was developed by Sigma Designs has been acquired by Silicon Labs ([here](#)). This could widen the adoption of Z-wave since Silicon Labs is a major player in IoT chips. More than 2,400 certified, interoperable Z-Wave devices are available from the Z-Wave Alliance of more than 700 manufacturers and service providers worldwide.

The official Bluetooth marketing material from the Bluetooth standard organization advertises that Bluetooth 5.0 has four times the range, two times the speed, and eight times the broadcasting message capacity of older versions of Bluetooth. Again, these improvements apply to Bluetooth Low Energy, ensuring devices can take advantage of them while saving power.

With Bluetooth 5.0, devices can use data transfer speeds of up to 2 Mbps, which is double what Bluetooth 4.2 supports. Devices can also communicate over distances of up to 800 feet (or 240 meters), which is four times the 200 feet (or 60 meters) allowed by Bluetooth 4.2. However, walls and other obstacles will weaken the signal, as they do with Wi-Fi.

Looking at their functionalities, following picture shows the overlap of these connectivity stacks when applied to OSI model.

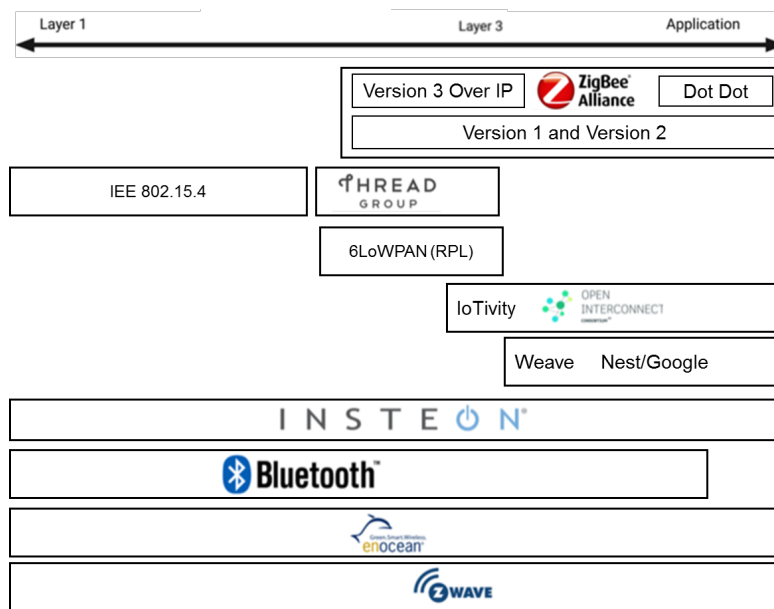


Figure 4: Connectivity Stack Comparison

Albeit these connectivity solutions will continue to exist and evolve within their silo, there is a clear need to come up with an “abstraction layer” that is protocol agnostic and provides an industry standard Data Model to enable those silos to interact with each other, and this is analogous to traditional IT normalization initiative in data models and applications. Having a normalize and coherent model to allow the creation of IoT applications that are agnostic to the underlying protocols and frameworks within the Smart Home. Like the efforts around “Semantic Web & Technologies” standardization in the last few years, IoT could leverage similar concepts to further improve its capacity to understand things' data and facilitate their interoperability.

There is a clear trend of moving the network layer to an all IP layer based on 6LowPan. Thread is a good example of that and device constraints of the latest version of Zigbee followed that trend with Zigbee/IP. This opens possibilities for easier integration with this layer in container technology, since the network stack is typically embedded in the firmware seen the close relation with the hardware drivers.

1.3.2.1 Messaging Protocols

There are many messaging protocols standards each with advantages and disadvantages. There is no panacea protocol to address a universal solution for IoT, however we are seeing some of these protocols dominate specific use cases. The IoT Platform will need to flexibly support multiple protocols to address many use cases. Below is a comparison of the messaging protocols being used by IoT solutions providers.

Messaging protocols are unique from other protocols since IoT device are resource constrained, are data centric, always on and have security considerations like getting through firewalls. Messaging protocols differ on several criteria such as communication model, message size, syntax and QoS mechanisms. The Hyper Text Transport Protocol (HTTP) which the most widely use internet protocol is a Request/Response protocol versus a more model data centric protocol such as Message Queue Telemetry Transport Protocol (MQTT). (R. Fielding, 2014) More recent introductions of HTTP-like protocols such as Constrained Application Protocol (CoAP) constrained nodes and networks are still request response but address size.

Table 4 : Messaging Protocols Comparison

	MQTT (OASIS, 2015)	HTTP/S (R. Fielding, 2014)	CoAP (C. Bormann Universitaet Bremen TZI, 2018) (Z. Shelby of ARM, 2014)	AQMP	XMPP	STOMP
Standard	IISO/OASIS	IETF	IETF rfc8323	IETF rcf2119	IETF rfc6120	
IP Type	TCP-based	TCP-based	UDP-based	TCP-based	TCP-based	TCP-based
Message Type	Publish/Subscribe	Request/Response	Request/Response	Transactional	Transactional/PubSub	Request/Receipt
Syntax	Simple Noun/Verb?	Verbs/Status Codes	HTTP Like		HTTP like	HTTP like
Size	Small: 2 Bytes	Large: ASCII	ASCII	8 Byte Header, Variable Ext Header and Variable Frame Body.	Verbose, XML	1KB - 10KB
QoS	3 levels	No mechanism	Confirmable requests			No
Reliability	Avoids packet loss on client disconnect via keep alive	No mechanism	“Observer”, “Response back”			No
Security	SSL/TLS, user/password in connect message	SSL/TLS	Datagram TLS	TLS/SASL	SSL/TLS	SSL
Other	MQTT 5 - Enhancements		rfc7252 Resource discovery			
Good For	Cloud Scale & Small Footprint	Non-event based applications	Improves simplicity of HTTP		Messaging, presence detection, signaling plane	

Conclusion there is no solution that fits all and thus as an NSP must deal with heterogeneous network of IoT protocols. However, effort should be made to minimize the number of protocols in the network.

1.3.2.2 Message Broker/Gateway Function

A key function is the message broker/gateway in the platform layer for the Service Provider. This function needs to support multi-protocol messaging interfaces, message load balancing, high-availability and management. A few open source message brokers are available in the market (RabbitMQ) but the Carrier Grade requirements for this critical function will require optimized and hardened platforms.

Messaging infrastructure which typically provides publish/subscribe semantics have been used for a long time and complement the more traditional Request/Response semantics of the current internet

communication paradigm. They have been used to decouple monolithic solutions towards a more decouple system that is easier to evolve. Additionally, these networks are always on and require real-time capabilities to enable a more event driven architecture.

With the advent of the cloud a new generation of cloud-native solutions have emerged for IoT which hence are more capable of scaling to a global infrastructure, while still providing resilience, High availability and guaranteed message delivery.

These new cloud-based solutions have followed the same evolution as their database counterparts moving from SQL to NoSQL solutions, but are therefore also lacking some maturity and have been simplified in terms of requirements in favor of their scalability requirements.

Typically, IoT platforms support different transport protocols like MQTT, Websockets, or proprietary protocols, which are chosen because there are lightweight and hence can easily be integrated in resource constrained devices. Also, the CoAP protocol starts to follow this trend with the recent extension of the protocol towards pub/sub capabilities, but this is still in a very early stage.

Today's solutions are using this infrastructure layer as a control plane for use cases such as command/control, notifications at large scale, configuration management, presence detection and even for the signaling plane for communication protocols like WebRTC. They typically also providing the necessary adaptors to other systems as queuing systems, streaming systems and rule engines.

Today all major cloud providers are offering this messaging services as part of their IoT offering, with the caveat that they provide out of the box integration with their own solutions and therefore create a risk of locking in their ecosystem. Some other companies have focused on providing alternative solutions in this space, even with capabilities of running the solution on premise. PubNub is a notable example of that evolution.

In short, careful consideration need to be made in the selection of these technology, with respect of locking-in, cost, global availability, size of the ecosystem and developer's community, protocols supported, adaptors, and the messaging semantics they provide.

1.3.2.3 IoT Device Management

An IoT Platform is the lowest layer at which the IoT devices connected to the system can be viewed and managed on a system-wide basis. This makes the platform the ideal place to manage those devices.

Device management activities can be classified into three distinct groups:

- Hardware Management
 - Inventory
- Software and Configuration Management
 - Device Modeling
 - Authentication of cloud/backends
 - Registration
 - Entitlement
 - Software upgrades, OTA updates
 - Configuration management: determination, verification, backup, reset
 - Policy creation and application
 - Off-boarding and device retirement
- Monitoring

- Centralized log collection and management
- Fault tolerance, failing safely
- Issue alerting
- Troubleshooting, diagnostics and remote reboot

Much of IoT Device Management overlaps with broadband device management that NSPs excel at. There are two areas in which IoT Device Management may expand on existing broadband device management functions in an NSP:

- **Scale:** It is likely the number of connected devices and sensors across the customer base has or soon will exceed the number of broadband gateways in the NSP network. This will require additional scaling and automation beyond what exists today.
- **Device Modeling and Offline Representations:** The requirement to support smaller embedded devices, devices with lower power, and intermittent internet connectivity will force IoT management platforms to seamlessly manage devices whether those devices are online at the time of a change or management action. Leading platforms are supporting this through a cloud-resident abstraction of a remote physical device. These abstractions are commonly called device shadows or device twins, which allow the IoT device management platform to execute changes and actions that are cached centrally until the device becomes online.

1.3.2.4 BSS/OSS Integration

Most NSPs have OSS/BSS systems that have evolved over decades. In such a brownfield environment, it is critical that any new service introduced by the NSP (e.g. IOT service in this case), integrates seamlessly into the NSPs current OSS/BSS systems.

For OSS integration, NSPs have mediation platforms that act as brokers for OSS integration. The mediation platforms map the operational workflows via a standardized interface within the NSP domain to provide configuration and provisioning of customer data, operational parameters /bounds for the service, health monitoring and analytics associated with the service. Typical integration technologies used are SOAP/XML, RMI/RPC ORB, EAI/CORBA.

For BSS integration, NSPs have customer management systems that provide billing & charging, NSP CRM systems (Product Catalog, Order Management, service order fulfillment etc.) as well as end user presentation systems e.g. customer account management portal, service management portal as well as a set of associated NSP branded mobile apps. Like OSS integration, the BSS integration on the backend is also done via mediation platforms that use integration technologies like SOAP/XML and RPC. However, the customer presentation systems are usually bespoke to the service being offered e.g. a set of mobile apps dedicated to the service to provide customized UX for the service.

1.3.3 Service Layer

The service layer deals with the end user application and service offering that is consumed by the end user. It deals with all the service logic, intelligence (ML, AI, Analytics) application interaction and services experience by the end consumer. The service layer deals with presentation layer (portal or mobile app) to the end customer and manages all the service logic for the end user or IoT devices. This layer must integrate with service activation, billing and customer care BSS systems in the NSP network.

Service or application execution environment have been virtualized and thus can run in Cloud in a centralized model or distributed to Edge Compute.

1.3.3.1 Edge Compute

Many challenges have emerged as the number of IoT solutions in the Smart Home market grow. Privacy, latency, bandwidth constraints, and reliability, among others, present challenges that cannot easily be overcome in cloud-only models.

Edge compute is a term that generally refers to the ability to perform enhanced or additional processing in the CPE. This processing uses higher-performance CPUs and additional RAM and may utilize virtualization or containerization technologies or may take place directly on the existing device operating system. Because processing via edge compute takes place in close-proximity and well-connected to the consumer, it is an attractive and useful tool to combat the challenges created by many Smart Home service offerings. For example, processing data in the consumer's home without sending data to the cloud is more private and less reliant on fast and reliable internet connectivity.

Privacy – which is only one of the concerns reduced by successfully leveraging edge compute – has become a major concern for consumers in IoT markets like US and Europe. Regionally, European markets have instituted greater regulatory protections over privacy than the US. This is best exemplified by the implementation of the General Data Protection Regulation (GDPR) which takes effect May 2018. The GDPR significantly changes how companies handle EU citizen data privacy. GDPR places requirements for managing EU citizen data, such as a 72-hour notice to citizens after first having become aware of a data breach. Other aspects of the regulation require data erasure, data portability, and reporting. In general, personal data about identity, habits, speech and video will become a growing concern. Edge compute can enable new service architectures where personal data is processed locally to minimize the exposure of personal data.

Edge compute can also maintain service operation during a network failure. On-device, service-specific processing that is enabled by edge compute can act as a buffer during a network failure and then synchronize data and state with cloud-based processes when internet connectivity is restored. Some sensor applications have additional redundancy requirements that may include using a battery backup to operate during a power outage. Edge compute, coupled with a battery backup, creates a robust platform for offline data processing.

Edge compute is also a useful tool to service providers looking to reduce operating costs by shifting processing away from expensive cloud providers. Utilizing edge compute to relocate data processing from data centers into consumer's homes decreases the cost associated with cloud provider compute resource consumption.

Over the last two years we have witnessed the shift towards edge computing to complement the first generations of cloud centralized IoT solutions, mainly for cost and latency sensitive solutions and allowing for autonomous intelligence at the edge in absence of internet connectivity. The concept of intelligence at the edge is not new. Similar concepts include mobile cloud computing (MCC), mobile edge computing (MEC), mist computing, and cloudlets (fog nodes). Although the edge ranges from private cloud, to NSP core network and the to endpoints, they contain similar architectural building blocks that designed for their respective resource environments. In the context of Edge Compute for IoT there will not be one solution fits all, but the true challenge is to provide a Network Operating System that can support an application execution environment to run IoT application and services on the Edge of the

Network. Fog compute can be confused with edge computing, but they are different. For example, Fog Node may be a service that runs on Edge Compute on the Homeware CPE.

One of the most important elements of this shift is the focus on offering a developer friendly environment for NSP that enables self-service and e2e control of a software workloads using container technologies. Containerization enables a leap forward in productivity with a modern workflow, and cloud-like agility to offer a new service, collecting user feedback and a fail fast attitude that eventually will disrupt the current status quo w.r.t. Embedded development.

1.3.3.2 Containers for Embedded Implementations

The introduction of containers in the open source community have their origin in Linux development projects. Linux began providing containers (LXC) in Release 4.x and OpenWrt projects began around 2015 (Rel. 15.x) for both LXC and Docker. Docker is a superset of LXC and will be discussed below in more detail. Containers are self-contained execution environment with their own, isolated CPU, memory, block I/O, and network resource which are share the kernel of the host operating system. This is different from virtual machines which are running many duplicate instances of the same OS and thus heavier weight. Another form of virtualization is process containers a.k.a serverless programming.

In this case a runtime binding is mapped into a container that offers an execution environment that contains the basic set of libraries necessary for a specific language (Python, JavaScript, Java, Go, etc). This offers a more lightweight option w.r.t. the container size, but creates some additional language specific dependency management problems. The key advantage here is that the need to manage all the security aspects and operational burden is taken away from the internal container environment and hence the name serverless was coined.

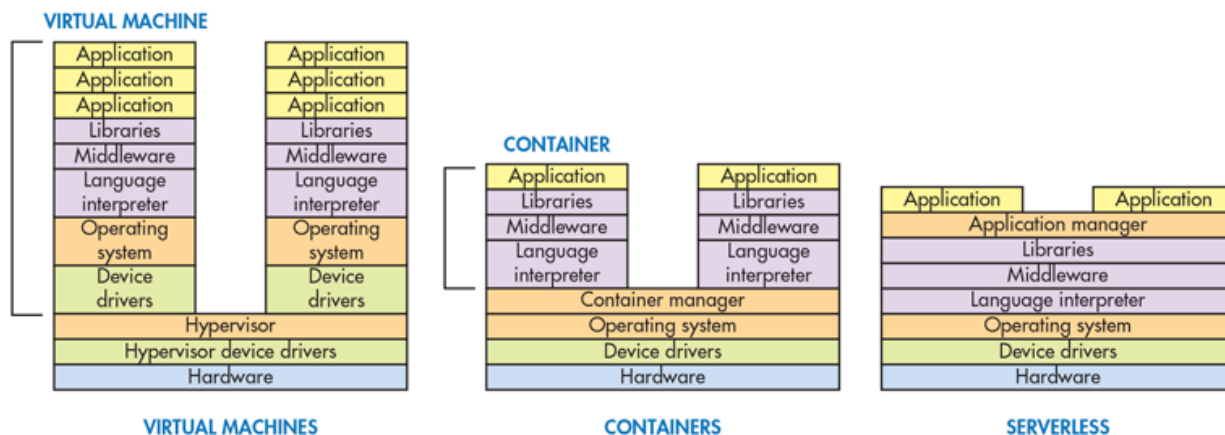


Figure 5: Virtualization of IOT services

The Linux Container (LXC) features and capabilities continue to evolve for embedded systems with limited resources and different IoT use cases. Containerization and light weight sandboxing tools are available in open source to develop a framework for secure application execution environments. This is particularly important in residential or connected home CPE, Wi-Fi APs and Extenders offered by service providers.

Orchestration and Portability of containerized applications is an important design requirement for Service Providers to move applications across different device hardware and resource capabilities. Running containerized or serverless application on customer premise equipment that the NSP curates will be a

challenge but represents and innovation that may allow them to have better service delivery than OTT solutions competing for the same consumer in the Smart Home market.

1.3.3.3 API Gateway Function

As discussed in earlier sections, many (if not all) Smart Home IoT solutions have leveraged the Application Programming Interfaces (API) based on Representational State Transfer (REST) style resource-oriented architecture (ROA). REST APIs allows application integration between IoT product silos and/or 3rd part application development and automation between these products or to integrate into other systems (eg. (backend DB, Analytics, BSS/OSS etc.). In addition to REST API which leverages HTTP protocol, an increasing number of IoT solution can support IoT protocols (MQTT, etc.) over Websockets. This shows that APIs are critical aspect of the Service Layer and NSP need to understand key aspects of API Management, Security, Privacy and Consumer Experience.

1.3.3.4 IOT visualization

Due to large number of datasets available with IOT application, the task of aggregating and rendering the datasets to the end user in an intuitive way is key for good user experience (UX). As an example, the figure below shows a smart farm IOT application that has LoRA based wireless IOT sensors that gather all the farm sensor data and aggregate it in a useful way that it can be layered upon a real time video feed and rendered to the end user. The screenshot below shows the temperature and humidity in the grain silos, soil humidity as well as health related information on the farm animals.




Figure 6 : IOT smart farm Augmented Reality

1.4 IoT Architecture Approaches

IoT architectures have traditionally been public or private cloud based, where all the IoT devices sent actionable telemetry directly to a virtualized cloud IoT gateway or via a physical IoT gateway on-prem. IoT gateways aggregate all the IoT events and feed it to an event processor which would correlate the events to actions and triggers. Triggers might be to send dynamic control messages to IoT devices and/or to notify the end subscriber of the events. Such a traditional architecture is shown below:

Tenets of a such an architecture as shown above are:

- IoT devices at premise

- Broadband gateways at premise that support Wi-Fi and other IoT specific low power 802.15 radios (Zigbee, BLE, Z-Wave, etc.)
- Virtualized cloud gateway (Public or Private Cloud)
- Stream Event processors: To process all the incoming IoT telemetry
- Scalable databases: To manage IoT device identities and states.
- Control systems: To analyze incoming IoT telemetry, map it business logic, and send control signals to actuators in IoT devices if needed
- Analytics: Create Actionable insights from IoT data and state for OSS and/or end user presentation.
- Subscriber Portal / Mobile UX: Cloud based portals and mobile apps to provide a GUI to end user to interact with the subscribed IoT services and manage notifications

However, in the last 18-24 months, users have been lot more concerned about privacy and security related to IoT services. Interaction with NSPs echoes those concerns as well. In recent discussions with a leading US MSO, this topic was front and center for their IoT platform requirements. This has led to evolution of a hybrid architecture that allows for user data to kept private and on premise while doing other non-sensitive data processing in the cloud.

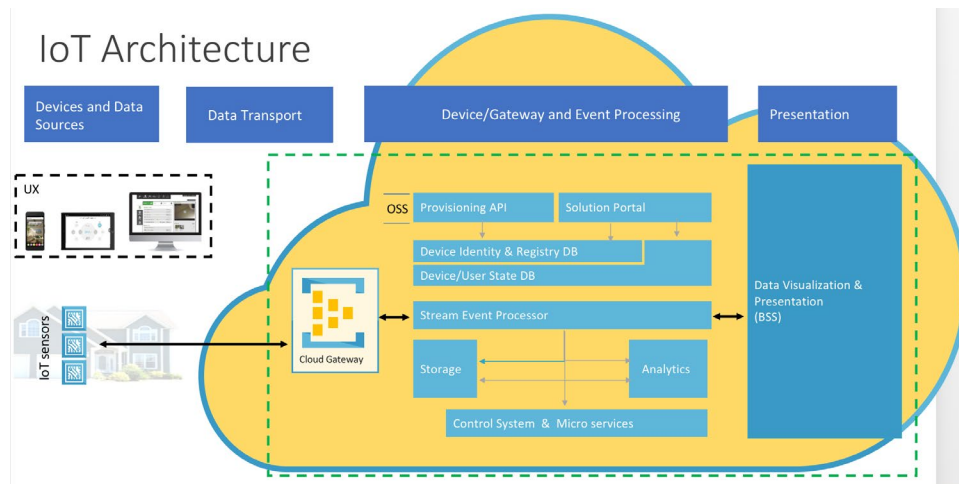


Figure 7: IoT Platform Architecture – Cloud Approach

Such an architecture uses a Broadband/IoT Gateway device at premise that has edge-compute capabilities to store privacy sensitive user data (telemetry) as well as be able to do local processing of such IoT telemetry. Once local processing is performed, only anonymized triggers are sent to the cloud backend for further processing. Based on the intent of the service, IoT user data may or may not be sent to the cloud. As an example, a 'peace of mind' security service could have business logic where a motion sensor at home triggers a video capture device (IP camera) on motion detection. Then the video clip is sent to the user's mobile app as a notification; for user privacy the captured video clip is sent directly from the edge-cloud capable IoT gateway to user's mobile app using a secure IP transport from Hybrid GW to the mobile app. Video clips are never stored in the cloud backend thus mitigating any privacy related fears (as well as addressing regulation like GDPR) related to cloud based storage.

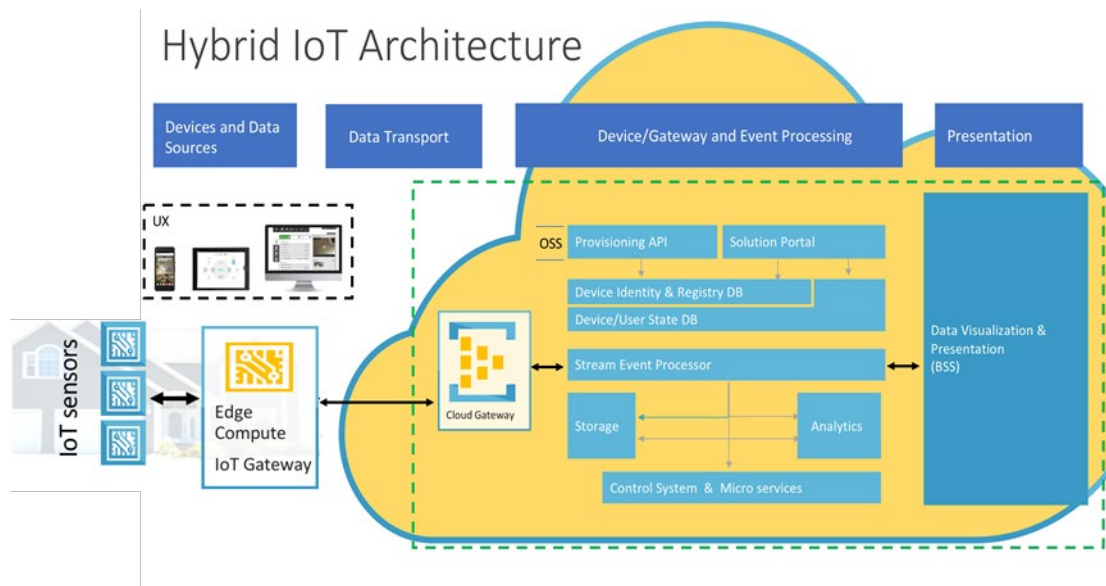


Figure 8: IoT Platform Architecture – Hybrid Approach

As shown above, the hybrid architecture utilizes an edge-compute capable Gateway on premise that is capable of localized processing of IoT events, thereby keeping user data and IoT telemetry local and private on premise. Typical Edge-Compute capable Gateways have Dual core or Quad core low power processors as well as local storage (Flash and RAM) that is more than traditional broadband gateways.

1.4.1.1 IoT Gateway Function

IoT hubs such as Logitech Harmony Hub, Samsung SmartThing Hub and others are focusing on integration at the connectivity layer to provide customers a "universal remote" experience. In addition to the capabilities of bridging a myriad of connectivity protocols, the industry has recognized the importance of a local execution environment on these Hub's referred to as Edge Compute. Edge Compute addresses several challenges related to latency, network bandwidth, reliability and security, which cannot be addressed in cloud-only models. Containerized execution environment running on Edge Compute provides agility in rapid development and deployment of new capabilities, driving cloud cost down and providing more autonomous operations of these devices without having to rely constantly on the availability of cloud services. Although this was the initial promise of OSGI, containerization and virtualization are providing a more flexible and secure alternative, and in this respect, we see quite some movement in this space with offerings from Amazon, Microsoft (Lambda's and Functions) or fully dockerized environments like what resin.io is providing.

Table 5: IoT Gateway Stack

IoT Gateway Functions	
Messaging & Data Management (MQTT, HTTP etc.,)	Network Management
Connectivity & Inter-Networking (ZigBee, Z-wave, Wi-Fi)	
Edge Compute (LXC, Docker, Serverless Programming)	
OS/RTOS (Linux, .NET)	

Being on the demarcation point between the NSP's network and the home network, we believe that a residential GW is at an ideal place to tap into Edge Compute. Edge Compute processing giving NSP's and end-to-end capability to balance network functionality and in-home service and being able to finally break the current IoT siloed solution challenge. It gives them the opportunity to improve security, home networking and fundamentally change the user experience in an agile way. For more details on Edge Compute go to section 2.2.8 Service Layer.

IoT gateways that are curated by the NSP can provide a single point of Wireless IoT connectivity interfaces, device and data management and edge compute. With edge compute resources core network functions such as the API Gateway could have a local instance and sync with core network when needed. This would address important privacy and resiliency requirements for IoT services. We have already seen the emergence of IoT API layer moved into the local devices execution environment. For example, Amazon allows you to develop application in the AWS Greengrass environment but allow it to execute local. Azure and Google are all moving to this local compute environment.

1.4.2 Platform Layer

The Platform layer deals with all the important management layers to provision, network, scale and manage IoT services as an integrated service offering. NSPs contemplating ownership of all or pieces of this layer require careful platform design of key integration points into existing provisioning, operational and business systems. The introduction of new a IoT device management framework and potential integration will be core function. Across these layers NSPs will have to consider how each layer meets security and privacy requirements, and how to facilitate by network infrastructure at scale. At the top of the stack important consideration of how the platform interacts with services and applications that are 3rd party or organic needs to be determined. The complexity of the Platform layer requires NSP's to look at vendor solutions that address these function areas, and different deployment models. The different layers of the IoT Platform from an end-to-end perspective is presented below and detailed in following sections.

Table 6: IoT platform layers

		IoT Platform Functions
Security & Privacy	Infra. Compute, Storage, Processing	Service Enablement Set of well-defined APIs, interfaces and tools between management layer and services layer.
		BSS/OSS Integration Run business logic for IoT service, operational integration, visualization, customer care etc.
		Entitlement Manage different services and policies including billing, bundles etc. This layer interacts with service catalogue and other BSS functions.
		Data Processing & Event Management Process data real-time via rule engine, events and notifications
		IoT Device Management Device management (on-boarding, monitoring, updating, replacing), Backup.
		IoT Messaging Management Message Broker/Gateway, Messaging, Queueing, Security...

As a point of reference, according to IoT Analytics ([here](#)) there are +450 IoT platforms on the market, of which 32% focus on industrial applications. IoT Analytics also states that “We believe, only 7% of the 450 IoT Platform companies generated revenues more than \$10M with their IoT Platforms in 2016. Furthermore, more than half of all companies made less than \$1M, most of them smaller startups. The firms leading the pack are mainly made up of large cloud players, legacy device management and connectivity backend platforms as well as a handful of heavily backed Silicon Valley startups that are scaling faster than most of their counterparts around the world.”

1.4.2.1 Messaging Management

This layer interacts with IoT devices and is typically referred to as a Messaging Broker or Gateway. The messaging must be securely transported from IoT device to the management layer, must be scalable messaging protocol to support millions of devices, must support maintain resiliency and integrity of the

messages, and must intelligently distribute message to the correct processes in the management core layer.

1.4.3 Cloud IoT Platform

Amazon, Microsoft, Google and IBM all have IoT infrastructure and services that are similar but with some key differences. They all are essentially offering their Platform-as-a-Service. For the most part these vendors focus on infrastructure and service to facilitate IoT applications and leverage other cloud services (compute, storage, analytics etc.) they offer. They are largely cloud-based platforms, but they do provide device level components on premise with local processing (e.g. AWS Greengrass, Azure IoT Edge). AWS and Google are most aggressive in enabling IoT platforms to complement their smart home devices push (Google Home, Alexa Dot), but AWS and Azure are more mature than Google. IBM is more focus on Artificial Intelligence (AI) as a service, and Microsoft Azure has refocused its efforts on telemetry and data collection of IoT devices. Despite the big 4 developing broad end-to-end approach and aggressive roadmaps, the market and technology is still in early phase of maturity. Both Amazon and Microsoft did not release their IoT platforms until late 2015, and Google's was made available in 2017 and is still in Beta.

Table 7: Infrastructure-as-a-Service Cloud IoT Frameworks

	Edge Products	IoT Core Product	IoT Core Services	Pricing (US)
AWS	IoT Device SDK Greengrass – Local Lambda (edge compute) Snowball – Local storage	IoT Core	Connectivity: Message Broker Device Management: State - Device Shadow Registry Rules Engine IoT Analytics -- filters, transforms, add meta-data	Connectivity: <ul style="list-style-type: none"> \$0.080 per million minutes of connections Messages monthly message volume: <ul style="list-style-type: none"> Up to 1 billion messages \$1.00 Next 4 billion messages \$0.80 Over 5 billion messages \$0.70 Device Shadow and Registry: <ul style="list-style-type: none"> \$1.25 per million operations Rules Engine: <ul style="list-style-type: none"> Rules per million triggered \$0.15 Actions per million executed \$0.15 *AWS Free Tiers (see AWS details below)
Azure	IoT Device SDK IoT Edge – local processing of Azure modules (edge compute)	IoT Hub	Connectivity: Message Broker Device Management: State - Device Twin Registry Provisioning Monitoring - IoT Suite Maintenance – IoT Suite	IoT Hub (all messages metered in 4KB blocks, max message 256KB): <ul style="list-style-type: none"> S1 tier: 400,000 messages per day per IoT Hub \$50 /mo. S2 tier: 6 Million messages per day per IoT Hub \$500/mo. S3 tier: 300 Million messages per day per IoT Hub \$5,000/mo. IoT Device Provisioning: <ul style="list-style-type: none"> S1 tier: General Availability Price: \$0.10 per 1,000 operations
Google	Google Cloud MQTT Client Opt. Brillo/Weave*	Cloud IoT Core	Connectivity: MQTT/HTTP Bridge Device Management	Don't allow flexibility with MQTT to address scale.
*Brillo became Android things, and weave got abandoned in favor of Nest weave which is a total different protocol. **				

A functional and feature comparison of the big 3 IoT frameworks is provided below. As you can see in the table AWS has the most comprehensive IoT offering, but there are many similarities between AWS, Azure and Google.

Table 8: Cloud IoT Framework Comparison

	AWS	Azure	Google
Client SDK / Language	Android-Mobile, Arduino, Embedded C, C++ , iOS-Mobile, Java, JavaScript, Python	.NET, Embedded C, Java, Node.js, Python	GCP - Java, Python, NodeJS, Ruby, Go, .NET, and PHP
Messaging Protocols	MQTT, HTTP, WebSockets	MQTT, AMQP, HTTP	MQTT and HTTP 1.1 (not 2.0)
Security Transport	TLS	TLS	TLS 1.2
Authentication	Per-device with SAS token	X.509 certificate client authentication, IAM Service, Cognito Service	Per-device public/private key (asymmetric) device authentication and JSON Web Tokens (JWTs RFC 7519)
Device Management	Registration Configuration State	Registration Provisioning State Monitoring & Maintenance	Registration Provisioning State Monitoring
Edge Compute & Services	Greengrass -	IoT Edge – Stream Analytics, Machine Learning, Azure Functions (custom code)	No
Data Ingestion & Processing	Kinesis	Event Hub	Cloud Pub/Sub
Stream Event Processing	Kinesis analytics	Stream Analytics	Cloud Data Flow
Data Storage (DB)	S3 DynamoDB RDS	Azure Blob Storage Azure Cosmos DB Azure SQL DB	Cloud Storage (object store) Cloud Bigtable BigQuery
Data Visualization	QuickSight	PowerBI	Cloud Datalab/Data Studio
Analytics	IoT Analytics	HDInsight	Cloud Analytics
Machine Learning	Sagemaker	Azure ML	CloudML
Notifications & Alerts	SNS	Azure Notification Hubs	Firebase Cloud Messagin

1.5 Harmonization of Standards

1.5.1 Connectivity Harmonization

Currently there are different opinions on how to tackle the connectivity challenge. Some are advocating, open data models, or a new generic opensource protocol, while we see also movements from vendors like

Amazon that is promoting de-facto API's like what happened in the cloud space for most of these services.

At connectivity layer there are several alliance/stands initiative to create common standards. Examples Alljoyn and OCF, you have half the industry standards. Wireless standards are also coordinating common standards. OCF/Alljoyn is focused on the lower layer abstraction...

At layer 3 we see consolidation around IP vs. proprietary.

At the application layer, silos remain a main issue among alliance groups. There remains a gap in the standardization of the application layer. For example, when Alljoyn was absorbed by IoTivity, They were directly competing with the established standards like Zigbee, Zwave and bluetooth, and despite all the efforts they have still not established a major footprint of compatible devices. Therefore, an abstraction layer will still be required in the architecture to stitch these APIs together.

1.5.2 Data Model Harmonization

The IPSO Alliance which is now part of Open Mobile Alliance (OMA) has been one of the first standardization committee's, identifying the problem around interoperability issues with the different connectivity protocols and has originally started to create a cross industry attempt to harmonize the different data models in different industries into a more generic set of data models.

The IPSO Alliance is actively developing an entirely new approach to resolve this data representation and scalability issue. They call it the Node Metal Model. It defines a unique method that allows smart objects to interoperate with each other.

This new meta model is the only known approach that universally sets out how all things should be defined, so that each specific thing, including its objects and resources, no longer needs to be predefined and preregistered.

1.5.3 Service Layer Harmonization

At the service layer, there are standardization and best common practice initiatives to develop harmonization of IoT service layer APIs. This is analogous to traditional IT normalization initiative in data models and applications. Having a normalized and coherent model to allow the creation of IoT applications that are agnostic to the underlying protocols and frameworks within the Smart Home. Like the efforts around "Semantic Web & Technologies" standardization in the last few years, IoT could leverage similar concepts to further improve its capacity to understand things' data and facilitate their interoperability and device constraints

By enhancing these standards with an abstraction layer could be leverage by Semantic Web of Things and the opening the different device ecosystems at the application layer with APIs. In this framework our IoT enabled products could provide the necessary bridging capabilities between the "old world" design methodologies that guarantee standardized interop at the device communication layer and the application and services layer which is e2e. This allows an agile introduction of new functionality focused on consumer interaction and experience. This effort will only be successful with the right partnerships and the creation of the right business incentives to be able to finally unlock the business potential of the future smart home.

Conclusion

This paper covers an end-to-end view of an IoT architecture from a NSP perspective. Depending on what layer the NSP wants to own or provide value will determine the importance of conclusions drawn from this paper. The most ambitious of NSPs will want to own up to the platform layer, and to enable their own and third party services to ride on top of that platform.

In order to achieve that objective from a device layer perspective, the key is to minimize the number of SKUs to be handled and kitting to be done for curated sensor solutions. An example was given in section 1.3.1 of the synthetic “super sensor” work done at Carnegie Mellon, that could fulfill this minimization objective.

From a connectivity layer perspective, the NSP must own/drive the requirements for the IoT hub in the smart home because this is a key interworking and management point of presence in the home. The IoT hub must support multiprotocol connectivity, messaging and framework capable and have enough resources and a network operating system that can support containers. Highly coupled to the device layer, the ability to be conservative on the device type needs will allow the NSP to be focused and concrete on IoT radio requirements in the gateway or hub.

It is at the platform layer that the most critical work must be done, in order for an NSP to realize the objective of creating an open and inviting service environment for consumers and third parties. Along with a significant commitment in time and resources to realize the objective, historic tools such as TR-069 must be abandoned in favor of a careful selection of modern orchestration and device management methods which are evolving in the could native DevOps community.

Abbreviations

API	Application Programming Interface
NSP	Network Service Provider
IoT	Internet of Things
CAGR	Compound Annual Growth Rate
OTT	Over The Top
EMI	Electromagnetic Interference
RSSI	Received Signal Strength Indication
BLE	Bluetooth Low Energy
LoRA	Long Range Wireless
OTA	Over The Air
CPE	Customer Premise Equipment
GDPR	General Data Protection Regulation
LXC	Linux Containers
MQTT	Message Queuing Telemetry Transport
OCF	Open Container Forum
OMA	Open Mobile Alliance

Bibliography & References

Alliance, Z. (n.d.). <http://www.zigbee.org/>. Retrieved from Zigbee Alliance.

- C. Bormann Universitaet Bremen TZI, S. L. (2018). *CoAP (Constrained Application Protocol) over TCP, TLS, and WebSockets*. Internet Engineering Task Force (IETF).
- Gierad Laput, Y. Z. (2017). *Synthetic Sensors: Towards General-Purpose Sensing*. Pittsburgh, PA: Human-Computer Interaction Institute, Carnegie Mellon University.
- OASIS. (2015). *Message Queuing Telemetry Transport (MQTT)* . Burlington: OASIS.
- R. Fielding, E. A. (2014). *Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content*. Internet Engineering Task Force.
- Z. Shelby of ARM, K. H. (2014). *The Constrained Application Protocol (CoAP)*. Internet Engineering Task Force.

Network Convergence

A Technical Paper prepared for SCTE•ISBE by

Jennifer Andreoli-Fang, Ph.D.

Distinguished Technologist
CableLabs
j.fang@cablelabs.com

Alon Bernstein

Distinguished Engineer
Cisco
alonb@cisco.com

Aeneas Dodd-Noble

Principal Engineer
Cisco
noblea@cisco.com

Elias Chavarria Reyes, Ph.D.

Software Engineer
Cisco
elchavar@cisco.com

Curt Wong

Senior Director
Charter Communications
curt.wong@charter.com

Bernard McKibben

Distinguished Technologist
CableLabs
b.mckibben@cablelabs.com

Table of Contents

Title	Page Number
Table of Contents	2
1 Introduction.....	4
1.1 Motivation	4
1.2 Key focus.....	4
1.3 Why now?.....	5
1.4 Where is the party?	5
2 5G Reference Architecture.....	5
3 Mobile and Cable, Commonalities and Key Differences.....	7
4 Mobile Backhaul	9
4.1 Benefits	9
4.2 Architecture	9
4.2.1 Converged Transport Network QoS.....	10
4.3 Example use case	11
5 5GC Control of Devices in the Home.....	12
5.1 Benefits	12
5.2 Architecture	12
5.3 Example use case	13
6 5GC Control of The Home Gateway	14
6.1 Benefits	14
6.2 Architecture	14
6.3 Example use case	15
7 5GC Control of DOCSIS Network as an Access Network	15
7.1 Benefits	15
7.2 Architecture	16
7.2.1 Interworking Model.....	16
7.2.2 Integration Model	17
7.3 Example use case	19
7.4 Brief Overview of Other Forms of ConvergenceUnified Infrastructure for Cable and 5G.....	20
7.4.1 Benefits	20
7.4.2 Architecture	20
7.4.3 Example use case	21
7.5 Unified Manageability.....	21
7.5.1 Benefits	21
7.5.2 Architecture	21
7.5.3 Example Use Cases.....	22
7.6 Business Level Convergence.....	22
8 Conclusion.....	22
9 Abbreviations.....	24
10 Bibliography & References.....	24

List of Figures

Title	Page Number
Figure 1 – 5G Reference Architecture	6
Figure 2 – DOCSIS For Mobile Backhaul	11
Figure 3 – 3GPP control of devices in the home	13
Figure 4 – 5G control of the home gateway.....	15
Figure 5 – Convergence of 5GC and DOCSIS Functions for Legacy Devices.....	17
Figure 6 – Convergence of 5GC and DOCSIS Functions	18
Figure 7 – Cloud Native Landscape.....	20
Figure 8 – ONAP Architecture [17]	21
Figure 9 – ONAP Service Lifecycle [17].....	22
Figure 10 – Layers where mobile and cable can converge	23

List of Tables

Title	Page Number
Table 1 – Comparisons Between Cable and Mobile.....	7

1 Introduction

1.1 Motivation

Many cable operators operate a combination of wireline (DOCSIS), Wi-Fi, and mobile (LTE, and soon to be 5G) access networks. Each access technology is supported by separate cores with overlapping capabilities. An operator that provides both mobile and wireline services will need to purchase and maintain separate core network elements to manage their corresponding access networks – 5G Core (5GC) for next generation radio access network (NG RAN), and cable core for DOCSIS access. Converging these overlapping capabilities will provide savings for cable operators, while allowing the operators to offer new network services across uncoupled access networks.

Today, mobile and wireline core networks have some obvious disparities. They use

- different credentials to authenticate and authorize devices
- different data management
- different accounting and billing systems in the back office
- different policies to instantiate and manage data sessions

In spite of this, targeted convergence can reduce or bridge many of these differences without the need for massive replacements of customer CPEs. Appropriate updates of infrastructure features can enable the efficiencies realized through core convergence.

While the operators manage their multiple networks, the users should be network agnostic. With converged core networks, the user does not need to know which access network (e.g., fixed, Wi-Fi, cellular) s/he is on, but can expect seamless and consistent user experience.

Finding the sweet spot to converge the functionalities of the two different cores will benefit the operator's capex and opex, while providing the end user with seamless and consistent service experience.

1.2 Key focus

Core network convergence has been discussed in various forums and occasions in recent years. Despite its popularity, we are not aware of any previous work to systematically approach the subject. In this paper, we will approach it in a comprehensive way by exploring various convergence scenarios and architectures. The architectures range from a less involved interworking model, a semi-convergence case to a fully committed integration model.

- Interworking model: Interworking between the HFC network and 5G core is placed in the network. No change to today's DOCSIS CPE is required.
- Integration model: DOCSIS is a type of 5G access network being managed by the 5GC. Some modification to today's modem may be required.

We will address the benefit of each use case in the context of optimizing OPEX, synergy with the new cloud architectures, and implementability by equipment vendors.

1.3 Why now?

LTE EPC uses a point-to-point architecture, where each EPC component talks to another with a dedicated interface and procedure that is defined by the 3GPP. The 5GC, in contrast, uses more modern service based architectural techniques to provide better scaling and network configuration flexibility in comparison to the EPC point to point approach.

For example, while significant work has been done to the LTE protocol to support IoT devices, the new 5G core architecture better addresses the sheer number of devices (e.g., IoT) connecting to network while meeting the key performance indicators (KPIs) of some 5G services, e.g., ultra-reliable low latency communications (URLLC). The 5GC has been armed with the modern techniques to provide flexibility to meet the requirements of different vertical industries.

Even as the mobile operators might be slowly upgrading portions of the EPC to 5GC and not replace the legacy EPC wholesale, the old EPC interface protocols can be containerized and provided as a microservice to allow any network functions to use it.

The trend in the mobile world is synergistic with cable, as cable is also moving towards a cloud native architecture. As we can see, the recent evolution of the 5G mobile core from a point-to-point architecture to a service-based architecture provides a rare opportunity to converge mobile and cable cores.

The international mobile standardization body, the 3GPP, has been keenly working on specifying techniques to take advantage of the new 5GC. The 3GPP has an active study item on wireless-wireline converged core, with the specifications planned to be completed by the end of 2019. Vendors will be providing products soon after. This provides an immediate opportunity to insert HFC operator requirements for convergence in the global mobile ecosystem.

1.4 Where is the party?

The Broadband Forum (BBF), a telco-led organization, is actively studying FMC models, driven by operators and vendors such as Huawei, Nokia, Ericsson, etc. The study is being documented in SD-407 of BBF TR, where several models of how the 5GC and fixed broadband cores interact are being considered. Some solutions have been developed, but the TR is still in process. The BBF has been liaising with the 3GPP SA2 group to communicate the impacts to 5GC, and SA2 is working on incorporating the impact to the Stage 2 5GC specs as part of the “5WWC” study item. 3GPP intends to complete its convergence study in Q4 2018 and develop complete specifications by the end of 2019. There is an opportunity to insert cable industry requirements for convergence into this global scale effort within this time frame.

2 5G Reference Architecture

The 5G architecture is defined in 23.501 [7] and it describes a number of new network functions, some of which are evolutions of 4G functions, while others are new. There are also a number of new tenets, one of which is that non-3GPP access can seamlessly connect to the 5G core. The architecture does not though define what the backhaul requirements are, and choosing the appropriate solution is left to vendors and operators.

For 5G, 3GPP redefined the access network and made a clear separation between the access management and session management, from both node level and logical levels. The UE connects to the access network and the Access and Mobility Management Function (AMF) and from there can establish one or more PDU sessions with the Session Management Function (SMF). The N1 NAS mobility management is

handled by the AMF, while the session management of the UE's connections is handled by the SMF. In release 15 of the 3GPP specification, only untrusted non-3GPP access, i.e. connection over an IPsec tunnel over any WLAN access, is possible, however in release 16, support for additional access technologies is being defined in the 5WWC study item in TR 23.716 [11].

The AMF is limited to managing the access layer, which can include:

- New Radio (NR): the 5G air interface
- Non-3GPP Interworking Function (N3IWF): the IPsec gateway for untrusted access in the current release or in release 16, the addition of WLAN
- Wireline 5G Cable Access Network (W-5GCAN): the DOCSIS access network

Authentication and registration of the devices that connect to the access network is performed through the AMF, which communicates with the Unified Data Function (UDM) to get the profile and the AUSF where the connection is authenticated.

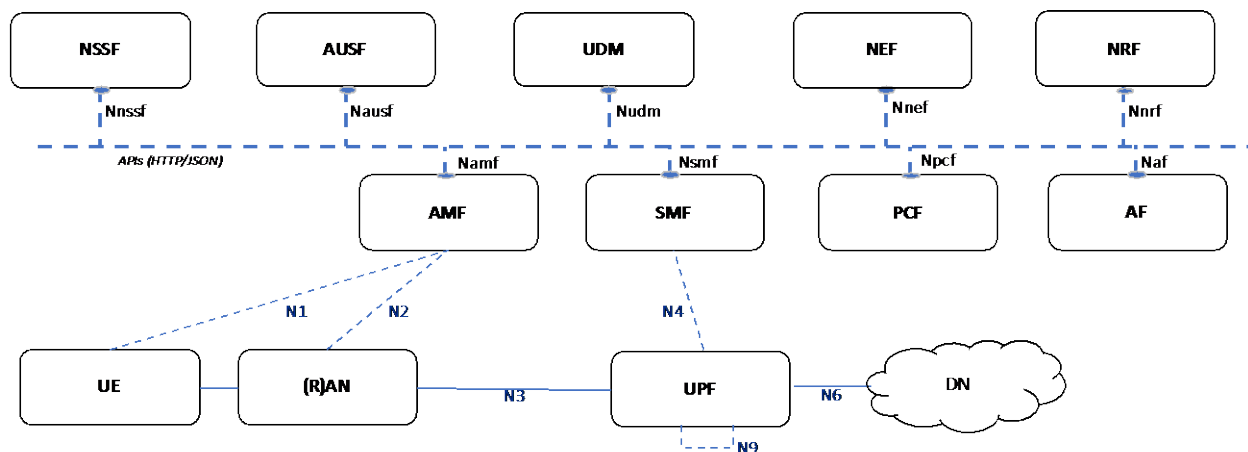


Figure 1 – 5G Reference Architecture

Figure 1 shows the fundamental components of the 5G core. The functions that are applicable to Cable Access are: AMF, SMF, User Plane Function (UPF), UDM, and Policy Control Function (PCF). The Radio Access Network (RAN) is the mobile equivalent to a CMTS in that it is the access part of the network. In addition, the UPF implements some of the functions of a CMTS (such as QoS control) although there are some subtle differences, which are explained in the following sections.

A major goal of wireless and wireline convergence work item is to unify the user experience. This goal can be achieved by applying the same policies from the PCF when the UE connects to either 5G new Radio (NR) or non-3GPP access such as cable. In 5G, the policies such as QoS, content filtering, SD-WAN, etc., are applied to the PDU by the PCF and in some use cases by the Network Exposure Function (NEF) and Application Function (AF).

The N1 interface is between the UE and the AMF and is relayed over the access network, i.e., the NR in 3GPP or cable network. This is a major change compared to 4G which means the 5GC can communicate seamlessly over non-3GPP access. If there are two simultaneous connections, then there are multiple N1 instances. The N1 Non-Access Stratum (NAS) (24.501 [14]) specification applies to untrusted access in release 15 and will be updated to address release 16 converged access requirements. N1 NAS provides authentication and authorization mechanisms, communicates the allowed slices and allows the UE to set

up PDU connections with associated IP addresses and policies, including traffic steering, offload, and routing policies. Furthermore, static and dynamic access selection rules to help the UE how best to use both access networks simultaneously are provided by the network.

The N2, also referred to as NGAP [11], interface performs Access Network configuration to associate the AN with the AMF and then manage the UE's connection at the AN. This includes PDU session management, session context which includes the Security Key, Mobility Restriction List, UE capability, etc., and transports N1 NAS messages. In release 16, there may be a new specification to capture the non-3GPP aspects.

The N4 interface, described in TS 29.244 [16] is between the SMF and the UPF. We could consider part of the cable access network as a UPF. The SMF associates with the UPF and exchanges capabilities and configuration if necessary. The PDU set up involves IP address management, policy, charging and service configuration. The specification may need to be updated to accommodate the cable network.

The N3 interface is for user plane traffic and is based on GTP [12] and [13]: The AN and the UPF use a tunnel to encapsulate the UE's datagram to allow the UE to have any IP address or even Ethernet or other non-IP traffic. QoS is marked in the packet so that the AN can prioritize the traffic based on policies applied in the UPF.

3 Mobile and Cable, Commonalities and Key Differences

The following tables outlines the key commonalities and differences between the cable and mobile cores.

Table 1 – Comparisons Between Cable and Mobile

	Cable	Mobile
Goal	Manage access line to home end-users	Manage the end-users directly
Scope and complexity	<p>DOCSIS is an access network protocol that defines the interface between CM and CMTS.</p> <p>The CMTS is a subset of the 5GS. Many 5GC control plane concepts do not apply in the cable environment, such as connection and mobility management.</p> <p>The CMTS is also a “god box”, a single entity that implements many functionalities ranging from communicating with the CM on the RF side, and policy enforcement point for a SF.</p> <p>With DAA, at least the RF portion is modularized and some split options in 5G are similar in concept to those in Remote PHY.</p>	<p>3GPP specifies an end-to-end architecture capable to manage registrations, connections, sessions, mobility, reachability, access authentication, policy, as well as charging and authorization.</p>
Subscriber	CMTS manages a CM, which is a connection point to a household	5GC manages a UE which is an individual end-device, similar to a CPE in the cable world.

	Cable	Mobile
Provisioning	CM is provisioned via TFTP config file server. DHCP hands out IP address. A CM can have only 1 IP address.	UE gets IP address from the SMF, or via the SMF acting as a DHCP relay, and UE can have multiple IP addresses, one per data network.
Identity, authentication	Each CM is identified by its MAC address. Uses “implied” authentication, because the CM is provisioned by the backend. Although the IP address of devices directly behind the CM are visible to the CMTS, individual user identity in the home is not visible to DOCSIS in a typical deployments because of the presence of gateways and NATs. Only the cable modem identity is established	SIM is used for identity; each UE has a separate SIM and may have multiple SIMs in the near future. Note this is the end user identity, not the “access device identity” which is the case in cable.
Policy	Focused on providing bandwidth policy / QoS for a service flow in order to provide an SLA. QoS policy is defined per service flow. CMTS implements network admission control, i.e., it is focused on bandwidth to the home and its allocation, rather than end user. In principle, a policy per device behind the CM can be defined. However, since the home devices are typically behind a gateway, a direct mapping is not easy (especially in the case of NAT). For FMC, cable may borrow some of the mobile policy concepts, with the understanding that they apply to access to the subscriber rather than the subscriber him/herself.	Encompasses resource and traffic management for subscription tiers and applications. QoS policy can be applied per PDU session.
Accounting and charging	Operators account for data usage. IPTV and voice are accounted but is zero-rated. Billing is per household, not per device / user. CMTS reports accounting records on bytes used. Billing is done by the back office.	Charging can be done per device, or 3GPP can link accounting records of multiple subscribers, i.e., “shared” billing plan, which is similar to cable. 3GPP has collapsed online and offline charging into a single Charging Function.
Vendor ecosystem	Because the CMTS is a combination of complex RF level implementation as well as user and control plane management all in one, this limits CMTS vendors to the traditional 4 players.	The network infrastructure side supports an ecosystem dominated by cellular vendors. However, with innovation going on to disaggregate the eNBs and the gNBs, MNOs, particularly MSOs who are also MNOs, will have the opportunity to deploy

	Cable	Mobile
	Particularly the space of silicon vendor for the RF components is extremely limited.	interoperable components from multiple, smaller vendors.
	“Cloudifying” the CMTS and converging with mobile will allow new vendors to emerge and enable product differentiation.	On the core side, as the 5GC architecture moves to containers and microservices, the new 5GC may start to enable modularity and interoperability between vendors.

Although the table shows that the concepts and ecosystem of 5G and cable are different, we are now at the cusp of a wave of innovations enabled by virtualization and disaggregation of components. Now is the time to explore the convergence between 5G and cable. Let us start with some obvious areas of convergence, such as the area of bandwidth management and QoS, and how they can be managed by a 5G toolkit.

4 Mobile Backhaul

4.1 Benefits

One of the major costs incurred in building a mobile network is the mobile backhaul. Indeed, among the major cost components in deploying a small cell network, including spectrum, network equipment, site lease, and power, backhaul alone accounts for 61% of the cost per GB [1]. The existing DOCSIS networks have been enhanced [2] to provide a backhaul service that is comparable to fiber performance but at a significant economic advantage [3]. Key use cases for mobile operators are:

- A cable operator building a mobile infrastructure
- A cable operator leasing backhaul capacity to a mobile operator
- In-home small cell, with “inside out” coverage

In [3], the authors explain in detail the market opportunity for a DOCSIS mobile backhaul, as well as options to optimize the DOCSIS protocol for mobile backhaul applications.

This paper focuses on the system level integration of DOCSIS backhaul in a 5G ecosystem.

It is worth noting that the DOCSIS technology can effectively backhaul, mid-haul, or fronthaul (for most splits) mobile traffic. This is contrary to what the fiber optical vendors have been telling the mobile and cable industries. Fronthaul is no longer synonymous to Common Public Radio Interface (CPRI), a closed specification created for macrocells by traditional radio vendors. DOCSIS has been shown to support lower layer splits such as MAC-PHY or intra-PHY [4] (Option 7 defined by 3GPP). Below that, the capacity requirements may become too high for DOCSIS to support. For the sake of brevity, we will focus on the backhaul use case in this paper.

4.2 Architecture

The 3GPP specifications focus on the RAN and the packet core. There is not much discussion about the backhaul as they are generally thought of as being part of the N3 transport interface. In Figure 2, the DOCSIS network is embedded inline with N3 (see Section “5G Reference Architecture” for more detail on N3). The UE has a PDU session on the RAN which is backhauled over the DOCSIS network. The UE

may also have hybrid access capability which means that it can reach the data network via a traditional macrocell network that is backhauled by an MNO's traditional transport.

To provide backhaul service comparable to fiber, the DOCSIS network needs to match the QoS on the backhaul to what is required for the PDU session. This can be done either dynamically after the session has been set up, and actively allocating resources on the DOCSIS side as proposed in [3], or semi-statically during session setup. The advantage of the latter is that it provides flexibility to the backhaul operator to define policies on when to admit / reject a PDU session request from the UE. For example, depending on how the DOCSIS backhaul is configured, it may be able to provide better or worse quality of service (QoS) compared to the other access networks such as the macrocell network. One open question is how can the DOCSIS system signal to the 5G ecosystem that there is excess resource, or the lack of, on the DOCSIS backhaul side? The following section outlines a possible solution based on an interworking function (IWF) between the DOCSIS backhaul and the 5GC.

4.2.1 Converged Transport Network QoS

To perform network resource optimization in order to achieve system-level QoS, one potential solution is a converged transport QoS framework which includes an IWF that talks to the wireless and the DOCSIS networks, and works with the AMF to determine best path for the UE's PDU session to the data network. It also provides session information to the AMF and therefore the UE, should a handover be initiated when an alternative path is deemed more optimal.

The IWF is a resource coordinator, shown in Figure 2. The IWF acts as a controller that reads various utilization information from the gNBs and with this information, it makes decisions on how to optimally balance resource usage across the wireless network. In order to do that it communicates with the AMF, SMF, and PCF. The IWF needs to talk to an AMF to discover the SMF associated with a UE's PDU session. The IWF can then retrieve the PDU session QoS information from the SMF. The information it needs in the mobile backhaul case is the utilization of the DOCSIS backhaul link and some policy associated with it (e.g. move UEs if the utilization crosses a certain percentage).

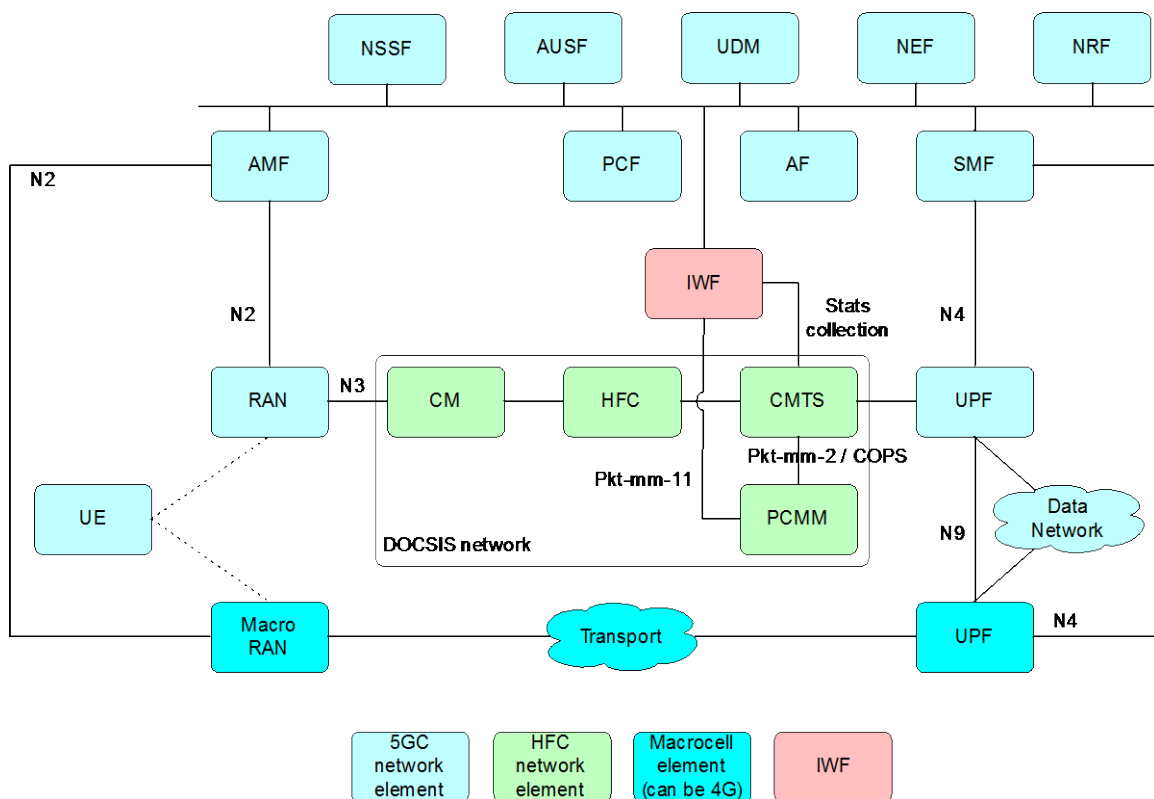


Figure 2 – DOCSIS For Mobile Backhaul

As depicted in Figure 2, the IWF can help deciding when to carry UE traffic over DOCSIS backhaul vs. other mobile options (e.g. macro RAN).

The interfaces pkt-mm-11 and pkt-mm-2 are standard Packet Cable Multimedia (PCMM) interfaces as explained in [6]. The “stats collection” interface can be any existing interface that the CMTS has to report service flow utilization, e.g. SNMP/IPDR etc.

Rather than a separate entity, the coordinator may be implemented as part of the CMTS control plane, although this might require tighter integration between the backhaul and the 5GC.

The IWF decisions can be very fine grained and based on a session-by-session input. Additional capacity can be added or subtracted as needed on the DOCSIS side based on session information on the mobile side. Another simplified and coarse-grained solution is to view DOCSIS as a static link, e.g., the DOCSIS backhaul presents a static capacity, and in that sense behaves in a similar way to any static peer-to-peer link, such as microwave or fiber. The IWF monitors the utilization on this fixed link, and as it approaches a certain threshold IWF can take a corrective action.

4.3 Example use case

A simple use case is one where UEs keep joining a small cell with a DOCSIS backhaul. At some point (depending on how the backhaul is configured) the utilization of the DOCSIS backhaul link may get too high. Monitoring such utilization is easy for the CMTS because it has service flow statistics and the backhaul traffic is carried over one or multiple service flows. The utilization statistics can be communicated to the IWF in the form of percentage. The IWF can have a local policy, for example, “take

action if the DOCSIS backhaul utilization exceeded 90%”. The form of the action is contained in the IWF, and the associated policies are outside the scope of the DOCSIS backhaul. For example, the IWF can start moving calls to an adjacent cell or a macrocell. The elegance of such a solution is that all the DOCSIS system has to do is to report the backhaul link utilization and from that point on the IWF can take over.

5 5GC Control of Devices in the Home

5.1 Benefits

If an operator considers merging policy of the two network domains, then a common use of policy appears to be an obvious choice. One option for 5GC is to control the connected devices in the home as if they were regular mobile devices.

The benefits of using 5GC to control devices in the home are:

- Consistent policy definition framework
- Consistent policy enforcement point (device and/or GiLan for the most part)
- Easy handover from a wired/Wi-Fi cable network to a mobile network and consistent behavior across these fixed/mobile handoffs

5.2 Architecture

In Figure 3 the devices in the home access the data network via a fixed line access network (AN). All the other surrounding 5G components are kept so that now we have a consistent control framework for both 5G and fixed AN.

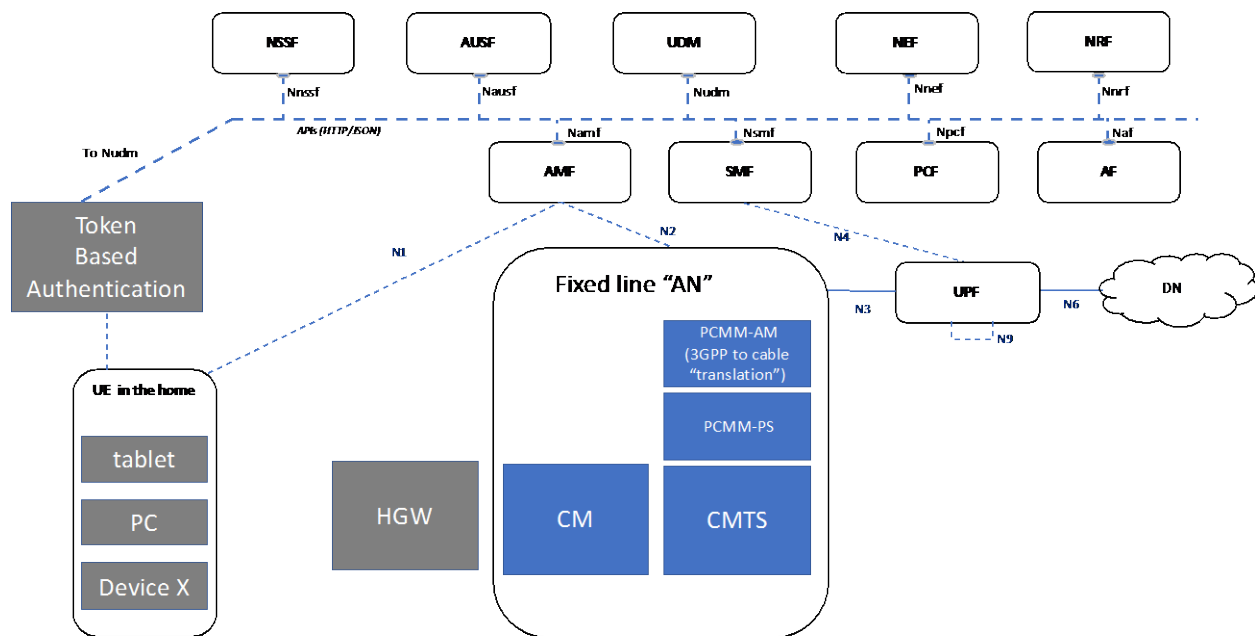


Figure 3 – 3GPP control of devices in the home

Note that we are not replacing anything within the DOCSIS system itself, i.e. the cable modem registration, policy control of CMTS and CM are still all DOCSIS based. This is because DOCSIS system is still the best set of standards to handle cable access at this time. However, when it comes to control of devices in the home, which is outside the scope of DOCSIS system anyway, it makes sense to align with wireless.

Because the home devices can access the internet without a SIM, an alternative identifier may need to be selected for the devices. An identifier can be communicated with the UDM, and from that point on the normal 5G processes can be executed. Note that token-based authentication is one of multiple solutions being explored. An alternative may be selected for standardization.

Token-based authentication is fairly common in commerce these days. It typically involves sending a unique code to a trusted device (e.g. a mobile phone) that a user can then type into an authorization page in order to prove that she is indeed the user authorized to use the device.

5.3 Example use case

Consider parental control as a policy example. In the current ecosystems, parental control rules that are defined on a mobile device do not automatically apply in the home. Therefore, a website that cannot be accessed from a mobile device might be accessed in the home over Wi-Fi – unless a separate policy is defined on the home gateway to block it. With the approach outlined in this section, a single parental control policy would apply to both mobile devices using SIM-based authentication and the home Wi-Fi environment.

Consider two separate use cases to illustrate this in more detail:

Use case 1: A mobile phone authenticated with a SIM is subject to the mobile policy rules. Once the same phone is in the home it will switch to Wi-Fi and be subject to the set of rules in the fixed access network. It is the same device with the same SIM so identity is easy to establish and we just have to make sure the policies are coordinated across the mobile and fixed access networks.

Use case 2: Certain parental controls are defined for a child. The same child can own several devices. One of them will have a SIM (the mobile) and therefore, a clear identity and a parental control policy associated with it. Other devices that may be used by the same child, e.g. an iPad in the home, that do not have a SIM, and token-based authentication is used to establish identity instead. Once identity is established it can be associated in the UDM to the primary SIM based identity and the same policies will be applied across several devices.

6 5GC Control of The Home Gateway

6.1 Benefits

Instead of controlling the home devices directly with the 5G infrastructure, it is possible to control the home gateway, and the home gateway in turn controls the devices in the home.

The benefits of using 3GPP to control devices in the home through the home gateway are:

- No need to change anything about the home device operation (i.e. token-based authentication) because they are not directly controlled by 5G
- Centrally managed policies

6.2 Architecture

A common protocol for controlling the Customer Premise Equipment (CPE) is TR-69. With this architecture, the 5GC is controlling the ACS server as if it were a UE. This will require the ACS server to have middleware written for it to translate N1 messages into a format that the home gateway will recognize. Note that TR-69 is nothing but a secure transport, it does not mandate the actual message format. And since every home gateway vendor might have different definitions, each home gateway vendor will need to create their own “N1 to ACS translator”. The actual business arrangement and implementation issues are outside the scope of this paper.

For this use case, authentication is required because the home gateway does not have a SIM. A token-based authentication, similar to the one used for the “5GC Control of Devices in the Home” can be used as well.

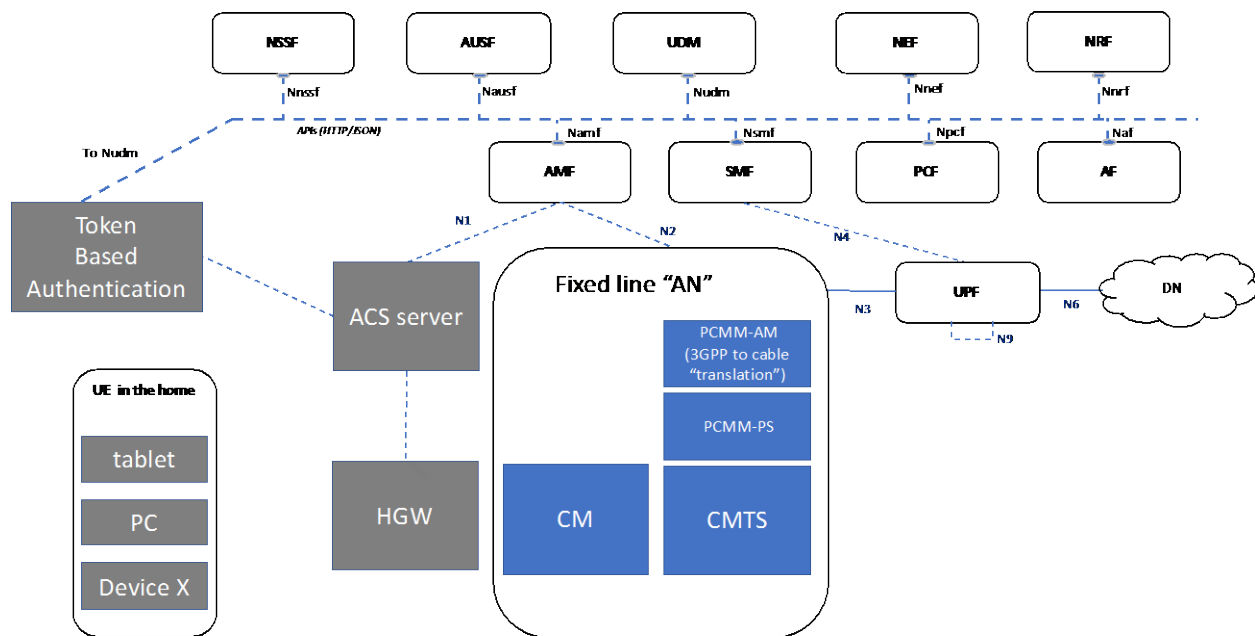


Figure 4 – 5G control of the home gateway

It is possible to have a hybrid model of N4 control of both HGW and GiLan. But such a model is out of scope for this paper for simplicity.

6.3 Example use case

The same “parental control” use case as in “5GC Control of Devices in the Home” applies to this use case as well.

One advantage the home gateway control method has is that devices in the home are controlled by the home gateway, so there is no need for a token-based authentication on the home devices. However, this means that someone has to configure the home gateway to associate identities to devices (and it is likely that many of the readers of this paper are the “IT manager” at home and are up to the task of doing exactly that). This also means that if you have visitors to your home, an identity will have to be established for them, otherwise their “parental control” rules will be subject to the default policy on the home gateway.

7 5GC Control of DOCSIS Network as an Access Network

7.1 Benefits

If cable and wireless were developed together from ground zero, we could have modeled the CMTS as a type of UPF and AMF, and the CM as a type of a UE. The HFC could have been considered a type of access network. This is in fact what the 3GPP System Architecture 2 (SA2) group has been working on

since early this year, in collaboration with the Broadband Forum (BBF). After all, the NG RAN, the LTE air interface, Wi-Fi, and in our case, the cable HFC are all access networks that push bits to the end user (and receive bits on the return path).

The emergence of 5GC with a service-based architecture enables us to converge selected network functions with less dedicated infrastructure, and to exploit new 5G network services across uncoupled access networks. Later in this section, we discuss a list of potential use cases.

7.2 Architecture

7.2.1 Interworking Model

As stated in “Introduction”, there are two models for converging the cable network and the 5GC. The less impactful interworking model is shown in Figure 5. Fundamentally, the key principles include:

- No impact on CPE
- Burden of interworking is largely placed on the HFC infrastructure
- HFC authenticates CMs with existing mechanisms, registrations posted into the UDM

With an interworking function implemented by the cable operators and vendors, some enhanced services can be enabled without any impact to the HFC network or changes to the 5GC components:

- Converged QoS: this has already been discussed in “Mobile Backhaul” earlier
- PCF links into the north side of HFC policy platforms
- HFC charging records can be updated for better correlation with 5GC records
- HFC infrastructure may communicate with NSSF, NEF and NRF for 5GC services

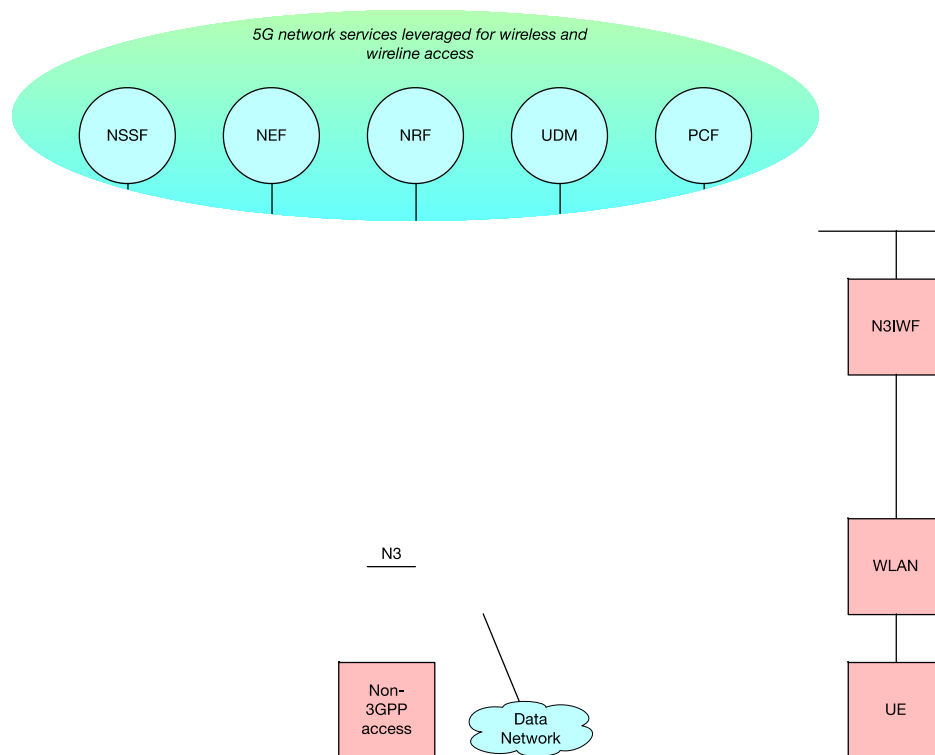


Figure 5 – Convergence of 5GC and DOCSIS Functions for Legacy Devices

7.2.2 Integration Model

In the longer term, the operators should evaluate the business case for a fully integrated model. A possible architecture is shown in Figure 6, in which the CM appears as a 3GPP UE to the 5GC. This means the next generation CM needs to support a “profiled” version of N1 interface, labeled as N1’. The HFC network on the other hand will need to support a “profiled” version of N2 interface, which links into the AMF.

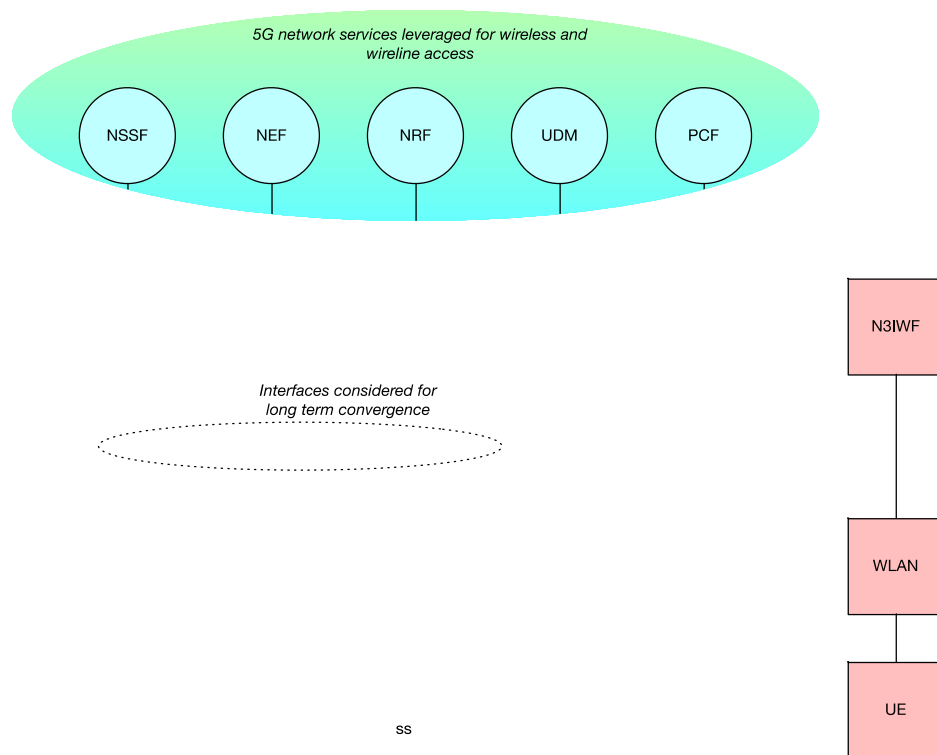


Figure 6 – Convergence of 5GC and DOCSIS Functions

While N1 and N2 interfaces encompass a large number of services, some of these services, such as mobility support, are not applicable to fixed AN. So the additional 5GC functionalities that need to be supported by the HFC network can be tailored and reduced.

Since the AMF and the UPF are control and user plane functions, respectively, and the CMTS already has control plane and user plane separation, the CMTS can implement a profiled AMF, labeled as AMF', and a profiled UPF, labeled as UPF', without exposing the N2 and N3 interfaces externally.

In 5GC, the AMF discovers the SMF that is associated with a PDU session, in order to obtain the set of QoS policies associated with the session. In the collocated case where the CMTS implements the SMF', these will become internal procedure to the CMTS.

The HFC network could be updated to generate 5GC accounting records. The AMF' and the SMF' will have to appear to the NSSF, NEF and NRF with 3GPP standard network element interfaces, so that the HFC can directly query these functions.

The network functions NSSF, NEF, NRF, UDM, PCF could be leveraged to enable converged services between cable and mobile, to enable:

- Integrated fixed (CM) and mobile subscriptions
- Converged accounting record system
- Fixed access aware network slicing
- Fixed access aware network exposure function for higher value data intelligence
- Converged service discovery
- Fixed network aware 5G traffic steering and splitting

Cable access and mobile NG RAN will manage its own

- Customer devices
- PHY and MAC layers
- Credential and authentication centers

It may be beneficial for the operators to start with the interworking model, while evaluating the benefit and the effort in the long run for the integration model.

7.3 Example use case

Some of the example scenarios that can be enabled by the integration model include:

- Correlation between fixed and mobile networks for an end-to-end user service (unified subscription, QoS and charging)
- End-to-end network slice that applies network resources across fixed and mobile segments for residential or industrial customer applications
- End-to-end data exposure
 - Improved data about your network and subscriber data patterns
 - Opportunity to monetize the data
- Converged traffic steering and traffic splitting across access networks (5G, 4G, Wi-Fi via HFC)
 - Steering and splitting native features in UDM, PCF, SMF and UPF
 - Continuous services across access networks

While the concept of integrating the 5GC and the DOCSIS network enables systematic convergence at the network element level, some of the use cases that are of interests to the operators can be enabled individually without a full-on convergence between the 5GC and the DOCSIS network.

For example, for the case of uniform parental control through “5GC Control of Devices in the Home”, all we need is for the cable operator to provide an interworking function that can link the token-based identity (or alternative identity) to the IMSI on the UDM. As such, the use case can be enabled as part of the interworking model, without a complete integration between fixed and mobile network elements.

It is important for the operator to identify high value use cases and analyze the short-term and the long-term benefit in light of the cable network updates.

7.4 Brief Overview of Other Forms of Convergence Unified Infrastructure for Cable and 5G

7.4.1 Benefits

With the advent of virtualization technologies, both ETSI-NFV and cloud Native approaches, it's possible to host both cable and wireless network functions and cloud functions on the same physical hardware because there are cloud/virtual versions for both. Fortunately, more than the hardware can be shared. Many of the software infrastructure components for these functions are open source and are common to many different use cases, from firewalls to virtual routers. In this world of common HW and common SW infrastructure cable and mobile are simply two different types of applications sharing many resources and expertise across different domains. Furthermore, operational methods, such as devops can now be shared across these functions so that an expert can easily handle devops in one domain and with little training master the other.

7.4.2 Architecture

Figure 7 is a high-level capture from the cloud native foundation (cncf.io) of the various tools used for developing cloud applications. An implementor may choose any subset of the tools here. See Ref [5] for an interactive map.

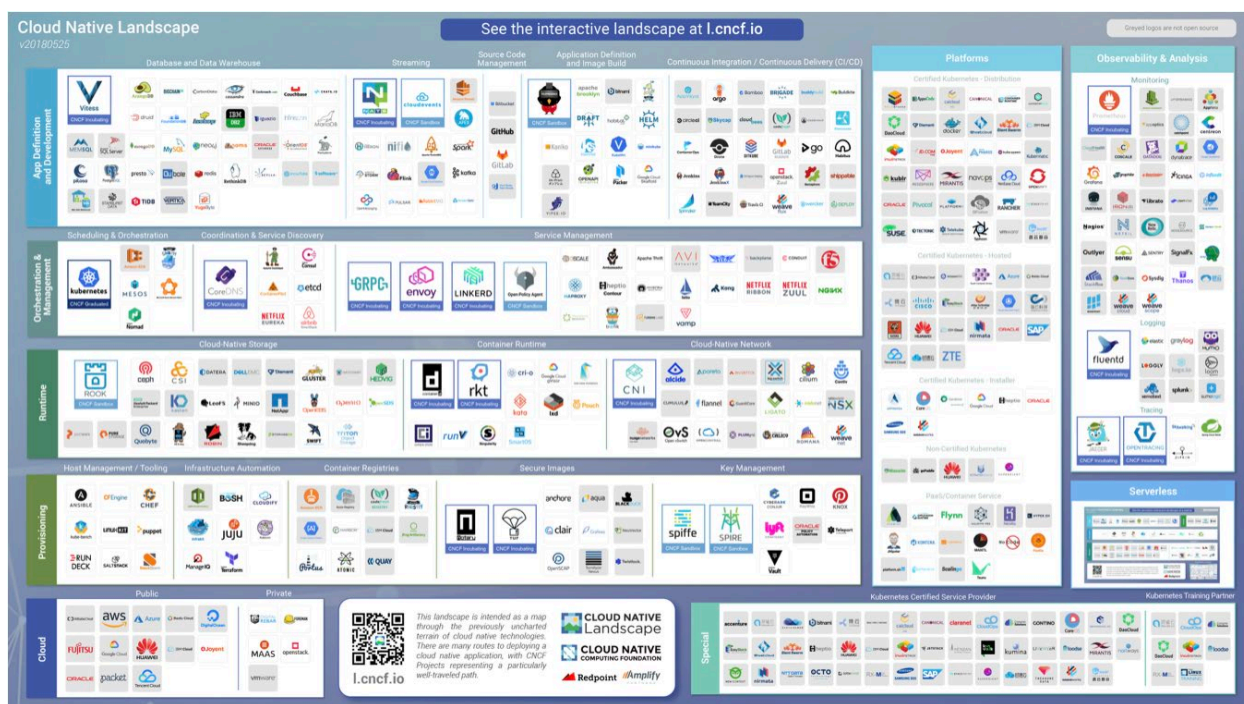


Figure 7 – Cloud Native Landscape

As can be seen from these diagrams there is a huge number of tools, from CI/CD tools to loggers, databases, buses and more. Functional components such as the packet core and the CMTS can be applications built on top of this infrastructure.

7.4.3 Example use case

- Mobile use tends to be high during the day. When subscribers are home fixed line usage goes up. With a common infrastructure compute resources can be moved by time of day
- Same devops process for SW upgrade for both mobile and cloud CMTS.

7.5 Unified Manageability

7.5.1 Benefits

Similar to the unified infrastructure case, the same set of tools can be used to manage both the mobile side and the cable side. For example, a stats collector polls managed objects and should not care if these managed objects are part of a mobile system or a wireless system.

7.5.2 Architecture

One option for unified management is ONAP (Open Network Automation Platform), see [17].

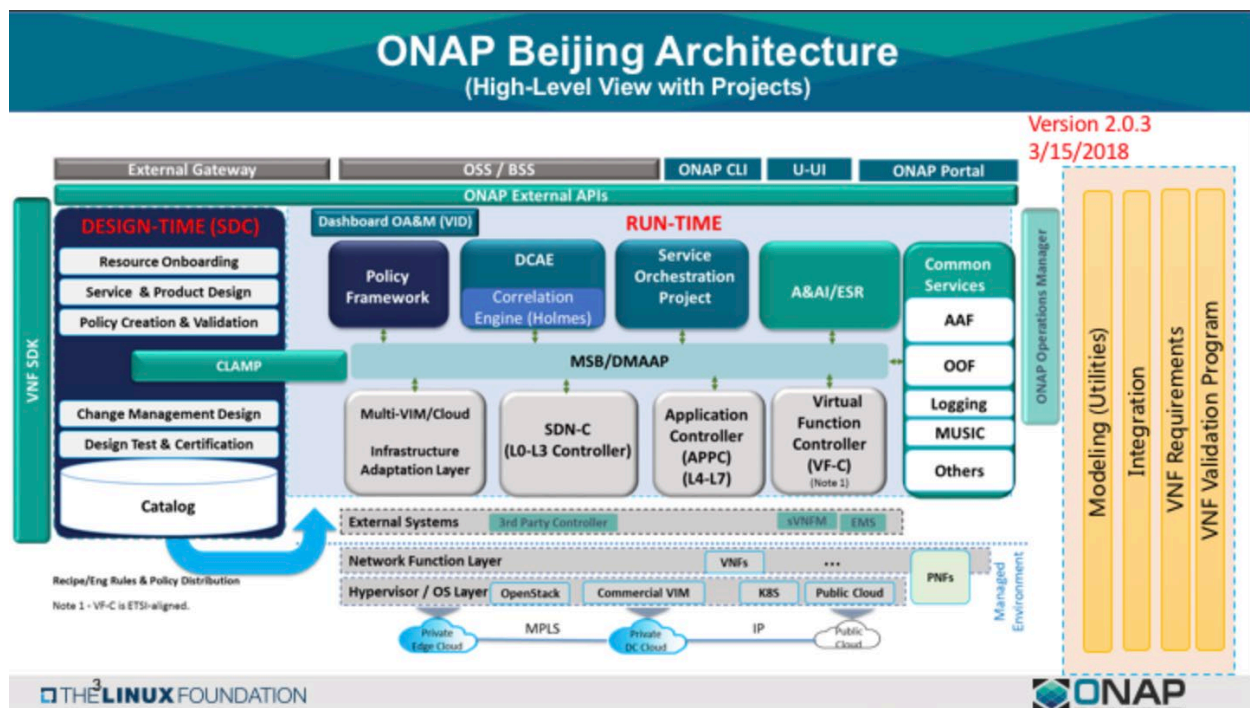


Figure 8 – ONAP Architecture [17]

Detailing ONAP is outside the scope of this paper. However, at a high level it's a full network management suite, currently optimized around virtual appliances but can apply to physical appliances as well. It can help the full automation life cycle that is depicted below:

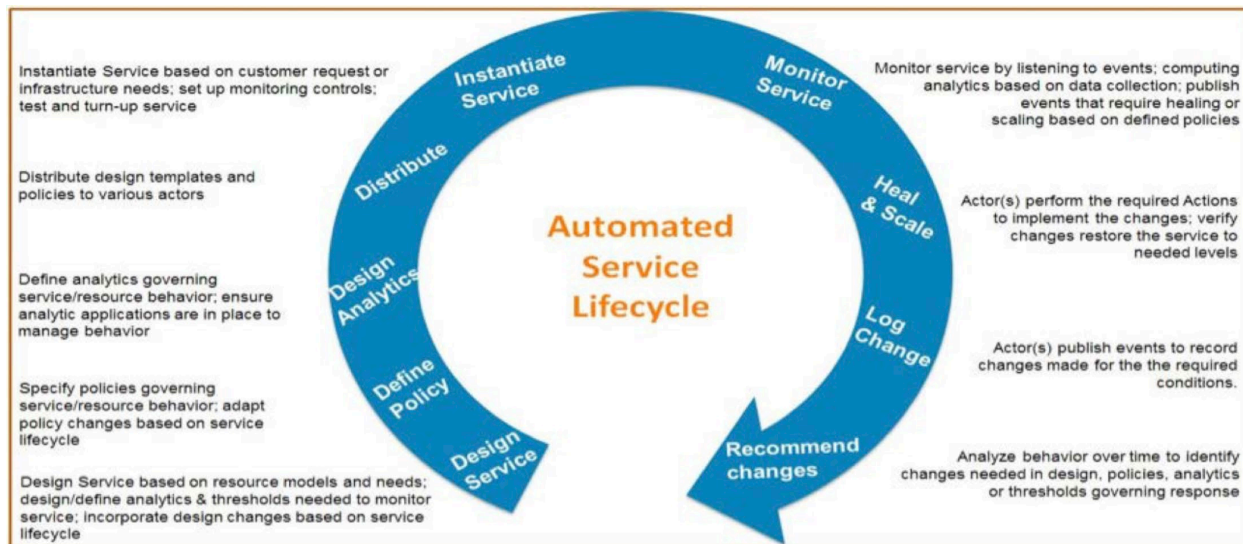


Figure 9 – ONAP Service Lifecycle [17]

Note that at this high level the underlying access technology can, and should, be abstracted. In the bigger picture the ability to build and support services is one of the more OPEX intensive tasks, and one where convergence can be achieved as it is all designed to work above the network layer.

7.5.3 Example Use Cases

An example use-case may be the definition of a VPN service that can work over mobile or wireless and eventually reaches the same firewall regardless of the access technology.

7.6 Business Level Convergence

It is possible to converge only at the business layer and leave all the technical layers below separate. With business level convergence, the customer receives a single bill for both the wireless and the mobile subscription, regardless of the technical implementation, e.g. converged, non-converged, an MVNO agreement with a mobile operator or a native wireless network. The one technical aspect this type of a convergence might have is on the provider icon displayed on the mobile, e.g., if cable company ABC sells a mobile service it will need to show ABC on the mobile phone. Cable company ABC will also need to issue its own SIM cards and have the ability to populate the customer database on the UDM.

8 Conclusion

This white paper presents a look ahead of emerging wireless-wireline convergence trends within the global ecosystem.

An access network is a multi-layered system as depicted in Figure 10. The layering itself is similar for both cable and mobile because the underlaying design is common to many communication systems. In this paper, we have provided a brief overview of convergence at the different layers, with our main focus and contribution at the network element layer. We are at a critical time where the 3GPP, the BBF, and the cable industries are innovating synergistically. This is a rare opportunity to converge at the network element layer.

Specifically, recent work to introduce service based architectures in mobile cores and wireline cores presents an opportunity to leverage core convergence in order to:

- Reduce OPEX
- Apply new 5G network services across mobile and wireline access networks

Can operators who own both mobile and HFC networks exploit these recent trends to provide improved subscriber services at lower costs? How deep into the network does convergence need to penetrate in order to produce savings and improved services? Solutions for light convergence and fully integrated convergence are now being identified. The 3GPP will standardize its view of core convergence by the end of 2019. Operators of mobile and HFC networks can work together with their technology suppliers to evaluate the convergence level solutions that best meets their needs. Now is the time to insert HFC operator requirements for convergence into the global ecosystem.

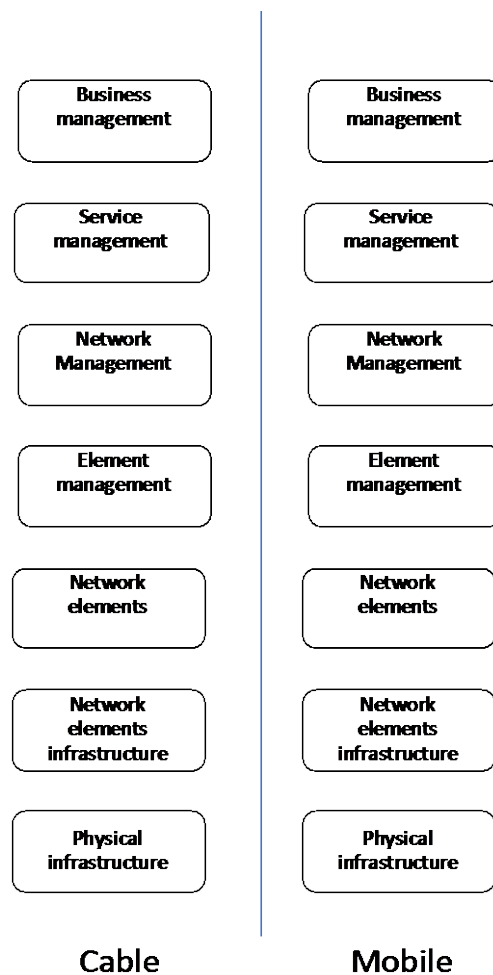


Figure 10 – Layers where mobile and cable can converge

9 Abbreviations

3GPP	Third Generation Partnership Project
5GC	5G Core
5GS	5G System
AF	application function
AMF	access management function
AN	access network
BBF	Broadband Forum
EPC	evolved packet core
FEC	forward error correction
FMC	fixed mobile convergence
HFC	hybrid fiber-coax
HD	high definition
Hz	hertz
ISBE	International Society of Broadband Experts
KPI	key performance indicator
NEF	network exposure function
NG RAN	next generation RAN
NR	new radio
NRF	network repository function
N3IWF	non-3GPP interworking function
PCF	policy control function
RAN	radio access network
SA2	System Architecture working group 2
SCTE	Society of Cable Telecommunications Engineers
SMF	session management function
UDM	unified data function
UPF	user plane function
URLLC	ultra-reliable low latency communications
W-5GCAN	wireline 5G cable access network

10 Bibliography & References

- [1] Kyung Mun, “Spectrum – Will its intrinsic value decrease in a small cell world?” October, 2016. <https://www.fiercewireless.com/wireless/mun-spectrum-will-its-intrinsic-value-decrease-a-small-cell-world>
- [2] “Synchronization Techniques for DOCSIS Technology Specification,” D-01 Draft 2018-07-30. CableLabs.
- [3] John T. Chapman and Jennifer Andreoli-Fang, “Mobile Backhaul over DOCSIS,” Proceedings of SCTE Fall Technical Forum, October, 2017. <https://www.nctatechnicalpapers.com/Paper/2017/2017-mobile-backhaul-over-docsis>
- [4] “Creating an ecosystem for supporting non-ideal fronthaul,” Telecom Infra Project white paper, February, 2018. https://telecominfraproject.com/wp-content/uploads/VRAN_WP_final.pdf

- [5] Cloud native toolkit. Cloud native toolkit. <https://github.com/cncf/landscape>
- [6] PacketCable Multimedia Specification (PCMM). I-07, 2018-11-11. CableLabs. <https://apps.cablelabs.com/specification/packetcable-multimedia-specification>
- [7] 3GPP TS 23.501, “System Architecture for the 5G System”,
<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>
- [8] 3GPP TS 23.502, “Procedures for the 5G System”,
<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3145>
- [9] 3GPP TS 23.244, “Wireless LAN control plane protocol for trusted WLAN access to EPC”,
<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=1062>
- [10] 3GPP Document 23.716
- [11] 3GPP TS 38.413, “NG Application Protocol (NGAP)” www.3gpp.org/DynaReport/38413.htm
- [12] 3GPP TS 29.281, “General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)” www.3gpp.org/dynareport/29281.htm
- [13] 3GPP TS 38.415, “PDU Session User Plane Protocol” www.3gpp.org/dynareport/38415.htm
- [14] 3GPP TS 24.501, “Non-Access-Stratum (NAS) protocol for 5G System (5GS)”
www.3gpp.org/dynareport/24501.htm
- [15] 3GPP TS 24.502, “Access to the 3GPP 5G Core Network (5GCN)
via Non-3GPP Access Networks (N3AN)” www.3gpp.org/dynareport/24502.htm
- [16] 3GPP TS 29.244, “Interface between the Control Plane and the User Plane nodes”,
www.3gpp.org/dynareport/29244.htm
- [17] ONAP, Open Networking Architecture Platform: <https://www.onap.org>

Network Planning Automation Using Big Data

An Analytical Platform for Infrastructure Capital Planning

A Technical Paper prepared for SCTE•ISBE by

Ted Boone

Senior Director, Network Engineering
Cox Communications, Inc.
Atlanta, GA
404-269-4771
ted.boone@cox.com

Jignesh Patel

Principal Architect
Cox Communications, Inc.
Atlanta, GA
404-269-6898
jignesh.patel@cox.com

Rob Ames

Executive Director, Network Engineering
Cox Communications, Inc.
Atlanta, GA
404-269-5451
rob.ames@cox.com

Kyle Cooper, Cox Communications, Inc.

Chaitanya Vasamsetty, Cox Communications, Inc.

About Cox Communications

[Cox Communications](http://www.cox.com) is a broadband communications and entertainment company, providing advanced digital video, Internet, telephone and home security and automation services over its own nationwide IP network. The third-largest U.S. cable company, Cox serves approximately 6 million residences and businesses. Cox Business is a facilities-based provider of voice, video and data solutions for commercial customers, and Cox Media is a full-service provider of national and local cable spot and digital media advertising. Cox is known for its pioneering efforts in broadband, voice and commercial services, industry-leading customer care and its outstanding workplaces. For eight years, Cox has been recognized as the top operator for women by Women in Cable Telecommunications; Cox has ranked among DiversityInc's Top 50 Companies for Diversity 13 times. More information about Cox Communications, a wholly owned subsidiary of Cox Enterprises, is available at www.cox.com and www.coxmedia.com.

All ideas, concept, principle, processes, and, method of operation described and detailed in this paper are original work of Cox Communications and are subject to copyrights.

Table of Contents

Table of Contents	3
Introduction.....	4
Network Planning Automation.....	4
Business Challenge.....	4
1. Big Data ecosystem	5
2. Data preparation.....	6
3. Analytics	6
4. Business Rules.....	7
5. Publication.....	7
Conclusion.....	7
Abbreviations and Definitions.....	8

List of Figures

Figure 1 - Network Planning using Big Data Infrastructure.....	6
Figure 2 - Network Transformation Visualization of select market	7

Introduction

Cox Communications is currently transforming its network infrastructure to support gigabit symmetrical speed offerings. This capital-intensive project needs a strategic plan providing prioritization using long term forecast accounting risk and macro-economic factors. Further, adding business constraints, budgetary restrictions, and other operational limitations we produce an optimal actionable plan.

The key business challenges with the current process of producing prioritized plan are scaling, repeatability, and traceability. Today, there are thousands of nodes in the network, but we will scale to a few hundred thousand nodes to meet network and customer growth. This increase made it difficult to continue with the prior manual process of building a strategic network plan. Data preparation relied on a manually extracting data from multiple sources to create a factual view of network bandwidth consumption. Business rules were applied in spreadsheets using macros.

Automation of this process had opportunities on several fronts, notably, providing consistency, repeatability, and modernization with the use of data science algorithms on an enterprise Big Data platform. Business requirements are made more transparent, and configurable, allowing planners to run multiple what-if scenarios faster to assist with the strategic decision process.

Network Planning Automation

Business Challenge

Last year, busy hour traffic on Cox's access network grew by almost a third. Downstream growth has softened a bit recently, dropping from near 50% in years past. Much of the drop coincided with re-encoding of video content by large over-the-top providers like Netflix since 2016. Upstream has been stable; only slightly off the 26% 5-year trend at 25% year over year growth.

Customer demand for higher internet speeds also continues to push higher fueled by service provider competition and newer technologies.

Most of this traffic is currently downstream. Currently 17 times as much data is delivered downstream during the busy hour as upstream. On the downstream side, traffic is predominantly over the top video sites like Netflix, YouTube, and Hulu. Future downstream growth drivers include the transition to higher definition video content like 4K/UHD which is about 3 times as much traffic as a high definition stream. We'll also be fueling some growth as we continue our own IPTV transition.

On the upstream side, traffic is uploading video to sites like Facebook and YouTube. Also, online storage sites like Google and Microsoft show up. Also, Cox Business subscribers drive a disproportionate amount of upstream traffic. (Not so much downstream). Upstream traffic growth drivers will include the growing Internet of Things streaming data to the cloud. With higher definition video content coming from a rapidly growing number of these devices. Ultra-High Def security cameras streaming continuously to the cloud is a trend we're watching.

Given all the network traffic growth drivers, Cox is in the midst of a network transformation that supports gigabit symmetrical speed offerings and manages long-term network congestion. This transformation will scale our network from thousands of nodes, to hundreds of thousands of nodes over the next 10 years. A capital-intensive project of this magnitude requires a comprehensive strategic planning capability that

balances demands from engineering and construction, manages congestion, delivers new product capabilities, incorporates geospatial, subscriber, business value, and budget data; and delivers an actionable plan that ensures we put the right investment in the right nodes at the right time to deliver the most value to the business, on budget.

Summarized Overview

Workflow



Challenges

<ul style="list-style-type: none"> Storage shortage Manual process 	Improve using machine learning Apply AI technology	Manage rules in macros Spreadsheets unable to handle 20K nodes	Limited what-if analysis Unable to analyze sensitivity
--	---	---	---

Solution requirements

<ul style="list-style-type: none"> Big Data platform Real-time data Telemetry from 	Timeseries and Ensemble forecasting Optimization for prescriptive results	Use decision tables Fast updates to rules Multiple scenario runs More analysis time	Enterprise BI platform Multi user support Track forecast Realtime analysis
---	--	--	---

1. Big Data ecosystem

Cox's Enterprise Data Services team provides end-to-end application development and analytics infrastructure mostly using Hadoop ecosystem infrastructure.

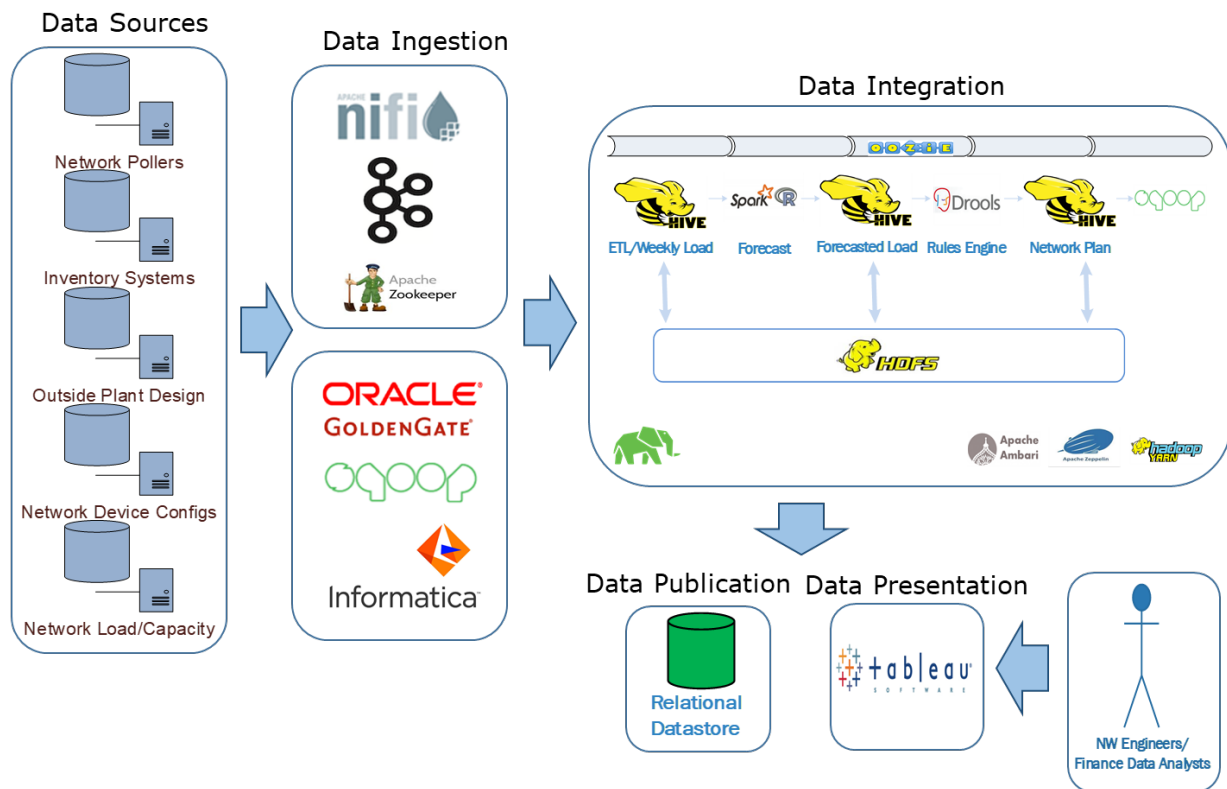


Figure 1 - Network Planning using Big Data Infrastructure

2. Data preparation

A network node-level capacity management and transformation plan requires several types of data such as including network topology, telemetry, physical, geographical, and subscriber data. Extracting, cleaning, preparing, and integrating these large data sets on a continuous basis can be the most time consuming and daunting task. The goal is to clean existing data for errors, combine data feeds from multiple systems, and create a single version of truth for analysis. Sqoop and HIVE are the primary Hadoop tools employed.

3. Analytics

An accurate node-level traffic load forecast is necessary to make savvy long-term business decisions about capacity management and technology transition. Accurate traffic forecasting must integrate long-term market level growth trends along with recent node-level dynamics. Combinations of exponential smoothing (Holt-Winters), ARIMA, and ensemble machine learning methods are compared, and combined, and the most accurate method selected.

Multivariant techniques, such as clustering, are employed with congestion, geospatial, subscriber, business value, and budget data to balance demands from engineering, outside plant construction, product and finance to ensure the right investment is made in the right nodes at the right time to deliver the most value to the business on budget.

Forecasting and clustering analytics automation is supported by Spark R on Apache Hadoop infrastructure.

4. Business Rules

Once the long-term node-level traffic load is determined, complex business rules engage to determine the optimal path for managing the technology and capacity evolution of the node. For example, spectrum updates, node-splits, N+0 build outs, and full-duplex DOCSIS steps must be qualified and sequenced.

Automated business rule application is enabled by the Drools core Business Rules Engine (BRE) using JBoss on the Apache Hadoop platform.

5. Publication

A capital-intensive project of this magnitude has numerous stakeholders, including engineering, outside plant construction, critical facilities, metro and backbone network, finance, field engineering, customer support, and executive steering. Data is transferred to data driven stakeholders via Sqoop. Data visualization is accomplished via Tableau.

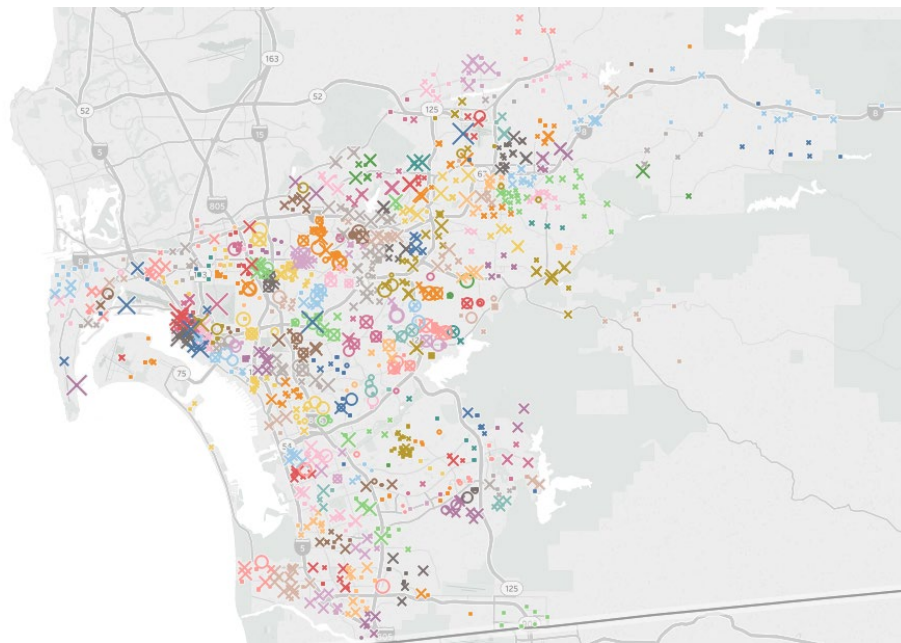


Figure 2 - Network Transformation Visualization of select market

Conclusion

This comprehensive end-to-end solution using a modern Big Data platform produces faster and repeatable results.

- Real-time input data processing and preparation combining multiple network topology and telemetry data sources (Hive, Spark on HDP Platform)
- Data Analysis (Forecasting and Clustering) using an advanced statistical programming language (Spark R on HDP)
- Enterprise business rule engine that can process declarative rule sets (JBoss Drools engine)
- Specialized data visualization dashboards to drill-down and assist detailed analysis (Tableau connected to Hive)

Abbreviations and Definitions

DOCSIS	Data Over Cable Service Interface Specification (DOCSIS) is a telecommunications standard used to provide Internet access via a cable modem. It is important because it is a key element in providing modem manufacturers and network service providers a common method for products to work together in a predictable manner.
Drools	Drools is a business rule management system (BRMS) with a forward and backward chaining inference-based rules engine, more correctly known as a production rule system, using an enhanced implementation of the Rete algorithm.
FDX	Full Duplex DOCSIS 3.1
HDP	Hortonworks Data Platform
HIVE	The Apache Hive™ data warehouse software facilitates reading, writing, and managing large datasets residing in distributed storage using SQL. Structure can be projected onto data already in storage. A command line tool and JDBC driver are provided to connect users to Hive.
JBoss	Also, formerly known as JBoss Application Server, now known as WildFly, is an application server authored by JBoss, now developed by Red Hat. WildFly is written in Java and implements the Java Platform, Enterprise Edition specification. It runs on multiple platforms.
N+0	Node-plus-zero (N+0) architecture. N+0 means there are no amplifiers required between a node and a subscriber household
Spark	Apache Spark is a unified analytics engine for big data processing, with built-in modules for streaming, SQL, machine learning and graph processing.
Spark R	SparkR is an R package that provides a lightweight front end for using Apache Spark from R, supporting large-scale analytics on Hortonworks Data Platform (HDP) from the R language and environment.
Sqoop	Apache Sqoop (TM) is a tool designed for efficiently transferring bulk data between Apache Hadoop and structured datastores such as relational databases.
SCTE	Society of Cable Telecommunications Engineers

Network Programmability – A Reality Check And A Glimpse Into The Future

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

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. Why Network Programmability?.....	3
2. The State of Network Programmability	3
• Programmable Client Interfaces.....	3
• Programmable Line Interfaces	4
• Programmable Optical Carriers/Super-channels	5
• Programmable Service Activation	6
• Programmable and Virtualized Infrastructure	7
• Programmable Telemetry and Network Pulse	7
• Programmability of Service Restoration.....	8
• Programmable Inter-layer Service Setup.....	8
3. Network Programmability – A Glimpse into the Future	9
• Programmable and Highly Granular Optical Transmission.....	9
• Analytics and Machine Learning-triggered Network Capacity	9
• Proactive Network and Traffic Protection.....	9
Conclusion.....	9
Abbreviations	10
Bibliography & References.....	10

List of Figures

Title	Page Number
Figure 1 - Programmable Client Interfaces	4
Figure 2 - Programmable Line Interfaces	5
Figure 3 - Programmable Super-Channels	5
Figure 4 - Programmable Capacity through Sliceable Optics.....	6
Figure 5 - Dynamically Add, Modify, Move and Retire Optical Capacity with SDC	7
Figure 6 - Programmability of Service Restoration	8

Introduction

The fast-paced migration to the cloud, forthcoming 5G deployments and the proliferation of internet-connected devices (IoT) are fueling the evolution to a new era of intelligent networking that uses advanced analytics and machine learning to build self-optimizing, self-healing and highly autonomous transport networks. Moreover, emerging technologies like Blockchain that enable peer-to-peer distributed ledger for open and secure exchanges over the internet are, more than ever, putting the network at the heart of this resolution. Network programmability is one of the key building blocks for such successful evolution. However, its real-life implementation across the various moving parts of the network (network elements, layers, management platforms, data collection tools, etc.) often comes with technological challenges. This paper focuses on the “state of network programmability” today by highlighting where network programmability has been successful and how it translates into business and operational benefits. It also provides a glimpse into the future, including what the industry is currently working on to extend programmability further into the network and the necessary building blocks to do so.

Content

1. Why Network Programmability?

The proliferation of streamed content over the internet, including the IoT with billions of connected devices, new data exchange and validation models like Blockchain, and accelerated enterprise migration to cloud applications, has dictated a new level of task automation and programmability in order to keep up with the constant and unpredictable demand for bandwidth, streamline operations and eliminate sources of human error. In this new era of hyper-connectivity, where content must be delivered to hundreds of millions of users across the globe with the highest levels of quality, or machine-to-machine communications must be automatically and autonomously initiated without human intervention, network programmability becomes paramount.


The underlying network must be highly programmable and span the various constituent parts of the network to automate tasks, optimize network resources (e.g. optical spectrum, capacity/reach ratio) and make real-time network decisions.

2. The State of Network Programmability

Programmability has been making its mark across various and key parts of network components, layers and functions, from service delivery (client interfaces) to core network traffic management. The following highlight where it has been successfully implemented as well as the value it provides:

- **Programmable Client Interfaces:** The early days of transponders used to be service- and protocol-specific; for example, one transponder would be designed to deliver a predetermined and fixed client service such as OC-192, another type of transponder would be for OC-48 and so on. This often led to the quick exhaustion of a network element’s available service slots and required service providers to maintain a large inventory of circuit packs for each service. Transponders are now programmable, so each port can be set, by software, to operate over a specific client service type or protocol and cover a wide range of bit rates. All different ports of a programmable transponder can operate over any supported service type without restriction, allowing service providers to reduce footprint, maximize return on investment and decrease sparring costs. as depicted in Figure 1.

Protocol-specific client interfaces



Programmable client interfaces




Figure 1 - Programmable Client Interfaces

- **Programmable Line Interfaces:** Programmability has also changed how the network is designed, operated and scaled. Early optical line interfaces operated over a fixed frequency, then later they became tunable across the C-band, where a single line card could be used to transmit over any of the 88 C-band fixed-grid frequencies. The latest innovation in software- and hardware-enabled line interfaces is full programmability: not only can they tune over any of the fixed C-band frequencies, but they also operate over any specific modulation schemes (e.g. QPSK, 8QAM, 16QAM) or a specific baud rate (e.g. 17 GBaud, 22 Gbaud, 33 GBaud), as shown in Figure 2.

Programmable Line Interfaces




Figure 2 - Programmable Line Interfaces

- Programmable Optical Carriers/Super-channels:** While DWDM technology disrupted the telecommunication industry by enabling multiple optical carriers to travel in parallel on a fiber to increase fiber capacity, the latest innovation in photonic integration and digital signal processing raised the bar of optical performance and capacity by introducing programmable super-channels. 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 without increasing operational complexity, as depicted in Figure 3. Super-channels are designed to overcome three fundamental challenges: to scale bandwidth without scaling operational procedures, to optimize DWDM capacity/reach and to support the next generation of high-speed services such as 100 GbE and 400 GbE.




Figure 3 - Programmable Super-Channels

Super-channels are designed to be programmable as they support a capability often referred to as “optical sliceability” by allowing any super-channel to be sliced, so each 100G or N x 100G wavelength can be tuned across the C-band, modulated and then routed in a separate direction from the others to the appropriate destination over any open optical line system, as shown in Figure 4. This ability to “slice and dice” super-channels streamlines operations through programmability of the number and capacity of these bandwidth slices, and it also significantly reduces TCO (power consumption and footprint) while increasing network flexibility.




Figure 4 - Programmable Capacity through Sliceable Optics

- Programmable Service Activation:** Activating new capacity is often complicated, requires truck rolls, is prone to human error and takes months and months to complete. The complexity of this process translates into extended time to revenue, which has a direct and negative impact on service providers' top line. Programmability has also been extended to service activation through the implementation of software defined capacity (SDC) that brings the principles of SDN, which has primarily focused on the Ethernet and packet layers, to the optical transport layer to dynamically add, modify, move and retire optical capacity based on the real-time requirements of upper-layer applications. SDC provides instant software activation of additional capacity, creating a pool of bandwidth that can be dynamically allocated based on traffic demand (Figure 5). Software defined capacity/activation is a true game changer from both business and operational perspectives. It enables a perfect match between the timing of CapEx and service revenue, thus accelerating time to revenue from months to minutes. It also reduces OpEx by streamlining operations and eliminating truck rolls.




Figure 5 - Dynamically Add, Modify, Move and Retire Optical Capacity with SDC

- Programmable and Virtualized Infrastructure:** New advancements in software and hardware programmability allow for the creation of virtualized networks at the packet, digital and optical layers and across metro and core domains. Parts of the optical infrastructure can be logically partitioned based on each end customer's (e.g. an enterprise customer's) needs. Premium customers can have dedicated logical partitions of the network with complete visibility and control of their logical network and isolation from other customers, allowing them to customize connectivity and services around their own applications. Similarly, such virtualization capabilities allow network operators to maximize their return on assets and broaden their addressable markets without the additional capital often required to build private physical networks.
- Programmable Telemetry and Network Pulse:** Cloud and IoT networking rely heavily on streaming telemetry and real-time data analytics to assess network health and proactively avoid failures caused by network degradation. This flow of information between the various parts of the network and the upper-layer software tools and SDN controllers is programmable, leveraging open APIs such as gRPC, RESTCONF, NETCONF/YANG and other northbound interfaces to connect to upper-layer orchestration systems or same-layer intelligent software tools. This programmable flow of telemetry allows the coordination and optimization of resources across the network and the delivery of predictive and prescriptive real-time recommendations and actions for maximum performance.

- Programmability of Service Restoration:** Network downtime often translates into the loss of millions of dollars as well as irreparable damage to the network operator's reputation and brand image. A network failure can have a disastrous impact on cloud applications where the network plays a vital part in connecting users to the cloud or in machine-to-machine connectivity. Network operators often leverage intelligent software tools such as control planes to operate as the brain of the network, reacting to network changes in real time, without human intervention. These intelligent software tools have evolved to be programmable to increase network availability and protect against failures (e.g. fiber cuts, hardware fa) while maintaining stringent service requirements such as low latency, avoidance of shared risks (e.g. restored traffic going through the same impacted fiber conduit), etc. A programmable control plane can partition a network carrying mission-critical data by allowing specific links, wavelengths, subwavelengths or even nodes to be dedicated to a specified use, with preset thresholds for latency, bandwidth and resiliency. The programmability of service restoration makes networks autonomous, highly reliable and self-healing (Figure 6).

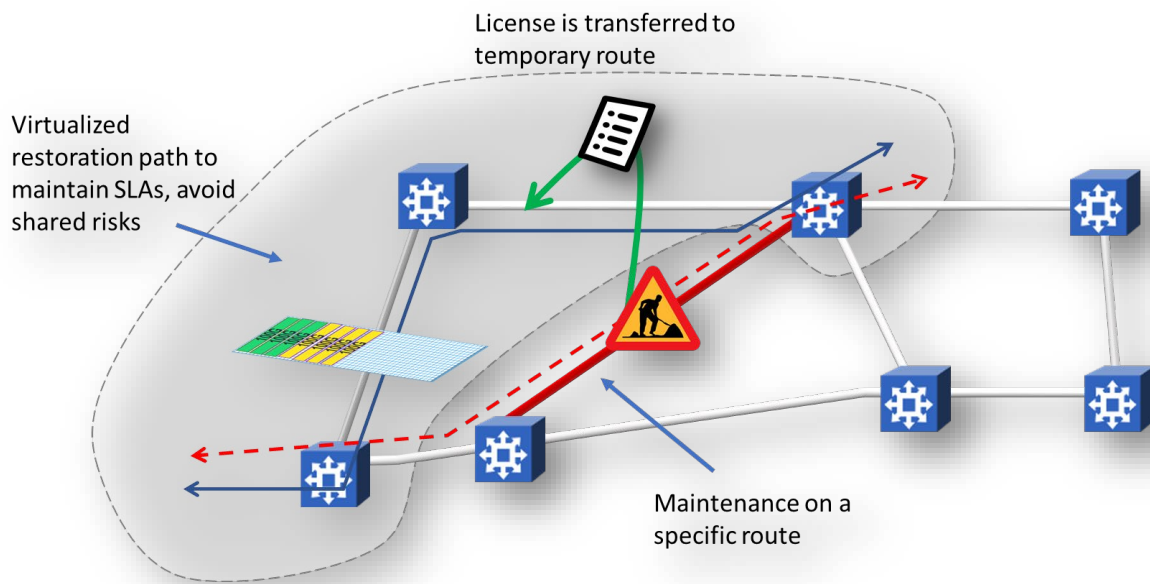


Figure 6 - Programmability of Service Restoration

- Programmable Inter-layer Service Setup:** Building automation and intelligence across all network layers and various operating tasks is central to underpinning cloud-era networking. SDN controllers and frameworks span network layers to provide programmable and intelligent capabilities to plan, monitor and conduct various network operations, such as real-time capacity planning, bandwidth on demand, network virtualization and many others without human intervention. For example, very sophisticated algorithms and data models can be used to build a microservices-based PCE. The PCE replaces manual offline route and capacity planning processes with highly automated, programmable, real-time service planning and activation over optimal routes across layers, to overcome multiple and often challenging fiber impairments.

3. Network Programmability – A Glimpse into the Future

As discussed earlier, programmability has been implemented across various and key parts of the network, including its components, layers and functions. Nonetheless, it's still poised to further expand and become entrenched in the next generation of connectivity, often referred to as cognitive networking. The following describe how programmability will touch many other aspects of networking, paving the way for highly automated, self-aware, self-organizing and self-optimizing networking infrastructure:

- **Programmable and Highly Granular Optical Transmission:** A practical way to implement higher-order modulation schemes (e.g. 128QAM) and baud rates (e.g. 100 GBaud) in optical transmission systems without a massive compromise in optical performance will certainly emerge. However, advancements in photonic integration and digital signal processing are leading toward a more programmable and highly granular optical WDM line rate through the implementation of hybrid modulations of super-channels and subcarriers. With the ability to program different modulation and dynamically adjust constellation shaping gains on each subcarrier, a variety of spectral efficiencies can be derived [1][2]. This provides higher fiber capacity and a better flexibility for diverse network applications.
- **Analytics and Machine Learning-triggered Network Capacity:** While SDC has already forever changed the way services are planned, provisioned and activated, its next phase will allow real-time network analytics, microservice-based engines and machine-learning algorithms to dynamically increase or decrease network capacity on specific routes based on past trends, spontaneous changes in traffic demand or an anticipated spike in capacity. This data-driven real-time traffic engineering will be fully autonomous, programmable and proactive to identify potential sources of failures before they happen and take the necessary steps to maintain the network at the highest levels of reliability and efficiency.
- **Proactive Network and Traffic Protection:** The use of control plane has significantly increased service and network availability by automatically restoring traffic after a failure over alternative paths while maintaining SLAs. Moreover, Layer 1 or Layer 2 traffic encryption embedded into optical transmission platforms has proven to be an effective way to protect in-flight data from intruders and hacking tools. The evolution of network and traffic protection could be the combination of control plane capabilities, real-time traffic and network topology engineering as well as traffic encryption - all functioning as parts of an intelligent and highly proactive network protection mechanism. Advanced, accurate and fast intrusion detection capabilities could trigger proactive measures to virtually isolate suspicious network areas and automatically encrypt the traffic over specific links. The same proactive capabilities can also be applied to minimize the impact of network maintenance operations by automatically and proactively isolating affected network areas and setting up backup plans during large-scale software upgrades or disruptive network operations.

Conclusion

Network programmability is one of the key building blocks for a successful and fast-paced evolution to the cloud and Blockchain. It's been entrenched across various moving parts of the network, from programmable optical client and line ports to highly sophisticated service creation and activation mechanisms, and it has proven to lead to significant business and operational benefits. Despite such progress, network programmability is still poised for further expansion and evolution by taking advantage

of the recent development in analytics and machine-learning tools to elevate the network to a new level of automation, flexibility and efficiency.

Abbreviations

API	application programming interface
CapEx	capital expenditure
DWDM	dense wavelength-division multiplexing
GbE	Gigabit Ethernet
gRPC	gRPC Remote Procedure Call
IoT	Internet of Things
NETCONF	Network Configuration Protocol
OC	optical carrier
OpEx	operational expenditure
OTN	Optical Transport Network
PCE	path computation engine
QAM	quadrature amplitude modulation (8QAM, 16QAM, 64QAM)
QPSK	quadrature phase-shift keying
REST	Representational State Transfer
SDC	software defined capacity
SDN	software-defined networking
SE	spectral efficiency
SLA	service level agreement
TCO	total cost of ownership
YANG	Yet Another Next Generation

Bibliography & References

[1] *Design Considerations for a Digital Subcarrier Coherent Optical Modem*; David Krause, Ahmed Awadalla, Abdullah S. Karar, Han Sun, Kuang-Tsan Wu, OCF 2017

[2] *Subcarriers having different modulation formats*; US patent pending; A. Awadalla, et al.

New Packet Network Design for Transporting 5G Fronthaul Traffic

A Technical Paper prepared for SCTE•ISBE by

Brian Lavallée

Senior Director, Portfolio Marketing

Ciena Corporation

2351 Boulevard Alfred-Nobel, Saint-Laurent, Québec, Canada, H4S 2A9

(514) 228-2300

blavalle@ciena.com

In memoriam to my esteemed colleague, Jim Kleinsmith.

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. Distributed Radio Access Network (D-RAN)	3
1.1. Backhaul Network	3
2. Centralized Radio Access Network (C-RAN)	4
2.1. Fronthaul Network.....	5
2.2. Opportunities.....	5
2.3. Challenges	5
3. 5G Mobile Networks.....	6
3.1. 4G and 5G Coexistence	6
3.2. 5G Fronthaul	6
3.3. IEEE 194.3 Radio-over-Ethernet (RoE) Encapsulation.....	7
3.4. IEEE Time-Sensitive Networking (TSN)	7
4. Converged Haul Transport.....	8
Conclusion.....	9
Abbreviations.....	10
Bibliography & References.....	10

List of Figures

Title	Page Number
Figure 1 - Multiple generations of mobile network technology deployed (source: GSMA)	4
Figure 2 – Backhaul and Fronthaul Networks (ref: EXFO)	5
Figure 3 – Existing 4G C-RAN vs. New 5G C-RAN Configuration Comparison	7
Figure 4 – Time-Sensitive Networking (TSN) Standards (ref: IEEE)	8
Figure 5 – Proposed Functional Splits (ref: 3GPP).....	9

Introduction

5G is not just about upgrading the handsets, radios, and antennas that together comprise the Radio Access Network (RAN). Offering 5G mobile services also requires substantial upgrades to packet-optical wireline networks that connect cell sites to each other and to data centers hosting accessed content, and everything in between. This means that for Mobile Network Operators (MNO) to achieve the 5G improvements over 4G LTE of 100x more devices, 100x faster data rates, 10x lower latency, and 1000x higher data volumes, everything in the end-to-end mobile service path must eventually be scaled and modernized. This applies to connect, storage, and compute resources resulting in multi-year modernization journey that will start in the RAN and network edge and steadily move inwards, which has already started in several countries.

Unlike previous introductions of mobile networking technology (2G, 3G, 4G) where the new generation of was intended to replace the old generation – but never did – 5G is not intended to replace 4G. 5G is intended to complement and coexist alongside 4G (and 2G and 3G mobile services in many cases) meaning they must coexist by sharing as much connect, storage, and compute resources as possible if MNOs are to support multiple generations of mobile services in a cost-effective manner. 4G will also continue to evolve over time from existing Long-Term Evolution (LTE) deployed today, to future LTE Advanced and LTE Advanced Pro, which are enhancements to 4G that bring it closer to expected 5G performance.

Holistically speaking, a mobile network is a massive wireline network with radios hanging off its edges. In most cases, the only wireless part of the end-to-end journey of data flowing between users and accessed content is between from User Equipment (UE) and cell site antennae. The rest of the end-to-end journey is predominantly over packet-optical wireline networks, although wireless backhaul does exist.

In short, this means the move to offering 5G mobile services is about far more than just a wireless upgrade.

Content

1. Distributed Radio Access Network (D-RAN)

Traditional mobile networks were designed with multiple Radio Heads (RH) and Baseband Units (BBU) installed in the same location called a *macro cell site* or *cell site*. RHs were installed atop a tower, with each serving a sector of 120 degrees in the common 3-sector configuration. Early connections between RHs and BBUs was over electrical media (copper). The distance between RHs and the BBU installed at the base of a tower is typically around 200 to 400 feet or so in distance, which determines propagation latency.

Electrical connections between RHs and BBUs led to high electrical power consumption and associated energy costs. It also meant being susceptible to environmental conditions (lightning), Electromagnetic Interference (EMI), and Electromagnetic Conductance (EMC). These macro cell sites comprised of RHs and BBUs were constructed in a distributed manner intended to serve subscribers within a typical radius of around 20km to 30km. This network topology is referred to as Distributed RAN (D-RAN) and has been the primary method of deploying macro cell sites in most mobile networks around the world.

1.1. Backhaul Network

The network connection between D-RAN cell sites and the MNO Mobile Telephone Switching Office (MTSO) is called *backhaul*, since traffic from the former is *hauled back* to the latter. As newer generations of wireless technology offered faster speeds over the airwaves, alongside an increased number of subscribers, backhaul traffic soared, and it was realized that legacy, copper-based backhaul technology simply could not maintain pace. This is precisely why packet-optical technology became, and continues to

be, the best option for high-capacity, low-latency, and major economies of scale for mobile backhaul networks. Packet switching, transported over underlying optical technology, offers benefits associated with statistical multiplexing. The main benefit yielded is optimized bandwidth utilization for reduced costs and is why packet switching technology is ubiquitous in most parts of the global network infrastructure, from edge to core.

Most mobile networks were constructed using D-RAN throughout the world. As new generations of mobile technology were developed, new radios and antennas were installed on existing towers alongside previous generations of radios and antennas. This is because MNOs were unwilling (or unable) to turn off previous generations of mobile services because new generations of mobile services required new radios and antennas at both cell sites and within handsets of subscribers. This is illustrated below showing the mix of 2G, 3G, and 4G mobile network technology deployed around the world today, and into the future.



Figure 1 - Multiple generations of mobile network technology deployed (source: GSMA)

There is, and will continue to be, a mix mobile network technology, and is precisely why adding 5G needs must be seamless and cost-effective; easier said than done, no doubt. In most developed countries, 2G mobile services have already been, or will soon be, decommissioned. However, 2G will have a long life in many countries, as will 3G and 4G for the foreseeable future. This is why MNOs understand and demand that 5G not be intended to outright replace previous generations of mobile network technology. It also means that a single, converged infrastructure, wherever and whenever possible, is an obvious primary goal.

Not all use cases require 5G performance and is one of the reasons why sub-generations of 4G LTE technology (LTE Advanced, LTE Advanced Pro) are actively being deployed with vendors continuing to invest in associated product roadmaps for many years to come. Although multiple generations of wireless technology can and will coexist, multiple wireline overlay networks are simply too costly and complex. This is why there is a pressing desire to converge different generations on a converged wireline network.

2. Centralized Radio Access Network (C-RAN)

As mentioned above, initial D-RAN deployments connected multiple RHs atop a tower to BBUs at the foot of the tower using electrical technologies. Although this configuration served the industry very well for many years, optical networking technology has steadily advanced with notable leaps in performance and cost-effectiveness when compared to its copper-based brethren. Optical fiber-based media is also far less susceptible to environment conditions, which is another notable advantage. This has resulted in electrical connections between macro cell RHs and BBUs to be steadily replaced by fiber optics over time.

Optical fiber-based communications enable much farther propagation distances than electrical copper-based communications, and this was not lost on MNOs and equipment vendors alike. Why not move and centralize multiple geographically dispersed macro cell BBUs into one location and then connect to Remote RHs (RRH) over distances afforded by fiber optics? This led to *fronthaul*, which is the connection between

centralized BBUs and geographically separated RRHs. BBU functions are increasingly being virtualized and are moving into data centers leading a cloud-based C-RAN. Although C-RAN was first applied to 4G, it is a prime candidate for 5G as well, given that the latter will leverage the higher frequency millimeter wave spectrum. Propagation in this part of the spectrum yields shorter distances and more difficulties through obstacles resulting in a reduced coverage area. This means wide-scale 5G service coverage requires a significant densification of cell sites closer to subscribers, and more fiber to connect to these sites.




Figure 2 – Backhaul and Fronthaul Networks (ref: EXFO)

2.1. Fronthaul Network

The two main 4G fronthaul protocols are Common Public Radio Interface (CPRI) or OBSAI (Open Base Station Architecture Initiative), although the former is far more widely deployed than the latter. CPRI is not a formal industry standard; rather, it's a *public specification* that's been implemented in such a way that interconnecting RRHs to centralized BBUs from different vendors is challenging at best, and in most cases, simply impossible. Although CPRI works, and is deployed, MNOs are locked into a single vendor.

2.2. Opportunities

There are many advantages to C-RAN. This is why MNOs are increasingly investigating this relatively new configuration. For example, having multiple RRHs serving a broad geographic coverage area connected to centralized BBUs simplifies implementing Coordinated Multi-Point (CoMP), cooperative beamforming, and enhanced Inter-Cell Interference Coordination (eICIC), which are part of LTE Advanced. Moving once geographically dispersed BBUs into a centralized location allows for greater economies of scale leading to a lower cost RAN to own and operate. C-RAN facilitates hosting of virtualized mobile network functions (Serving Gateway, Packet Gateway, Home Subscriber Server...) of the Evolved Packet Core (EPC) by leveraging ongoing data center technology advances related to both storage and compute.

2.3. Challenges

We live in a world of compromise, and the adoption of C-RAN is no different. Although there are several advantages to connecting RRHs to BBUs, the assumption is that optical fiber is available. In many cases,

optical fiber is already available between macro cell sites and the MTSO used for backhaul purposes so adding RRHs to these existing cell sites and moving the BBUs into the MTSO is greatly facilitated. The challenge is related to maximizing the use of existing fiber, especially as some traffic carried on this fiber will be 2G/3G/4G D-RAN backhaul traffic and 4G/5G C-RAN fronthaul traffic, as multiple generations of mobile network technology are expected to coexist for many years to come. New RRH cell sites will require new fiber optic availability, which conjures up major challenges related to permits, rights-of-way, and the cost and time implications of trenching these fiber optic connections.

Another key challenge associated with C-RAN is that the original electrical connection between a RRH and BBU was designed from inception for a propagation distance, which dictates latency, as high as 400 feet. The upper limit of CPRI-based fronthaul is around 200us, which includes the latency associated with the propagation of light and latency incurred as CPRI traffic traverses intermediate network elements. Although the maximum distance between RRHs and BBUs in 4G C-RAN is approximately 20km, it is deployed over just a few kilometers, in most cases. Stringent CPRI latency limits coupled with the cost and right-of-way challenges associated with gaining access to optical fiber to connect to new RRHs in the quest to cell site densification has significantly limited wide-scale 4G C-RAN deployments, at least for now.

3. 5G Mobile Networks

5G promises 4G LTE improvements of 100x more devices, 100x faster user (man and machine) data rates, 10x lower latency, and 1000x higher data volumes. To achieve these aspirational goals, fiber and cell site densification will be required, along with the adoption of many new and emerging technologies. 5G will leverage as much of the existing packet-wireline network infrastructure, where possible, in the early stages to simplify and cost-reduce early 5G rollouts. This is evidenced by MNOs attaching 5G New Radios (NR) to the existing 4G EPC, referred to as the Non-Standalone (NSA) configuration, and is an elegant way to test and prove the performance of the 5G NR products, before wide-scale deployments can commence.

3.1. 4G and 5G Coexistence

From inception, 5G is not intended to outright replace 4G. This has profound consequences on the wireline network that connects 4G and 5G cell sites to each other and to data centers where access content is hosted. These data centers offer storage and compute resources and can be located anywhere from the base of a cell site tower to thousands of kilometers away, and anywhere in between. Moving the storage and compute resources closer to the network edge has led to such industry initiatives as Multi-Edge Computing (MEC). The location of MEC resources will be dictated by the applications and use cases they are expected to support leading to challenges for MNOs related to deciding where to place storage and compute resources. As virtualization continues to evolve, the ability to dynamically relocate resources is greatly facilitated by providing increased flexibility to dynamically orchestrate storage, compute, and connect resources.

3.2. 5G Fronthaul

CPRI was designed for 4G and simply cannot scale to expected 5G rates in its current form. This has led to the development of enhanced CPRI (eCPRI) targeted at 5G C-RAN. Standards-based transport of eCPRI traffic between RRHs and centralized BBUs is required, and must be open, scalable, and cost-effective. Although there are different ways of transporting this new type of traffic, such as Passive Optical Networks (PON) technology, Ethernet has once again come to the forefront as the protocol of choice for carrying all kinds of traffic, which has resulted in its near ubiquity. However, traditional best-effort Ethernet will not suffice given the latency-sensitive nature of 4G and 5G fronthaul traffic, so enhancements are necessary.

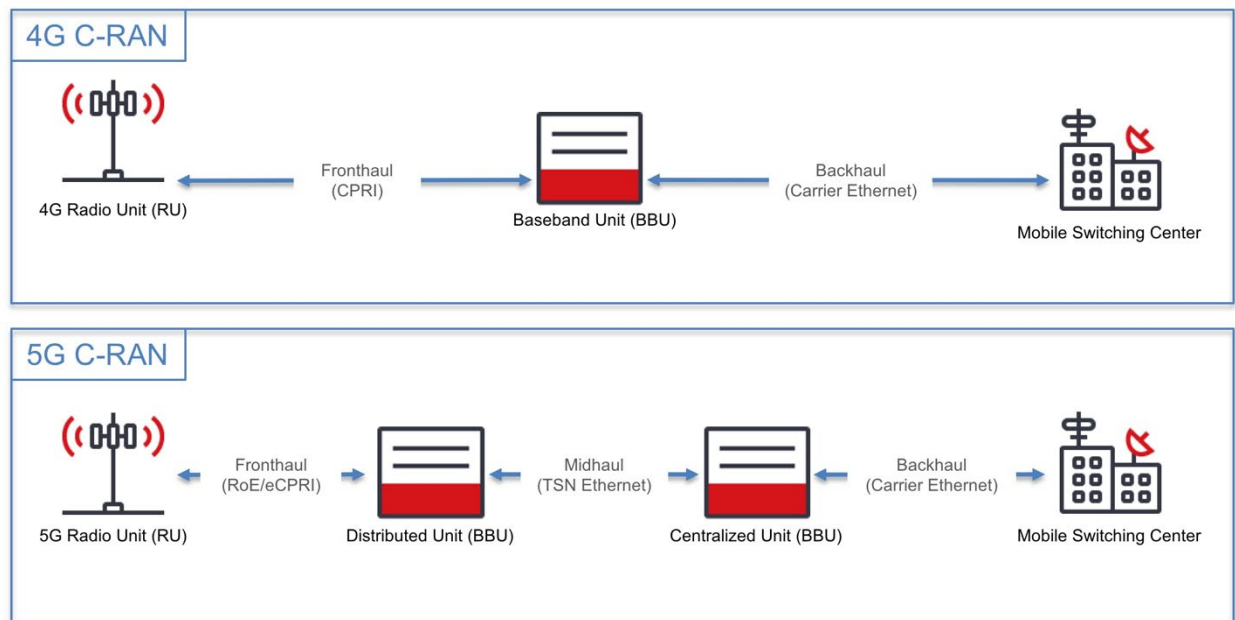


Figure 3 – Existing 4G C-RAN vs. New 5G C-RAN Configuration Comparison

3.3. IEEE 1914.3 Radio-over-Ethernet (RoE) Encapsulation

The IEEE 1914.3 standard defines how radio information, both data and control, is mapped into Ethernet frames using standardized Radio-over-Ethernet (RoE) headers. The standard supports the encapsulation of time-domain IQ (4G CPRI) and frequency domain IQ (5G eCPRI). Once radio information is packetized, it needs deterministic transport network mechanisms to ensure bounded latency and zero packet loss.

3.4. IEEE Time-Sensitive Networking (TSN)

There are three ways in use today to transport packet traffic. The first way is Constant Bit-Rate (CBR) that leverages legacy SONET/SDH or modern OTN to carry packet traffic offering such connection-oriented advantages as constant low latency and zero packet loss, albeit at the expense of locking of capacity whether it is being used or not. The second way is via traditional, highly cost-effective Ethernet leveraging statistical multiplexing for connectionless, best-effort transport resulting in less predictable latency and non-zero packet loss. The third way combines these two via deterministic Time-Sensitive Networking (TSN) that offers the best of both worlds, such as fixed paths for tightly bounded latency and zero packet loss.

TSN is not a new protocol; rather, it is a standards-based enhancement to traditional, IEEE standards-based Ethernet that ensures data can travel from network ingress to network egress in a highly predictable amount of time offering similar performance to Time-Division Multiplexing (TDM) options, such as OTN, albeit at a lower cost and complexity. The zero-packet loss and tightly bound latency capabilities directly address the latency sensitivity associated with CPRI and eCPRI-based fronthaul traffic between RRHs and BBUs.

Although deterministic traffic flows can be created using other technologies, such as MPLS, they don't offer the low latencies required 5G fronthaul. TSN gets right down to how packets are queued within the switch and how they're allowed to block or not block each other. While MPLS can carve a path through the network, traffic is still queued and buffered along the way and thus doesn't provide as tight controls as the Link Layer 2 techniques available with TSN, which is one level of tighter control, at the bit level.

The IEEE TSN Working Group has created standards to enhance existing, best-effort Ethernet such that it can properly support deterministic networking applications, such as fronthaul. These standards fall into four main categories of functionality; (1) Synchronization, (2) Reliability, (3) Resource Management, and (4) Latency, as illustrated in Figure 3.

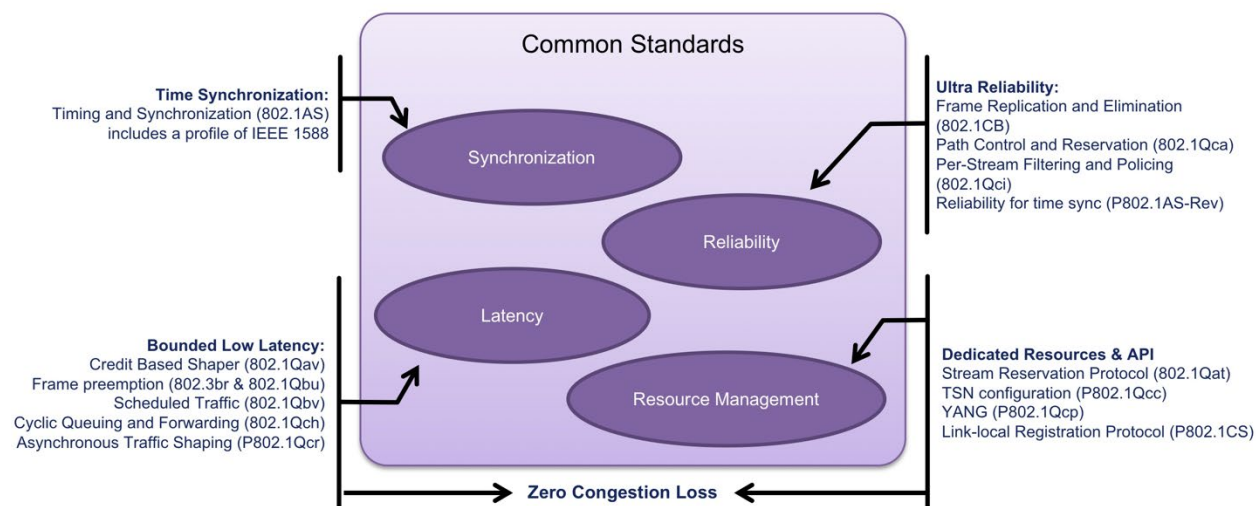


Figure 4 – Time-Sensitive Networking (TSN) Standards (ref: IEEE)

TSN has been used for decades in such applications as industrial Ethernet, audio-visual, and power grid automation. TSN-based Ethernet achieves deterministic transport of 4G CPRI and 5G eCPRI by controlling timing synchronization, traffic scheduling (forwarding, queuing), and system configuration of all nodes in the end-to-end traffic path. There are multiple IEEE standards associated with the TSN Working Group related to synchronized network elements, controlled/accountable latency, prioritization of different traffic classes (deterministic and non-deterministic), guaranteed bandwidth reservation, and enhanced redundancy and resiliency. These enhancements to standards-based Ethernet make it a prime candidate for 5G fronthaul transport, and since it is based on open, well-understood, and field-proven standards, 4G C-RAN fronthaul vendor lock-in is significantly reduced via a broader, open, and more secure vendor ecosystem.

4. Converged Haul Transport

Converged Haul transport refers to a common physical network infrastructure carrying multiple generations of backhaul traffic, fronthaul traffic, and what is being called *midhaul* traffic, which is related to a variety of *functional split* proposals of 5G fronthaul traffic. Fronthaul *High Layer Split (HLS)* options are best served by either IP routers or Ethernet switches, while *Low Layer Split (LLS)* options are best served by the better performance of Ethernet. By converging all traffic types *hailed* to and from the RAN via a converged packet-optical wireline infrastructure, MNOs benefit from increased economies of scale by reducing costly overlay networks for a simpler network to design, deploy, and maintain. Overlay networks are unnecessary.

Migration is underway with 5G NRs attached to existing 4G wireline infrastructure, but for 5G to reach its full promise, the wireline network must undergo significant modernization in terms of standards-based fronthaul transport topologies, increased scalability, fiber and cell site densification, virtualization, and the guaranteed end-to-end service performance enabled by Network Slicing.

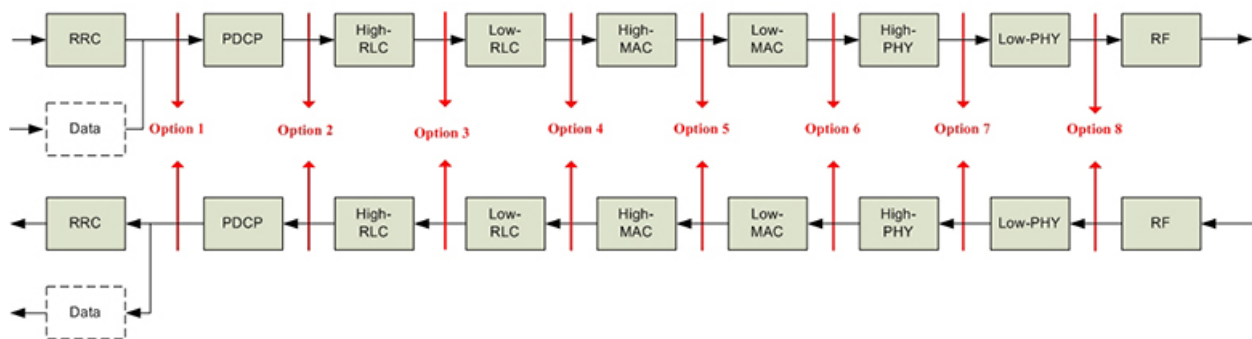


Figure 5 – Proposed Functional Splits (ref: 3GPP)

5G performance gains dictate that traditional network designs must be reevaluated and changed if the full promise of 5G is to be delivered to the masses, man and machine. For example, 5G network slicing will guarantee end-to-end performance across storage, compute, and connect (wireless and wireline domains), which is a monumental change from existing best-effort 4G networks. 5G also touts end-to-end latency of 10ms or less, which is in stark contrast to existing 4G network latency of hundreds of milliseconds.

5G requires software platforms for a virtualized and distributed architecture that pushes intelligence and functionality to the network edge to serve new and unique 5G use cases, such as self-driving cars. A highly virtualized and distributed core network is managed end-to-end by leveraging orchestration and analytics resulting in a more *adaptive network* that can self-configure, self-optimize, and even self-heal in a far more autonomous manner, compared to existing 4G networks, to best address ever-changing network conditions.

Conclusion

Mobile network technology, designs, and mindsets used for decades must be challenged and changed if the full promise of 5G is to be delivered and successfully commercialized. MNOs already know this and are actively developing and executing upon different strategies today. The 5G NR NSA specifications were recently standardized allowing MNOs to test the 5G NR technology in field trials and proofs-of-concept by connecting them to the existing 4G core wireline network. As MNOs gain increased confidence in new 5G NR wireless technology, and as 5G handsets are rolled out, major upgrades will occur in the RAN and the end-to-end wireline network, starting with the fronthaul, backhaul, and new midhaul network segments.

Fronthaul is a new battleground with a variety of proposed functional splits being debated in the industry because MNOs want to migrate away from closed, proprietary solutions to open, standards-based solutions. Ethernet transport is the frontrunner, especially when enhanced with TSN capabilities, and will allow MNOs to exploit the many benefits of this ubiquitous transport protocol that has permeated essentially all parts of the global network infrastructure – *why should the fronthaul and midhaul be any different?*

As 4G and 5G are expected to coexist, fronthaul working groups and associated standards will facilitate carrying 4G CPRI and 5G eCPRI over a common Ethernet-based wireline architecture that can also be used to carry backhaul, and the new midhaul traffic as well. This is the industry's chance to develop and deploy fronthaul networks based on open, field-proven, and standards-based technology – *the time to act is now!*

5G is so much more than just a wireless upgrade – the entire end-to-end network must be considered.

Abbreviations

2G	2 nd Generation Mobile Networks
3G	3 rd Generation Mobile Networks
4G	4 th Generation Mobile Networks
5G	5 th Generation Mobile Networks
BBU	Baseband Unit
CAPEX	Capital Expenditures
CBR	Constant Bit-Rate
CoMP	Coordinated Multi-Point (CoMP)
CPRI	Common Public Radio Interface
C-RAN	Centralized/Cloud Radio Access Network
D-RAN	Distributed Radio Access Network
eICIC	enhanced Inter-Cell Interference Coordination
eCPRI	Enhanced Common Public Radio Interface
EPC	Evolved Packet Core
IEEE	Institute of Electrical and Electronics Engineers
MNO	Mobile Network Operator
ms	milliseconds
NR	New Radio
NSA	Non-Standalone
OBSAI	Open Base Station Architecture Initiative
OPEX	Operational Expenditures
OTN	Optical Transport Network
PON	Passive Optical Network
RAN	Radio Access Network
RH	Radio Head
RoE	Radio-over-Ethernet
RRH	Remote Radio Head
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
TDM	Time-Division Multiplexing
TSN	Time-Sensitive Networking
us	microseconds

Bibliography & References

The Mobile Economy, 2017, GSMA
Radio Access Architecture and Interfaces, TR 38.801 v0.2.0, 3GPP
Mobile Fronthaul Testing Reference Guide, 2016, EXFO

Node Provisioning and Management in DAA

What Changes Are Coming for Techs

A Technical Paper prepared for SCTE•ISBE by

Robert Gaydos

Fellow

Comcast

1701 JFK Boulevard, Philadelphia PA

215-286-8737

Robert_gaydos@cable.comcast.com

Mehul Patel

Senior Principal Architect

Comcast

183 Inverness Drive West, Englewood, CO, 80112

303-658-7826

Mehul_Patel@cable.comcast.com

Joe Solomon

Principal Engineer

Comcast

401 Wynkoop St Ste 300, Denver, CO 80202

303-242-7037

Joe_Solomon@cable.comcast.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Overview of a DAA Control and Data plane Architecture	4
DAA-related Challenges Facing Field Employees	5
1. Network Access and Steering	5
2. Configuration	6
3. Service Routing	6
4. Node Management.....	6
Core Configuration	7
1. GCPP Configuration.....	7
1.1. Video Configuration.....	9
1.1.1. Logical Nodes	9
1.2. Linear Service Group	9
2. Aux Core Configuration	10
2.1. vCMTS Configuration Collection	10
2.2. Node ID to RPD Configuration Association – An App for That!	11
Node Management.....	11
1. Software Version Management.....	11
2. Secure Shell (SSH) Lockdown	12
3. Controlling Multicast Video by SDN Controller	12
4. How Do You Know It Worked.....	12
Conclusion.....	13
Abbreviations	13

List of Figures

Title	Page Number
Figure 1 Comcast’s DAA Architecture	4
Figure 1 - GCPP Event Flow.....	8

List of Tables

Title	Page Number
Table 1 – GCPP Common Configuration	8

Introduction

In traditional HFC deployments, the headend technician controls the services that are delivered to a fiber node because this is done via RF combining in the headend. CMTSs and Edge QAMs output RF channels at configured channel frequencies. This output is split and combined such that the appropriate services are delivered to the lasers going to the node; once the wiring in the combining network is complete, it is rarely changed.

By contrast, and as HFC infrastructure evolves toward a Distributed Access Architecture (DAA), consequent RPHY (Remote PHY) node combining is done virtually, with software. Each node must be virtually directed to appropriate service “Cores,” expressed by frequency plan per service, and video multiplexes to be joined. Service cores span DOCSIS flows, linear and on-demand video, and legacy, out-of-band information.

More specifically, each node must be software configured to listen to its appropriate QAM broadcast and VOD feeds, as well as connect to the correct CMTS, and legacy Out-of- Band (OOB) components.

These advances mean that the industry’s technical workforce needs to be able to program the node, along with Cores which also have to be configured to provide the right data. Pre-planning which physical node (and MAC address) will be installed at a given fiber location is improbable, because line technicians typically carry many nodes. As a result, mechanisms to map logical nodes with intended service configurations and physical instantiations are required.

At the same time, an increasingly intrinsic design goal for network design is to prevent “vendor lock in.” In order to encourage a competitive cost and innovation environment, operators prefer and require multiple sources of components and to be able to pivot to new resources easily. In addition operators must deal with the realities of supporting differing QAM video conditional access (CAS) systems throughout their footprint.

If an operator the size of Comcast used a CMTS to provide all services and manage all aspects of a DAA node, for instance, it could yield as many as 18 permutations of nodes, CMTS, and CAS to test and integrate. Clearly, this is not sustainable. This led to the notion of applying separate “Cores” for broadcast video, VOD, out of band (OOB), and high speed data flows. Cores allow for best-of-breed product selection. Also, keeping video out of the HSD cores simplifies CMTS operations; more importantly, any call to pivot to a new CMTS, or multiple CMTS providers, would sidestep the need to re-integrate video services across six permutations of nodes and CAS systems.

The work related to disaggregating service flows into multiple Cores presented the next dilemma: Deciding which “Core” to make the “primary,” or lead coordination Core. This led to the creation of a “primary core” that is, in essence, a vendor-independent orchestrator. This allows us to mix Cores and nodes at will, and to use our own internal software management processes and tools when needed. We call this software “GCPP,” which stands for “Generic Configuration Protocol Principle.” This internally-developed software performs the following functions, which will be discussed in the paper:

- Network access and steering – aligning the DAA node on the network with the appropriate Cores for configuration
- Configuration – providing QAM Video, VOD, legacy out-of-band configurations and other non-DOCSIS functions
- Service routing – orchestrating the multicast IP routing of video and out-of-band content to the DAA node across the IP network

- DAA node management – managing the versions of software and the code’s signed certificates used by DAA nodes in the field

Yet another design goal was to do as little upfront design as possible, relying on auto-discovery instead of complex, pre-drawn wiring diagrams to connect nodes with switches and photonic muxes. This relates to the concept of the logical vs. physical node. The logical node has a name or ID known to billing systems and GIS systems; it has a known channel map and frequency plan. The physical node is the hardware that hosts the logical node. The physical node can be replaced because of hardware failure or natural disaster. Our system uses “late binding” of physical to logical node mapping, meaning that it happens at the time of install (via an app), thus allowing any node in inventory to host the logical node. (This alone vastly simplified construction processes.) In addition, we tried to ensure that the node’s connectivity to a CMTS would be detected, rather than designed. This makes capacity management easier, ensures the databases are up to date, and gives greater visibility into the network for technicians..

This paper describes the software infrastructure used to manage the thousands of nodes that will be transitioning to DAA. The solution makes use of software defined networking (SDN), distributed cloud servers, and multiple Cores.

Overview of a DAA Control and Data plane Architecture

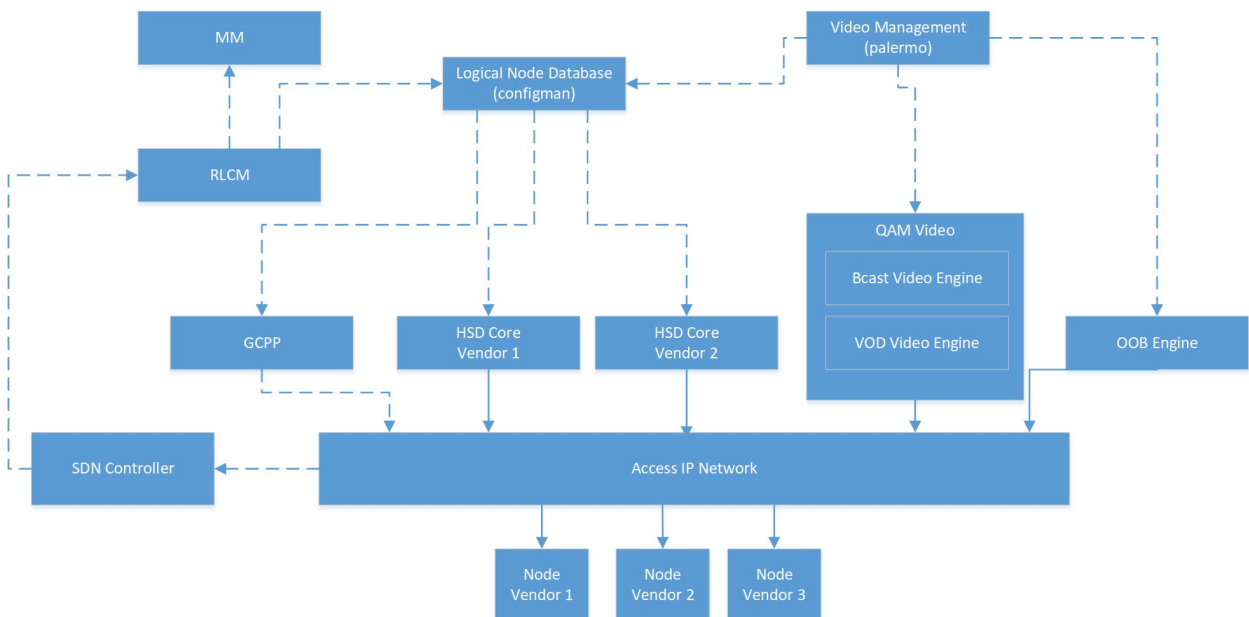


Figure 1 Comcast’s DAA Architecture

Figure 1 illustrates the DAA architecture we developed to overcome the challenges covered in this paper. While we don’t envision running Cores from different vendors in the same sub-section of the plant, it is possible. In fact, it is highly probable that nodes from multiple vendors will exist in the same sub-section of the plant. The following is a brief description of the components:

- SDN Controller – A software-defined networking controller that programs the access IP fabric for both multicast and unicast, detects new and dropped hosts in the network, and heals faults in the fabric.

- Access IP Fabric – A group of whitebox switches comprised of redundant leaf and spine switches that connect to a DAAS (Distributed Access Architecture Switch), which aggregates 10GE connections to remote nodes. There is at least one access IP network per headend and on average 10 individual networks controlled by their own SDN controller. Each network serves approximately 12,000 HHP (households passed.)
- GCPP – General Config Protocol Principle, an internally-developed software component that configures the RPD, receives alerts, and refers the RPD to auxiliary Cores, such as HSD Cores and possibly legacy OOB Cores¹ (55-2). GCPP receives dynamically-generated configurations for each node.
- RLCM – RPD LifeCycle Manager. This internally-created component is the workflow engine that performs the late binding configuration of nodes. It receives notifications from the SDN controller when a node joins the access IP network, which includes the switch and port number. The RLCM will also trigger a “match maker” function to pick an available V-CMTS Core and configure it. RLCM will also communicate with a server for 802.1x authentication, as well as an inventory management system to ensure the device isn’t cloned or otherwise outside of our equipment inventory.
- Logical Node Database – hosts the frequency map (both HSD and video) for a logical node. Maps the video multicasts to static pseudo-wires -- for example tunnels and the frequencies that they are placed on.
- Video Management – “Palermo” is another internally-created tool that tracks desired video configurations on a per node basis. It also must be aware of or program the configurations of video and OOB engines.

DAA-related Challenges Facing Field Employees

1. Network Access and Steering

Each DAA node connects to the CMTS in the headend via Ethernet optics, rather than the traditional HFC network. While 802.1x-based network authentication verifies that the remote-PHY device (RPD) has valid, CableLabs-provided certificates, it does not ensure that the device is one from the operator’s inventory, versus a cloned device. Additional measures to authenticate devices connected to the operator’s IP network will be needed. As mentioned earlier, we implemented an app that binds a physical node to its logical configuration. No device is allowed to enter the network, meaning that the 802.1x request will not be granted and the port will remain “off” if this binding is not done. Furthermore, any attempts to bind multiple physical nodes of the same ID (MAC address) to different logical nodes, i.e. a cloned device attempts to enter the network, will fail because the backoffice prevents such security breaches.

Once connected and authenticated, the DAA node has to be directed to the appropriate service-providing devices (HSD Core, Video Core, OOB Core, etc.). When a DAA node obtains its IP address from the DHCP server, the DHCP server also provides the IP address of its Principal Core. The challenge for the operator is to determine how to direct that DAA node to the correct Principal Core, among the several IP reachable Cores that service that node’s population. In other words, in a DAA world, several principle IP cores are IP-reachable, whereas before DAA only one CMTS was physically reachable, because of the

¹ Cores use l2tpv3 protocol to establish dynamic tunnels. It is a two-way control protocol; the RPD picks the tunnel ID. Engines multicast data using static tunnel IDs. They do not communicate via a control protocol with the RPD. Static tunnels are provisioned on the RPD via the GCPP.

wiring; what used to be 1:M, CMTS to nodes, is now N:M, depending on how many switches are connected together..

2. Configuration

Each Core contributes to the node's frequency plan, and for each RF channel in the frequency plan, the node has to be configured to join video multicasts and transmit the multiplex, or to stand up the channel as a DOCSIS service. Legacy video out-of-band (SCTE 55-1 and SCTE-55-2) upstream and downstream channels also have to be configured on the right frequencies. The channels that are present are determined by the geographical area the node serves (its logical configuration) -- but as the DAA node joins the network, it is impossible to ascertain its desired serving area.

When a new node joins the network, the only identifier that it provides is its MAC address. Again, pre-planning which physical node (identified by its MAC address) will be installed at a given location/connected to a fiber is improbable, because line technicians and construction crews typically carry several nodes on the truck at a time and should not be tasked with ensuring that node *A* is placed in location *Z*. For those reasons, a mechanism was required to map a logical node, with its intended service configuration, to its physical instantiation installed in the plant.

3. Service Routing

Because the service and control traffic is delivered via IP on pseudowires, the unicast and multicast routes must also be set up and managed on the network, between the Cores and the DAA nodes. The broadcast video content is multicast to multiple nodes, as they share video channel service groups -- but network capacity is not infinite. Therefore, operators need to ensure that video multicasts are present only on links that will use them. Pseudowires containing video-on-demand content may also be multicast to multiple nodes for efficient use of Edge QAM resources. Orchestrating this routing, especially as new nodes join the network, is complex.

4. Node Management

Once verified, the node must get the latest firmware version. Unlike RDK-based cable modems and gateways, where all possible drivers are distributed, or legacy set-top boxes with on-board agents that check for updates, the Primary Core must ensure that each node is running the latest software, and if not, instruct the node which software file to download and from where to retrieve it. The Principal Core needs to provide software details that are specific to the make and model of the DAA node, as firmware is vendor- and version-specific. Also, in a CI/CD (continuous integration/continuous development) software deployment, not all nodes will get the same software at the same time.

DAA nodes also need additional controls placed upon secure shell(SSH) access – because as shipped from manufacturers, these nodes have static admin usernames and passwords. Because the DAA node is an IP device in an unsecured location, stronger user authentication was needed and must be applied before the node goes into service. Updating thousands of DAA nodes on a service bench, before installation, is impractical, as nodes are shipped to warehouses across an operator's footprint. Ultimately, a solution was needed to allow the node to be updated in the field as part of the installation process.

Core Configuration

In our architecture, it is the responsibility of the Principle Core to configure the DAA node with the settings necessary to bring it up and to begin delivering video services. We developed a Principle Core that delivers no services, but instead supports such management via the generic control protocol, or GCP. The GCP Principle Core (GCPP) is responsible for providing the DAA node's initial configuration; the only interaction it has with the DAA node is via GCP. Once initial configuration is completed by the GCPP, the DAA node is "handed off" to an Auxiliary Core, which configures HSD services. That configuration process, and the tasks performed in each stage are described in the following subsections.

1. GCPP Configuration

The GCPP is responsible for configuring the following on the node:

- The DAA Node's Auxiliary Cores
- Non-service-specific operational configuration (precision timing server settings, RF ports, event management)
- Video services (described in section 1.1)

As described previously, the DAA node gets the IP address of the GCPP from the DHCP server when it receives its IP address. When the DAA node sends a configuration request to the GCPP, the GCPP does not know what configuration to apply to the DAA node, as it only has the node's MAC address and no other indication of where the node has been installed, or what service groups it serves. Essentially, the GCPP is not pre-provisioned with the configuration for a DAA node. Instead, the GCPP sends a request to a new back office element that serves as a repository for DAA node configuration files – the Configuration Manager. This request includes the MAC address of the RPD; the Configuration Manager, working with other back office systems, determines the appropriate configuration to apply to the DAA node. The process of associating a DAA node's MAC address with the appropriate configuration detail is discussed in more detail in the Video Configuration section.

While the video configuration details are determined by node's geographic serving area, the other items that are configured by the GCPP are the same for all DAA nodes it manages. Figure 2 depicts the GCPP's event flow, and Table 1 summarizes the configuration settings that are the same for all DAA nodes that the GCPP manages.

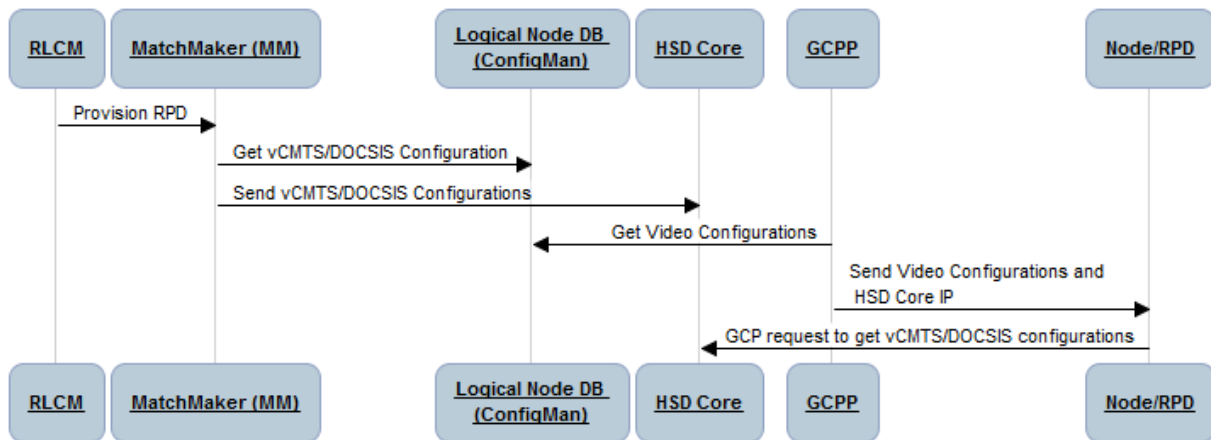


Figure 2 - GCPP Event Flow

Table 1 – GCPP Common Configuration

Settings	Purpose
Core details, including Aux Cores	Configures the GCPP Core details and the IP addresses of any Auxiliary Core to which the RPD will connect once the GCPP finishes configuration (e.g., CMTS Core, Video Out-of-Band Core, etc.).
Core Reconnect	Controls how long the GCP control plane can be idle before the DAA node considers the connection to the GCPP down; also determines the action the DAA node is to take upon connection failure.
PTP settings	Configures the DAA node's connection to the precision timing server.
Downstream and upstream ports	Configures the active RF ports on the RPD and associated power levels.
Carrier Wave (CW) tones	Configures the placement of CW tones in the downstream RF spectrum that can be used for RF leakage detection and automatic gain control.
Event reporting	Configures how DOCSIS-defined events are handled by the DAA node – logged locally or sent to the GCPP Core

The GCPP receives almost all of these configuration details from the Configuration Manager via a configuration file created for the specific DAA node. However, some of the configuration settings are local to the GCPP and Aux Core installation and are not known by the Configuration Manager or the back office – for example, the local IP addresses of the Precision Time Protocol (PTP) server and the Auxiliary Core. For these configuration parameters, the configuration file passed to the GCPP contains variables. The GCPP reads a local file, written in YAML (Yet Another Markup Language and/or YAML Ain't Markup Language), and replaces the variables with the values that are specific to the local installation. As part of the DAA installation process, the YAML file is updated with the appropriate values by the local technicians with access to that information.

1.1. Video Configuration

1.1.1. Logical Nodes

In the past, cable operators were only really concerned with the logical node since physical nodes were not addressable. The logical node ID is associated with the service addresses and is governed by the franchise agreement. It is this logical node ID that is referenced in billing systems, GIS mapping systems, and OSS tools.

The logical node ID associates all the configuration data needed to provision and manage services for a DAA architecture. The logical node ID is best conceived coincident with the initiation of the fiber design. The designer will use these logical node IDs when designing the node and fiber build out from the plant. From a design perspective, the logical node ID is used to represent the location of the physical RPD.

The logical node ID follows some rules and criteria from a node provisioning and management perspective, including:

1. The logical node ID can be an arbitrary value but must be unique across the entire footprint
2. The logical node ID is what will be used to associate with location/addresses
3. The logical node ID is what will be used to associate with subscriber accounts
4. The creation and assignment of the logical node ID must be completed prior to build out and activation of any RPDs
5. The association of the logical node ID must be one-to-one to a physical RPD MAC
6. The logical node ID will be used to create the Video and DOCSIS configurations needed to provision and manage services

1.2. Linear Service Group

The Linear Service Group (LSG) is a logical construct for grouping nodes that share the exact same binary multiplexes to their customers, i.e. they are in the same PEG zone, blackout zone, EAS zone, and ad zone. They all share a common channel lineup which is carried on the same channel map/frequency plan. A node can be in only one LSG. To create an LSG, a location and franchise-based community unit ID (CUID) is used as foundation, to help with identifying the specific video configuration to associate to a node ID and also helps with associating the right zones to the specific nodes.

The LSG will need to have configuration for the following:

1. Broadcast services - a broadcast services configuration, including channel lineups pulled from the video controller; the video controller is a possible resource for helping to identify the source IDs that can be used to create a LSG.
2. Public, Educational, and Government Access Channels (PEG) services.
3. OM (Out-of-Band Modulator) and VARPD – Downstream and upstream traffic configuration.
4. DSG Configuration

To help with collecting the configuration data for the broadcast and PEG services, we initially looked at the video controllers (DAC and DNCS) to help map the corresponding source ID and multicast feeds that feed the channel lineup for a downstream plant. The lesson learned was that for local access, the source IDs were duplicated and would cause issues when trying to build unique configurations for a specific lineup. For testing and trial purposes, the initial video configuration was incorporated into a spreadsheet, and a python script was created to build configuration files necessary for the GCPP and Node provisioning. The mapping of a logical node to an LSG was done manually and the LSG's configuration

information copied into the spreadsheet for every node in the LSG. By doing it manually first, we were able to quickly adapt as we learned which parameters were essential and which were superfluous.

The management of this configuration has now been placed into a video application internally called “Palermo,” which is a tool to manage LSG configurations, map logical nodes to LSGs, and finally push the video configurations to the configMan. We do not keep LSG information in the Access Network platform. Rather, we let Palermo manage the definitions of an LSG in case the definitions change. This isolates video business policy from the enforcement provided by the access network. To the access network, it looks like each node has its own video configuration.

A Video On Demand (VOD) Service Group is a logical construct of grouping nodes together that share all the same On Demand video feeds. A node can only be in one VOD Service Group, but a VOD service group can contain multiple nodes. The number of nodes and the number of QAMs in a service group depends on historical usage, number of homes passed per node, and the amount of frequency the operator can afford to provide for the service. We chose to use four nodes per service group, with a shared pool of four QAMs.

The question then became how to map a node to a service group. In the past, this was done in the RF combining network. The VOD system had no knowledge of nodes and this still holds true. Palermo only needs to provide the frequencies for the VOD QAMs. Four multicasts from the VOD engine are mapped to four ports on the Ethernet switches connecting the RPDs to the access IP network, i.e. the DAAS. When a node is connected to the DAAS, its VOD service group is implicitly selected. The RLCM is notified of this port connectivity by the SDN controller. The configuration for the VOD portion of the node is then completed and stored in the ConfigMan, including the multicast addresses that the node should map to the frequencies given by Palermo.

2. Aux Core Configuration

2.1. vCMTS Configuration Collection

The original CableLabs architecture assumed DOCSIS as the primary core, with the option to hand control for certain services to other Auxiliary cores. Our primary core is the GCPP, and all DOCSIS cores are auxiliary. The Aux Core is responsible for configuring the DOCSIS data paths for the RPDs. In our vCMTS architecture, each node is mapped to a single vCMTS instance, which is a set of containers managed by Kubernetes for that node only.

For configuration of the Aux Core, a “Remote PHY CRD” (Custom Resource Document) is created in a JSON format. One CRD is created for each vCMTS instance. The CRD is comprised of many parameters, some of which are dependent on logical node while others are standardized or static for all nodes.

A CRD template was created that maintains all the static values. For the components that required specific variables, the template was broken out and specific “configlets” were created. This allowed specific modifications of the CRD using site-specific variables. The following configlets were created to allow site- and node-specific configurations:

1. Downstream (DS) Channel Configlets – Allow modification of the number of channels and site-specific center frequencies.
2. Service Class Name (SCN) Configlets – Allow site-specific tiers of service flows, where the name would be standard but the speeds can vary by site and node.

The DAA node will inherit the configuration of its parent i.e. the analog node that served the same set of customers prior to be converting to digital. The backoffice will retrieve the the specific DS channel configuration, the center frequencies used for those channels, and the SCN profiles and specific values for downstream and upstream speeds from the CMTS that hosted the parent node.

2.2. Node ID to RPD Configuration Association – An App for That!

The challenge of mapping a logical node to its physical host was solved by building a smartphone app specifically for the line techs doing the initial node install. Each RPD comes with a QR code that contains the serial number and MAC address of the physical node. Once the technician arrives at the location where the RPD will be installed, geolocation is derived and used to provide a list of nearby intended logical Node IDs. The technician, via the handheld application, associates the specific RPD MAC to a unique node ID by scanning the QR code on the RPD. After the association is made, a trigger is then sent to the RLCM, which in turn triggers the “match maker” to pick an available vCMTS core, generating most of the configurations for video and vCMTS. The RPD is then activated using the configurations from Configuration Manager. The final configuration of the RPD is then stored in a database.

Node Management

1. Software Version Management

Another significant shift facing the industry’s technicians, is that DAA nodes require software, where analog nodes did not. Cable operators carry vast experience downloading code to embedded devices, such as set-tops, cable modems and gateways. However, there was always a “get out of jail free card” available if the download was not successful. That came in the form of the ability of the customer to restart the device. That’s not an option here: Restarting a node hanging on a strand would necessarily require a bucket truck.

As with CPE devices, the software that ships on node devices is usually out of date by the time it goes into service. Therefore, it’s almost a guarantee that the initial provisioning of a node will require its software to be upgraded. However, the latest version of software might be immature. The last thing anyone wants is to deploy new software to a new node. It would be impossible to ascertain if resultant customer issues were attributable to the new software, or the new plant changes. Therefore, the configuration management software must be capable of determining which code version a node should have, based not only on manufacturer and model, but also the logical node itself. At Comcast, this function is provided by the RLCM.

Ideally the initial push of a new software version would not go to nodes that have had recent or ongoing issues. Perhaps other system software upgrades, i.e. new guide firmware, was recently pushed and for related reasons, the system is at risk of having a negative customer experience. Or, perhaps the node recently suffered a fiber cut or other outage. Therefore the configuration management system should be able to use heuristics to determine the appropriate software version for a node. The DAA node is told the version of software it should have at its initial connection to the GCPP. Given that the value changes on a per-node basis, static configuration files are not a viable solution. The configuration manager must dynamically determine the version of software that a node should use. When a software upgrade is rolled out, the nodes need to be notified of the change in a throttled manner to reduce the risk of having too many changes happening to the network at the same time.

2. Secure Shell (SSH) Lockdown

One of the first challenges we encountered was that nodes ship with a default set of accounts and passwords, for debugging purposes. This struck us as an unreasonable security risk. As a result, we elected to prevent any device from accessing the nodes directly; only certain servers, with a given set of keys, are able to access the nodes. Our internal term for this system is “autobahn.” The nodes need the public keys of the “jump hosts” that are allowed to connect to the nodes. Upon initial provisioning, the default user name and passwords must be revoked or changed, and the current public keys associated w/ the jump host private keys must be installed.

The RLCM detects the presence of default accounts and absence of “autobahn” keys, and takes corrective action to change the access controls.

3. Controlling Multicast Video by SDN Controller

In our DAA architecture, we do not pre-plan to which DAAS or port a node is connected. We discover this information when the node joins the network. It’s at this stage that the VOD service group is determined. Based on the physical-to-logical node mapping provided by the aforementioned app, and the port and switch information garnered by the RLCM, a multicast path can be created. We do not use SSM (Source Specific Multicast) in the node, because we have multiple sources of broadcast video. We also chose not to use multicast control protocols e.g. PIM or MLD, as this would defeat the purpose of using inexpensive white box switches. Instead, we build the multicast paths, and then monitor them. If the SDN controller detects a failure in the path, it reprograms a new path. Also, the multicast redundancy app that we created for the SDN controller will monitor the bit rates of the multiple video sources, and can cause drops of all sources but one. In this way, the system controls the sources, destinations and internal port flooding of multicast switches, based on real time information and port discovery.

4. How Do You Know It Worked

In the current architecture, headend techs can tell if their combining is right by tapping into the RF combining network and hooking up set-tops and cable modems. This verifies not only that the video is working, but that the correct channel map, PEGs, and ad zones are being included.

With DAA, signal generation is done for the first time in the node. Therefore, there is no way for techs to check their work, in terms of the soft configuration of the system. We approached this problem in the following manner: The first DAA node that we turned up wasn’t actually done on a pole. We installed a node in the headend, which required installing its own power supply. Then we connected the RF-to-optical analog transmitters and receivers, which were connected to existing nodes in the field. This let us create a pseudo-mCMTS. It also let us roll back to an existing upstream and downstream port, on an existing non-DAA CMTS, as well as the legacy video RF combining network.

In the next phase of deployment (i.e. the first node on the pole), we built a portable test rack with a set-top box, cable modem, eMTA, etc., that can be carried on a truck. When the fiber to the node is spliced in, temporary power is applied and an RF cable is temporarily run to the truck. We dynamically provision the node, as described above, and ensure the node is working properly. Technicians can check that the video is correct. Then the node is spliced into the RF and power network. At this point, we are still somewhat blind as to what’s in the network. That’s why it is essential to apply real time telemetry, to see if the node is receiving and transmitting data on all of its channels. (That discussion is beyond the scope of this paper.)

Conclusion

The evolution of the industry's workhorse HFC architecture toward a Distributed Access Architecture brings with it several changes to node configuration and maintenance that will impact how headend, line and field technicians work.

In traditional HFC deployments, the headend technician controls the services that are delivered to a fiber node because this is done via RF combining in the headend. CMTSs and Edge QAMs output RF channels at configured channel frequencies. This output is split and combined such that the appropriate services are delivered to the lasers going to the node; once the wiring in the combining network is complete, it is rarely changed. Node combining in DAA is done virtually, with software; nodes are virtually directed to different service "Cores," expressed by frequency plan per service, video muxes, DOCSIS-related flows, linear and VOD, and legacy/out-of-band information.

These advances mean that the industry's technical workforce needs to pivot toward node programming, Core configuration and dynamic mapping, to link logical nodes with intended service configurations and physical instantiations. The learning curve will necessarily involve network access and steering, configuration, service routing, and node management.

Comcast worked extensively to develop automated, manageable, and monitor-able solutions for provisioning and managing DAA nodes deployed in the field. While the foundations for these systems are available via the standardized DAA protocols, additional work was required to simplify the provisioning and management of services on DAA nodes.

Abbreviations

CAS	Controller Access System
CRD	Custom Resource Document
CUID	Community unit ID
CW	Carrier Wave
DAA	Distributed Access Architecture
DAAS	Distributed Access Architecture Switch
DWDM	Dense wave division multiplexing
GCP	Generic Control Protocol
GCPP	Generic Control Plane Principal Core
ISP	Inside plant
L2TPv3	Layer 2 tunneling protocol, version 3.
LSG	Linear Service Group
Mcast	Multicast
Mplx	Multiplex – a collection of streams
OOB	Out of band
OSP	Outside plant
PTP	Precision Time Protocol
PW	Psuedowire – a data tunnel
RLCM	RPD Life Cycle Manager
RPD	Remote Phy Device
SCN	Service Class Name
SSH	Secure Shell

VOD-SG	Video On Demand Service Group
--------	-------------------------------

Operational Considerations and Optimization of OFDM Deployments

A Technical Paper prepared for SCTE•ISBE by

Christopher Topazi

Principal Engineer
Cox Communications
6305-B Peachtree Dunwoody Rd, Atlanta, GA 30328
404-269-1021
chris.topazi@cox.com

Michael Cooper

Senior Strategic Architect
Cox Communications
6305-B Peachtree Dunwoody Rd, Atlanta, GA 30328
404-269-4133
michael.cooper4@cox.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
1. Selection of OFDM parameters.....	4
1.1. OFDM channel placement within the RF spectrum	4
1.2. Power spectral density of the OFDM Channel.....	6
1.3. Windowing.....	6
1.4. Adjacent channel interference and profile bit-loading	7
1.5. Profile Assignment and MER thresholds	9
1.6. Channel metrics for monitoring the network	12
2. Early field trials and learnings	13
3. Advanced office and field trials	15
3.1. Office Trial	15
3.2. Field trials	16
4. Future Considerations	18
4.1. Desired Future Metrics.....	19
4.1.1. MER per subcarrier graph.....	19
4.1.2. Profile distribution to identify “trouble nodes”	19
4.1.3. MER Margin to Profile	20
Conclusion.....	20
Abbreviations	20
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 - Expanded Capacity Benefits of Converting Existing Spectrum to OFDM versus Roll-off Reclamation	5
Figure 2 - Spectrum Analyzer Display of OFDM with 3 dB Higher Power Spectral Density	6
Figure 3 - OFDM Sub-carrier MER Impact of SC-QAM channels directly above and below the OFDM channel.....	8
Figure 4 - OFDM Sub-carrier QAM bit-loading at the channel band edge for Cox profiles	9
Figure 5 - Spectral analysis of DOCSIS 3.1 modem in the presence of a data-only trap	14
Figure 6 - MER per subcarrier, collected hourly for one week, single modem	15
Figure 7 - Profile distribution in office trial using all recommended configurations.....	16
Figure 8 - Modem profile distribution after one week of field trial	17
Figure 9 - Modem profile distribution after three weeks of field trial	18
Figure 10 - An example of an MER per subcarrier graph used for troubleshooting	19

List of Tables

Title	Page Number
Table 1- Default guard band configuration from CableLabs DOCSIS 3.1 PHY Specification	7

Table 2 - Default MER thresholds from CableLabs DOCSIS 3.1 PHY Specification	10
Table 3 - Profile MER Threshold Test Performance Results.....	11
Table 4 - Profile MER Threshold Recommendations	12

Introduction

With the convergence of services to IP and the continued growth of bandwidth demands, the access network is stressed to greater and greater capacity limits, forcing cable operators to find the most efficient network configurations such that their networks are providing the largest capacity possible. Such configurations require the highest number of modems operating at the highest modulation order profiles in order to maximize the capacity provided by the networks RF spectrum. The introduction of DOCSIS 3.1 is one of the latest options for cable operators to optimize the performance and maximize the capacity of their networks. DOCSIS 3.1 provides a host of new levers for improving the bandwidth offered by the network. One such lever is the introduction of much higher order modulations than previously provided in earlier versions of DOCSIS. On the downstream, the DOCSIS 3.1 CMs and CMTSs must now support modulation orders up to 4096-QAM with options to support 8192-QAM and 16384-QAM. This is a significant increase over the limited SC-QAM modulations of 64-QAM and 256-QAM required in DOCSIS 3.0 and earlier, and offers as much as a 50% capacity improvement within the same spectrum.

With Cox's initial DOCSIS 3.1 deployments, we focused on a few select parameters in order to maximize the number of modems running the highest modulations. These parameters included:

- 1) OFDM channel placement within the RF spectrum
- 2) Power spectral density
- 3) Windowing
- 4) Adjacent channel interference and profile bit-loading
- 5) Profile assignment and MER thresholds, and
- 6) Channel metrics for monitoring the network.

This paper will explore the approach that Cox used to select each of these parameters and provide some details on the performance we were able to achieve with these configurations.

1. Selection of OFDM parameters

One of the key elements of the DOCSIS 3.1 PHY specification in general, and OFDM specifically, is the tremendous degree of flexibility afforded operators. In considering our OFDM deployment, it was necessary to select a set of parameters to be modified and define appropriate values for each that would result in optimizing the bandwidth of the channel. This section will detail the parameters that were considered and the decisions made regarding each.

1.1. OFDM channel placement within the RF spectrum

The acceleration of consumer bandwidth demands has driven Cox, as well as other operators, to expend significant resources to ensure that their HFC networks are operating as cleanly as possible. As a result, most networks are easily running 256-QAM across their full spectrum of downstream carriers with significant amounts of headroom. Excessive headroom means capacity is being left on the table. With 256-QAM as the highest modulation option in DOCSIS 3.0, and many networks running at over 40 dB MER, networks are running at sub-optimal configurations with DOCSIS 3.0 and prior versions. However, initial DOCSIS 3.1 deployments cannot convert all the RF spectrum to OFDM, as the population of 3.1 modems is relatively low, while the population of other legacy DOCSIS modems is high. As a result, spectrum will need to be slowly migrated from SC-QAM to the more efficient OFDM as modem populations move toward DOCSIS 3.1. A key question is in what portion of the RF spectrum should we begin that initial conversion from SC-QAM to OFDM channels.

As we explored this question, we recognized that certain regions of the RF spectrum (even the downstream spectrum) are less conducive to communications than others. For example, the 700 MHz frequency region commonly overlaps with the LTE band of cellular operation. As a result, these signals can often ingress into the cable operator's networks interfering with communications across the cable network within this spectrum. In addition, the upper frequencies (within 40 MHz of the upper band edge, e.g., 870 MHz, or 1 GHz) typically experience significant roll-off resulting in lower SNRs for channels nearer the band edge making them unusable for 256-QAM SC-QAM channels. Granted, while there are VHF and UHF interferers, the presence of LTE and roll-off is much more common across most nodes and often impacts a wider band of spectrum.

Because of several features of OFDM including long symbol periods, LDPC error correction, and wide channels in which to interleave symbols, it is a much more robust signaling protocol than legacy DOCSIS SC-QAM signals and as a result it is tempting to choose some of the most hostile areas of the frequency spectrum for initial OFDM deployment to take advantage of this robustness. While such a selection would increase the network's capacity, even more capacity benefits can be reaped by leveraging the high MERs from premium legacy SC-QAM channels where higher modulations could be run. For example, 32 SC-QAM channels running 256-QAM transitioned to OFDM running 4096-QAM yields about 500 Mbps of additional capacity within the same spectrum while enabling 42 MHz of spectrum at the top edge that was lost due to spectral roll-off, which would yield about 200 Mbps of additional capacity and require greater refinement in the bit loading of the profile. Using existing spectrum can also avoid the issue of significant performance issues with different cascade depths, which could be an issue in using the roll-off region.

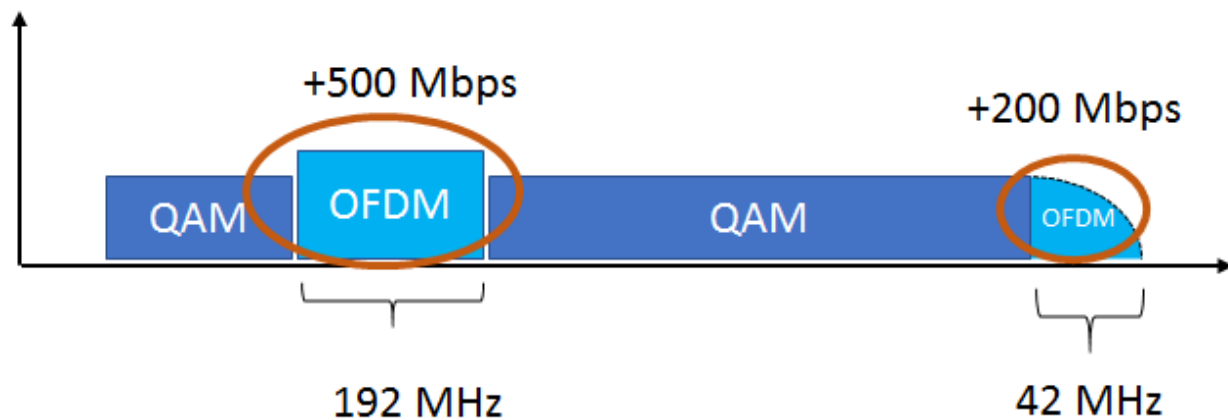


Figure 1 - Expanded Capacity Benefits of Converting Existing Spectrum to OFDM versus Roll-off Reclamation

As a result, Cox selected the lower frequency spectrum of 258 to 450 MHz (lowest frequency spectrum required by DOCSIS 3.1 for OFDM) for our initial OFDM deployments, likely allowing the highest modulations orders (and therefore highest bit rates) supported by the network. By targeting these premium areas of spectrum for OFDM channel deployments, the operator is able to maximize the benefit of DOCSIS 3.1 and OFDM. Spectrum where the existing margin for 256-QAM SC-QAM channels are at their highest should be considered the prime target for initial OFDM deployment and growth. By targeting these spectrum regions, the operator should be able to run the greatest number of modems at the highest modulations supported by the available spectrum. Subsequent expansion of OFDM will likely encompass reclaiming spectrum which is more hostile such as the LTE and roll-off regions.

1.2. Power spectral density of the OFDM Channel

Just as one must consider the placement of OFDM channels within the frequency spectrum, one should also consider the power level to run the OFDM signals as compared with that utilized with the SC-QAMs. The motivation for such a configuration change is similar to choosing the premium spectrum for OFDM, in that the elevated power of an OFDM channel yields even higher SNR allowing the modems to support higher modulations within the OFDM channel. For non-distributed access architectures (DAA), the optical margin within the analog link is a key limitation to consider as total power impacts the SNR that can be achieved on the operating channels. If all channels (OFDM and SC-QAM) are maintained at the same power level, the result is approximately equal SNR across all channels. Providing higher SNR for SC-QAM channels that are limited to running 256-QAM does not benefit the operator relative to network capacity; however, if the OFDM channels are raised in level, the added SNR on the OFDM channel improves the SNR allowing the channel to run a higher modulation.

In the case of Cox's network with our initial DOCSIS 3.1 deployments, we targeted a 3 dB increase in OFDM channel power as compared to a SC-QAM channel, effectively providing an additional single order of modulation benefit. With the evolution to distributed access architectures including RPD and RMACPHY, the specification only allows for 2 dB of variance of channels within the spectrum limiting the elevation that could be achieved with RPD type networks; however, this 2 dB is available. Because we selected a lower frequency for our OFDM channel and when considering spectral tilt, this additional 3 dB of elevated power has a minimal effect on the overall total composite power of the spectrum. As OFDM channels are enabled in the upper end of the spectrum, elevated signals will likely not be feasible as it will push optical links closer to compression; however the migration to DAA will remove the analog optical link limitation.




Figure 2 - Spectrum Analyzer Display of OFDM with 3 dB Higher Power Spectral Density

1.3. Windowing

Another consideration in efficiently using the bandwidth of the OFDM channel is selecting windowing parameters for the channel to be able to guarantee solid performance in a variety of plant conditions. An additional advantage of electing to use a lower frequency for the OFDM channel is that, generally, there

are fewer and less severe impedance mismatches that cause reflections in the plant. It was our intent to minimize the overhead in the channel incurred by cyclic prefix and roll-off, which are explicitly designed to compensate for these plant reflections. We theorized that we should be able to use the lowest available settings for cyclic prefix and roll-off and still have good performance in the plant.

Of course, verification of this theory was needed. This was performed in a variety of steps, beginning with lab testing to verify that the CPE and CCAP would support these settings and that there was no variation in performance across settings in a lab environment. Once this was established, a field test in multiple markets on varying amplifier cascade lengths was performed. We collected MER per subcarrier data and uncorrectable FEC, along with other metrics, across various cyclic prefix and roll-off settings and saw, in our particular case, no degradation when using the lowest possible values.

One downside of reduced roll-off was the potential for additional interference with adjacent SC-QAM channels. The CableLabs DOCSIS 3.1 PHY specification (Table 75 in Appendix V and replicated below for convenience) proposes addressing this by increasing the guard band between the OFDM and the adjacent channel. In fact, some CCAP vendors chose to make this the default behavior and automatically adjust the guard band to higher values when smaller roll-off values are used. While this is certainly a safe approach, we suspected that we would be leaving bandwidth on the table, and we worked with our CCAP vendors to implement an override function to be able to configure the minimum 1 MHz guard band on the channel edge.

Table 1- Default guard band configuration from CableLabs DOCSIS 3.1 PHY Specification

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 ¹
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.		

This required additional testing to determine the effect of the OFDM channel on the MER of the adjacent SC-QAM channels when reducing this guard band, especially given the previous decision to elevate the RF power level of the OFDM channel. In our testing, we determined that minimum roll-off, minimum guard band, with elevated power, resulted in adjacent SC-QAM channel MER readings in excess of 38 dB, and with non-elevated power, MER readings were in excess of 41 dB. Again, placement in the spectrum helped with this issue, as even the elevated power configuration has significant margin for the performance required for 256-QAM, especially in the lower portion of the spectrum.

1.4. Adjacent channel interference and profile bit-loading

Having determined that the effect from the OFDM channel on the adjacent SC-QAM channels was within acceptable limits, it was now important to determine the effect of those adjacent SC-QAM channels on

the edges of the OFDM channel and compensate for it. Fortunately, with the DOCSIS 3.1 specification, the ability exists to define modulation order on an individual subcarrier basis (bit loading) to address a situation such as this.

In our testing, we measured the impact of adjacent channel SC-QAM channels on the individual sub-carrier MER measurements of the OFDM channel. As shown in Figure 3, MER degrades near the edge of the OFDM channel. For the shortest roll-off configuration, the DOCSIS 3.1 PHY specification recommends increasing the guard band to 3.575 MHz at the edge of the channel (7.15 MHz total channel impact when considering both edges) effectively eliminating an entire 6 MHz QAM channel. A more efficient use of this spectrum would be to bit-load the individual sub carriers near the edge with lower QAM levels, thus providing capacity with the spectrum while maintaining equal MER margin across the entire band.




Figure 3 - OFDM Sub-carrier MER Impact of SC-QAM channels directly above and below the OFDM channel

Each modulation level requires the channel to meet a minimum MER threshold in order to meet a certain BER target. Based upon the sub-carrier MER and by applying appropriate QAM-level MER thresholds, Cox was able to define a bit-loaded profile for each of our desired modulation profiles (1024-QAM, 2048-QAM, 4096-QAM). That is, while a 4096-QAM profile is predominantly 4096-QAM for most sub-carriers, the sub-carriers at the edges of the channel would run lower QAM levels in order to maintain

equal MER margin as the subcarriers in the center of the OFDM channel. Equal margin means that there is not a particular area of the spectrum which is more susceptible to channel impairments. Figure 4 illustrates the refined profiles that Cox developed to maximize the capacity of the OFDM channel.




Figure 4 - OFDM Sub-carrier QAM bit-loading at the channel band edge for Cox profiles

The reader should note that if narrower OFDM channel widths are used (192 MHz is shown in the figure), the number of sub-carriers that are bit loaded near the channel edge will remain the same as the adjacent SC-QAM channel impacts the same number of sub-carriers regardless of OFDM channel width.

1.5. Profile Assignment and MER thresholds

The use of multiple profiles within an OFDM channel allows one to maximize the capacity of the channel by assigning cable modems to profiles based upon their received signal quality within the plant. That is, cable modems with better downstream channel quality can utilize higher modulations without errors

while cable modems with lower signal quality may utilize lower modulations to provide error free operation. Categorizing cable modems to appropriate profiles requires the establishment of appropriate MER thresholds.

The CableLabs DOCSIS 3.1 PHY specification (Table 46 in Section 7.5.12.1 and replicated below for convenience) establishes requirements for modulation performance for what is essentially error free post-FEC operation (10^{-6} PER (packet error rate) with 1500 byte Ethernet packets which is less than 10^{-10} BER). In addition, some CCAP vendors have chosen to utilize these performance requirements as their default thresholds to operate each of the QAM levels; however, a 10^{-10} BER requirement is extremely conservative, and while it would be a safe configuration, we believed that a quality customer experience could be provided with a lower BER requirement while at the same time, maximizing the capacity of the network. That is, such an extremely conservative requirement leaves bandwidth on the table.

Table 2 - Default MER thresholds from CableLabs DOCSIS 3.1 PHY Specification

Constellation	CNR ^{1,2} (dB) Up to 1 GHz	CNR ^{1,2} (dB) 1 GHz to 1.2 GHz	Min P _{avg} 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
128	24.0	24.0	-15
64	21.0	21.0	-15
16	15.0	15.0	-15

Table Notes:

Note 1 CNR is defined here as total signal power in occupied bandwidth divided by total noise in occupied bandwidth.

Note 2 Channel CNR is adjusted to the required level by measuring the source inband noise including phase noise component along with transmitter noise and distortion and adding the required delta noise from an external AWGN generator to achieve the desired CNR at the CM F-connector.

Note 3 Applicable to an OFDM channel with 192 MHz of occupied bandwidth.

As a result, Cox executed a series of lab tests to characterize acceptable profile PER performance against MER levels using an AWGN generator. These tests allowed us to better understand the benefits of LDPC error correction in an OFDM channel and better balance those benefits against utilizing the maximum QAM level supported while providing a high-quality customer experience. One thing that we were able to demonstrate during our testing is that the LDPC error correction algorithm is so powerful that we can see 100% correctable errors and still run error-free with some margin left before dropping the modulation. Table 3 provides a summary of our test results.

Table 3 - Profile MER Threshold Test Performance Results

AWGN Attenuator Setting (dB)	Average MER (measured via CM across all subcarriers) (dB)	2nd % MER (measured via CM across all subcarriers) (dB)	Highest Modulation Meeting Customer Experience Performance Requirements
35	39.0	37.3	4k QAM
34	38.3	36.6	4k QAM
33	37.5	35.8	4k QAM
32	36.7	35.1	4k QAM
31	35.9	34.2	4k QAM
30	34.9	33.2	2k QAM
29	34.1	32.4	2k QAM
28	33.3	31.5	2k QAM
27	32.2	30.5	1k QAM
26	31.3	29.4	1k QAM
25	30.3	28.4	1k QAM
24	29.4	27.4	256 QAM
23	28.4	26.4	256 QAM
22	27.4	25.4	256 QAM
21	26.4	24.7	256 QAM
20	25.4	23.7	256 QAM
19	24.4	22.7	256 QAM
18	23.4	21.7	None

When comparing our results with those thresholds defined in the CableLabs DOCSIS 3.1 PHY specification, our PER at the various QAM levels was performing acceptably when reporting an average MER that was nearly 6 dB below the CableLabs levels. Granted, one never wants to operate the network near the performance edges, so we chose to add approximately 1.5 dB of margin to the acceptable levels to establish the Cox thresholds. As a result, Cox's thresholds were 4 dB below the CableLabs levels as shown in Table 4 below.

Table 4 - Profile MER Threshold Recommendations

QAM Level	CableLabs Spec CNR (dB)	Cox Profile Threshold
4096-QAM	41	37
2048-QAM	37	33
1024-QAM	34	30

As another consideration, the ability of DOCSIS 3.1 OFDM channels to dynamically support multiple profiles where a cable modem may seamlessly move between profiles over time as plant environment changes is critical to the success of a multi-profile configuration and eliminates the need to run only the lowest profile supported by the entire modem population. Cox spent a significant amount of effort in early system testing to understand the behavior of this dynamic process and to assess its impacts on the customer experience before even considering deploying a multi-profile configuration. This testing included both downgrading the profile as well as the reverting back to higher profiles. By using a 1.5 dB margin on our thresholds, we assured that profile changes were made prior to a customer experiencing negative consequences. Because of the complexity of this dynamic process, early testing was quite beneficial for Cox as we were able to identify several software issues and work with our vendors to incorporate necessary changes. We would encourage other operators to do the same against the full spectrum of various DOCSIS 3.1 cable modems that they are expected to deploy within their network.

1.6. Channel metrics for monitoring the network

The final parameter that we considered for our initial DOCSIS 3.1 OFDM deployments was which metrics to add to support monitoring of our OFDM channels. OFDM channels introduced a number of significant changes from our traditional 6 MHz wide 256-QAM Reed-Solomon FEC-based channels. OFDM channels are wider ranging from 24 MHz to 192 MHz, with more common deployments expected to cover 96 MHz to 192 MHz. Historically, a 32 SC-QAM channels covering 192 MHz worth of spectrum would be providing a number of performance metrics (e.g., MER, Receive Power Level, FEC statistics, equalization, etc) for each channel. If we assume just 5 parameters per SC-QAM (which is certainly on the low side), we would be characterizing a 192 MHz portion of spectrum with 160 metrics (5*32). With an OFDM channel, we are now challenged to characterize and monitor that same 192 MHz portion of spectrum with what is likely a much smaller number of metrics.

In addition, with SC-QAM channels utilizing the less powerful Reed-Solomon FEC, which degrades more slowly when transitioning from corrected codewords to a condition of uncorrectable errors, the industry learned to use Reed-Solomon FEC corrected codewords as an early indicator that the network was beginning to operate near the edge with perhaps 2-3 dB of additional degradation margin before customer impacting conditions might occur. The movement to LDPC error correction in OFDM channels makes this much more complicated as LDPC offers significantly more margin (perhaps 5.5 dB) between the point at which corrected errors first manifest and where uncorrectable codewords are present. In order to maximize the capacity of our network by utilizing the maximum modulation levels possible, we will likely operate well unto the LDPC corrected codeword space. In addition, the margin between a small number of uncorrectable codewords (e.g., 0.1%) and a large number of uncorrectables (>2%) is very small and leaves little warning. While it is certainly important to collect these LDPC codeword statistics, Cox's sense was that their value maybe somewhat diminished with OFDM channels except to confirm customer-impacting conditions where uncorrectables are present.

Rather than 32 SC-QAMs, a 192 MHz OFDM channel has up to 8000 sub-carriers which is certainly too many MER values to monitor. The DOCSIS 3.1 specification combines those 8000 sub-carriers to produce just three MER measurements: Average MER, 2nd % MER, and MER standard deviation. The average MER and MER standard deviation are exactly what one would expect representing the average and standard deviation of each of the up to 8000 sub-carrier MER values. The 2nd % MER represents the MER level of the sub-carrier separating the highest 98% MERs from the lowest 2% MERs. While it may seem the 2% number would represent a good conservative metric to use to assess channel quality, it seems that this metric can be easily degraded while not being a good indicator of true channel performance. Specifically, Cox encountered this problem as a result of our bit-loaded profiles utilizing minimum 1 MHz guard band configuration. The subcarrier MER values near the band edges are significantly impacted by the adjacent SC-QAM channels resulting in significantly lower 2% MER values. However, because we have bit-loaded the profile with lower QAM levels for these sub-carriers, the low value doesn't actually indicate a poorly performing OFDM channel. Similarly, narrow band interferers can significantly impact the 2% MER and their effects are almost always mitigated as a result of the powerful LDPC. Even LTE would degrade 2% MER but is often compensated for by LDPC. As a result, we found that the better metric was the average MER for the channel.

Another useful metric for assessing channel quality for a particular modem is the profile it is using. The CMTS instructs the modem to make MER measurements and return the results using the OFDM Profile Test (OPT) mechanism defined in the DOCSIS 3.1 MULPI specification. From this information, the modem is assigned a set of profiles that it may use. This provides an excellent indicator of channel quality for a particular cable modem. Similarly, for assessing the overall quality of the channel, measuring the modem counts or percent of the total model population within each modem profile is a good indicator of the overall health of the channel within or across network segments.

2. Early field trials and learnings

When first deploying DOCSIS 3.1 OFDM channels, we determined that because of relative immaturity of the technology, it would be best for Cox and for our customers to begin with simpler configurations and work our way toward more advanced features. Having tested our recommended configurations in the lab, we selected a group of parameters to be part of our initial deployments and a set of parameters to be added later.

For first deployments, we chose to use placement of the OFDM channel in the premium spectrum and to elevate the power spectral density of the channel by 3 dB relative to the adjacent SC-QAM channels. We also selected the cyclic prefix and roll-off to be the lowest overhead values available on our CCAP. For modulation profiles we elected to use a limited number of very straightforward, simple profiles and used the default guard band settings as specified in the CableLabs standards. Rather than push the limits of the technology in the early stages, we employed only two modulation profiles - one which placed all subcarriers at 256-QAM and the other with all subcarriers at 1024-QAM.

The reasoning behind using the simplified configuration was to allow us to get some real-world experience with OFDM deployments, see our monitoring tools in action, train our workforce, and allow us to collect more engineering data to validate our previous lab testing. We also had a keen interest in determining the stability of MER for modems operating in a live network, thus trying to determine how often we should expect modems to require a downgraded or upgraded profile.

As expected, when this was deployed, virtually all modems were placed on the 1024-QAM profile and remained there. While this was the expected result, we did also discover a drawback to the selected placement of the OFDM channel. The specific spectrum that was selected was being filtered by legacy

data-only traps that still existed in our network. Fortunately, due to the design of our channel bonding groups, the modems were still able to function in a partial service mode using SC-QAM. This allowed us to react quickly to the situation and build tools around identifying the locations of these traps and have them removed. This was done through a combination of identifying potentially affected modems using scripting to poll the CCAP, followed by collecting spectral analysis data from the modem via the Proactive Network Maintenance (PNM) MIB and analyzing for the signatures of known traps. An illustration of the signature of one of these types of traps is shown in Figure 5.

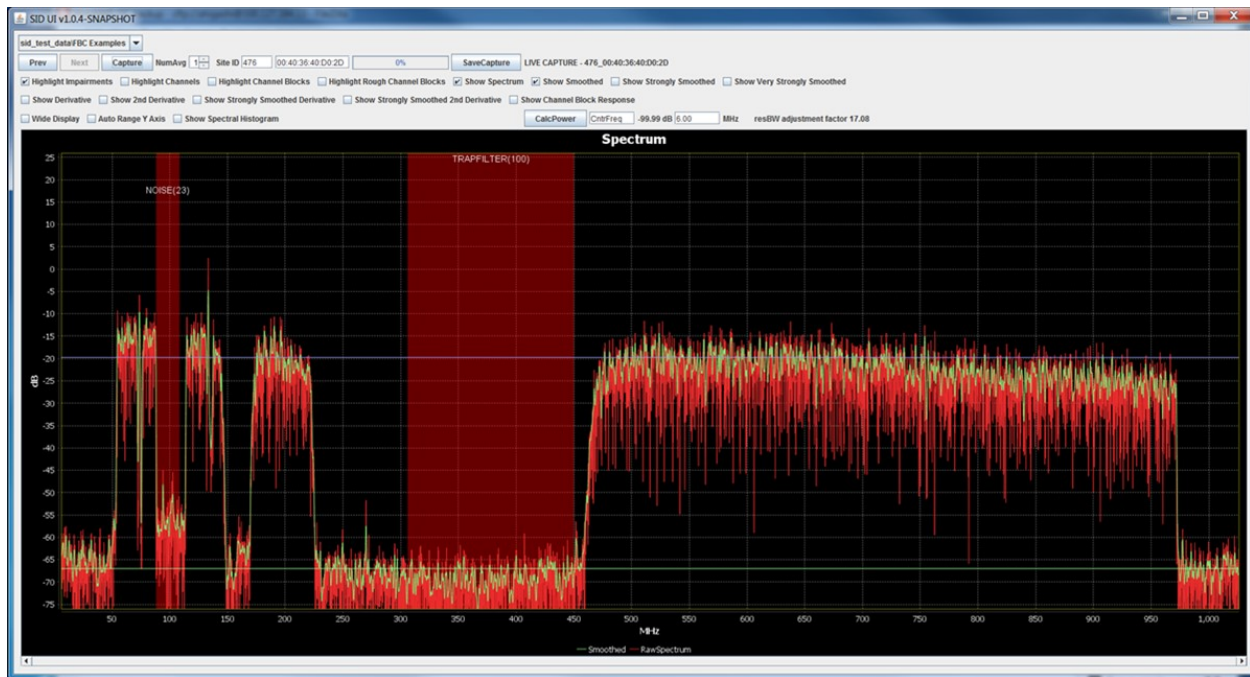


Figure 5 - Spectral analysis of DOCSIS 3.1 modem in the presence of a data-only trap

While the results were as expected for modem profile assignment at initial launch, we also were interested in longer-term stability. It would be considered an undesirable behavior for modems to be frequently changing modulation profile assignments, and we wanted to verify that we had left enough margin in our MER thresholds to allow the modem to remain on the profile to which it was initially assigned during normal operation, given no disruptions in the plant. To measure the stability of the performance over time, we elected to collect MER data per subcarrier for a sample of modems at hourly intervals and compare the results over a week. We intentionally chose areas of the longest amplifier cascades in the market where we were performing field trials in order to get the maximum RF receive level variation at the modem as well.

From a stability perspective, the results were quite good, with the MER varying by an average across subcarriers of roughly 2 dB over the course of 24 hours, and less than an additional 0.5 dB over the course of a week. This validated our assumption that leaving 3 dB of margin between the MER thresholds for initial profile assignment and the MER threshold at which the modem would downgrade profiles would be sufficient for stability purposes. An illustration of the hourly measurements of MER per subcarrier over the course of the week for a single modem is shown in Figure 6.

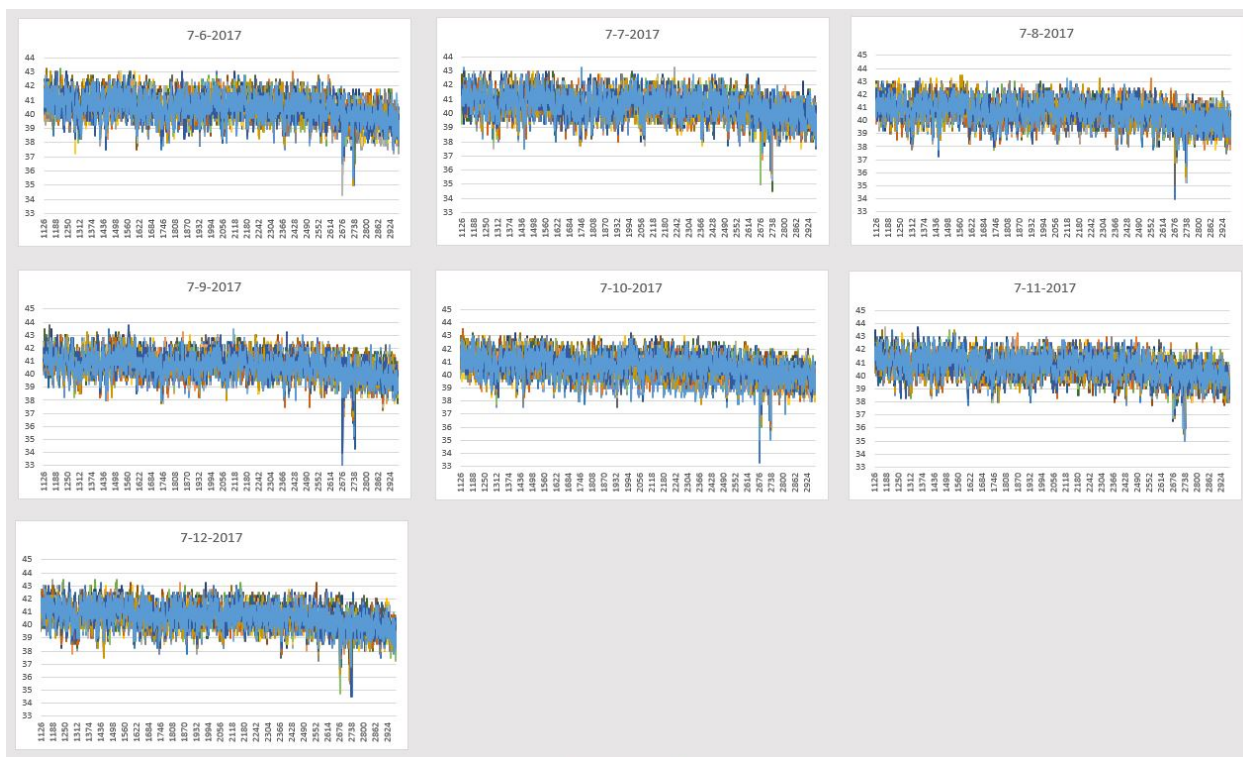


Figure 6 - MER per subcarrier, collected hourly for one week, single modem

3. Advanced office and field trials

Having established that we could achieve the expected performance with our initial, simplified set of OFDM parameters, it was now left to incorporate the remainder of our recommendations. As before, the process was to first perform lab testing of these new parameters to ensure support from both the CCAP and CPE devices. One early finding from this lab testing was that not all devices supported a mix of modulation orders on subcarriers within the channel, which was a critical piece of being able to reduce guard band. This was later resolved using firmware updates to both the CCAP and CPE, but underscored the relative immaturity of the technology and further validated our approach of beginning with a smaller, simpler set of parameters and progressing to the more advanced settings. This is also an example of why the cable operator should always be conservative in the rollout of newer technologies, especially when implementing the more advanced features of those technologies.

3.1. Office Trial

Once issues were resolved and lab testing was successful, the next logical step was to deploy these more advanced settings in a controlled office environment and verify the performance. The location for this test was our corporate office in Atlanta, with employees using modems at their desks as the test devices.

For this testing, guard band was reduced to 1 MHz on each side of the channel, and four profiles were created:

- Profile 0 = all subcarriers at 256-QAM
- Profile 1 = majority of subcarriers at 1024-QAM, edge subcarriers at 256-QAM
- Profile 2 = majority of subcarriers at 2048-QAM, edge subcarriers tapered at 1024-QAM and 256-QAM

- Profile 3= majority of subcarriers at 4096-QAM, edge subcarriers tapered at 2048-QAM and 1024-QAM

While the plant serving the office building is not a completely accurate simulation of actual plant, it would expose the changes to a larger user and device base. As shown in Figure 7, the results from the office trial showed a large majority of modems on profile 3, a smaller number on profile 2, and less than 5% of modems using profiles 1 or 0. Recognizing some of the unique challenges of an office environment, we believed that this would be a baseline result, and we could expect actual plant performance to be similar or slightly better.




Figure 7 - Profile distribution in office trial using all recommended configurations

3.2. Field trials

The encouraging results from the office trial showed great potential for deployment, but in order to be more conservative in our approach and to follow established processes, we elected to trial the advanced configuration on a subset of the network before deploying nationally. The nodes for the field trial were selected using the following criteria:

- Geographic diversity - we wanted all regions to be represented
- High concentrations of DOCSIS 3.1 modems
- Where possible, whole CCAP line cards were selected, for simplicity
- Areas with friendly customers / employees who could test the service
- A variety of amplifier cascade depths

Having selected a set of nodes from each market, configuration changes to incorporate the more advanced parameters were put into place, and we began monitoring the outcome. After the first week, we had approximately 250 DOCSIS 3.1 modems on the more advanced configurations.

The initial measurement made was a profile distribution percentage, similar to the measurement made during the office trial. The results were encouraging across the 250 modems of the trial, with nearly 90% of modems using highest modulation order profile, and over 95% using profiles at 2048-QAM or greater, as shown in Figure 8.




Figure 8 - Modem profile distribution after one week of field trial

The higher than expected percentage of modems on profile 0 was a concern, and we wanted to determine the cause of this. Investigation into those specific modems showed that in nearly every case, there was a data-only trap that was present. Despite efforts to remove these from the network, in this case, these were modems where they were missed.

The next steps were to add more modems and collect data over a period of time to look for changes in profile distribution. Another snapshot was collected two weeks after the first snapshot, and as expected, it showed similar results. This time, however, the number of modems was increased, and the sample represented 1500 modems. The results after three weeks of field trial are as shown in Figure 9.




Figure 9 - Modem profile distribution after three weeks of field trial

These results showed also that the distribution of modems over time was consistent, which further reinforced conclusions from previous testing that the margin in MER thresholds for determination of initial profile assignment were sufficient.

Another key parameter that was an indicator of stability was the number of modems that were currently in a downgraded state, or currently utilizing a modulation profile that was not the same as the initial assignment for that modem. Out of 1500 modems, only 9 modems were in a downgraded state. Less than one percent of modems experiencing an event that caused a profile downgrade also seemed to support the conclusion that the system was stable. We also verified that uncorrectable FEC codewords were not being seen except in the rare case when a modem was required to downgrade.

Of course, we are also tracking the extremely important customer experience metrics such as call volume and truck rolls in the field trial areas, but that data is not currently available at time of this publication. Based on the measurements that we were able to make and the feedback from employees in field trial areas, we expect that it will be favorable, and if so, we will roll out these configurations across the enterprise.

4. Future Considerations

Through this process, the configurations that were selected certainly appear to be achieving the objectives. However, there were also learnings from the process that we have not yet been able to incorporate, which we will discuss here.

4.1. Desired Future Metrics

As part of the testing process, we identified three metrics that we would find extremely useful. Two of these are already supported and simply require tools development, while the third currently requires more development in the DOCSIS 3.1 modem chipsets. These metrics will be discussed below.

4.1.1. *MER per subcarrier graph*

While we have developed proprietary tools using Microsoft Excel macros and/or scripts, it would be extremely useful to have an enterprise-level tool to plot the MER per subcarrier for a given modem or set of modems. One very valuable aspect of the graph is being able to see if there is ingress under an OFDM channel and if it has a signature that can help find the source. An example of where we used this capability in our lab is in tracking down an RF switch with poor isolation based on the plot shown in Figure 10.




Figure 10 - An example of an MER per subcarrier graph used for troubleshooting

This graph showed what looked to be 6 MHz wide SC-QAM underneath the OFDM channel, which allowed us to then start troubleshooting to find the source. This is just one way in which a graph like this could be useful.

4.1.2. *Profile distribution to identify “trouble nodes”*

It is now possible to use modem profile distribution data to identify nodes that are not operating optimally. One aspect of OFDM is that it covers a much larger portion of the spectrum than a single SC-QAM, and it is far less susceptible to narrowband interference. This makes it a useful signal as an indicator of the overall health of the downstream spectrum. While we are not yet at the level of having an enterprise tool for this, it is possible to create thresholds for percentages of modems on a given profile to

identify sub-optimal nodes. When tied into the customer location, modem profile distribution could also be used to localize and isolate issues.

4.1.3. MER Margin to Profile

This metric is defined in the CableLabs DOCSIS 3.1 OSSI specification, but is not yet supported in most CPE. We believe that this will be a useful metric when completing an install or truck roll to determine the quality of the installation or repair. It is preferred that when the technician leaves the house that there is margin to the lowest acceptable profile. While an estimate of this parameter can be made using average subcarrier MER, we feel that this metric is more accurate, and therefore, will be a valuable piece of information.

Conclusion

Having completed these lab, office and field tests, we feel confident in recommending the following for optimizing the use of OFDM:

- Choose clean spectrum for placement of the OFDM channel, if possible
- Where possible, use higher power spectral density for OFDM as compared to SC-QAM
- Use minimum windowing parameters unless plant conditions are extremely severe, and the use of minimum guard band is sufficient to avoid adjacent channel interference
- Use the flexibility of bit-loading to compensate for interference from adjacent SC-QAM channels
- The default MER thresholds for modulation orders can be reduced by several dB, depending on how much margin the operator prefers
- When using these recommendations, average subcarrier MER and assigned modem profile become the most important metrics, until more metrics become available

Use of these recommendations has demonstrated that greater than 90% of modems can operate in real HFC plant at 4096-QAM, resulting in nearly a 50% increase in efficiency as compared to using 256-QAM SC-QAM.

Abbreviations

AWGN	Additive White Gaussian Noise
BER	Bit Error Ratio
bps	bits per second
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Customer Premises Equipment
DAA	Distributed Access Architecture
dB	deciBel
DOCSIS	Data Over Cable Service Interface Specification
FEC	Forward Error Correction
GHz	GigaHertz
HFC	Hybrid Fiber Coax
Hz	Hertz
IP	Internet Protocol

ISBE	International Society of Broadband Experts
LDPC	Low-Density Parity Check
LTE	Long-Term Evolution
Mbps	Megabits per second
MER	Modulation Error Ratio
MHz	MegaHertz
MIB	Management Information Base
OFDM	Orthogonal Frequency Division Multiplexing
OPT	OFDM Profile Test
PER	Packet Error Ratio
PNM	Proactive Network Maintenance
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RMACPHY	Remote MAC PHY
RPD	Remote PHY Device
SC-QAM	Single Carrier - Quadrature Amplitude Modulation
SCTE	Society of Cable Television Engineers
SNR	Signal to Noise Ratio
UHF	Ultra High Frequency
VHF	Very High Frequency

Bibliography & References

CM-SP-PHYv3.1-I14-180509: *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1 Physical Layer Specification*; CableLabs

CM-SP-MULPIv3.1-I15-180509: *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1 MAC and Upper Layer Protocols Interface Specification*; CableLabs

CM-SP-CM-OSSIV3.1-I12-180509: *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1 Cable Modem Operation Support System Interface Specification*; CableLabs

Operational Practices for Energy Conservation/Sustainability Measures in the Cable Outside Plant

An Operational Practice prepared for SCTE•ISBE by the Access Network Efficiency Working Group in the SCTE•ISBE Standards Energy Management Subcommittee

Chair: Daniel Howard
Principal
Enunciant, LLC
2512 Parkdale Place NE
404-625-1593
Daniel.SCTE@gmail.com

Chris Day, Analog Devices

Kevin Gantt, CommScope

John Holobinko, Cisco

Rob Howald, Comcast

Dick Kirsche ConsultKirsche

Todd Loeffelholz, Alpha

Dan Marut, Comcast

Kathleen Miles, PG&E

Rene Spee, Coppervale

Dean Stoneback, SCTE•ISBE

John Ulm, Arris

Lamar West, LEW Consulting

Dan Whitehouse, Hitachi Consulting

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Access Network Energy Consumption, Conservation Measures, and Operational Practices for Improved Efficiency	5
1. Energy consumption growth with current access network trends.....	5
1.1. Next generation nodes and distributed access architectures	5
1.2. Impact of wireless access on the OSP energy consumption: small cells/5G, Wi-Fi, IoT, CBRS, LoRa, security and surveillance, and more.....	8
1.3. Impact of fiber deep and not moving power supplies	9
1.4. Summary of energy consumption trends for the OSP architecture	10
2. Energy Efficiency Improvements for OSP Equipment	10
2.1. New LPS technologies with greater efficiency	10
2.2. New plant actives with greater efficiency	12
2.2.1. GaN technology	12
2.2.2. APSIS™ functionality	13
2.2.3. Example of new MSO node requirements.....	14
2.2.4. Generic access platform (GAP)	14
2.3. Alternate energy for the OSP	14
2.3.1. Cable operator explorations of solar powering the OSP	14
2.3.2. Implications of feeding power into the 90 V quasi-square wave OSP power system.....	17
2.3.3. Lessons to learn from IEEE 1547-2003	17
2.3.4. What this means to the OSP.....	21
2.3.5. Implications for MSO partnerships with utility providers	21
3. Payback framework for access network energy conservation measures.....	22
4. Access network energy conservation measures and estimated paybacks	23
4.1. New plant actives/technology as an ECM.....	23
4.2. Early retirement of LPSs: Replacing Gen 1 & 2 power supplies with latest technologies; collocating with nodes	24
4.3. Running a lower loss line to power deeper fiber nodes	24
4.4. Alternate energy as an ECM for the access network.....	25
4.5. Machine learning/artificial intelligence as an ECM for the access network	26
5. Access network operational practices to improve energy awareness and efficiency	27
5.1. Measurement and verification (M&V) to prove energy savings	27
5.2. Longer term planning for dynamic power consumption (e.g. APSIS) and renewable energy usage in the OSP	28
5.3. Measurements to make during normal maintenance.....	29
Conclusion.....	29
Abbreviations	30
Bibliography & References.....	32

List of Figures

Title	Page Number
Figure 1 – ANE roadmap of technologies and energy conservation measures	4
Figure 2 – Breakout of energy bill components for cable operators	5
Figure 3 – Power consumed by edge facility and outside plant per 100k HP (from [ULM])	6
Figure 4 – OSP annual powering cost (from [HOL])	7
Figure 5 – Remote PHY annual powering cost (from HOL).....	7
Figure 6 – Total access network energy consumption models for conventional and virtualized architectures (from [BEL])	8
Figure 7 – Energy consumption growth with practical fiber deep deployment	9
Figure 8 – Utility savings from more efficient LPS technology at \$0.22 per kWh	11
Figure 9 – Envelope Tracking Concept (from [DAY])	12
Figure 10 – Active linearization concept (from [DAY])	13
Figure 11 – Customer solar powering of node concept (from [LGI])	16
Figure 12 – Feasibility of using customer diurnal excess solar capacity (from [LGI]).....	16
Figure 13 – Actual frequency response to a generating unit trip on Oahu [IEE]	18
Figure 14 – “Duck curve” showing increased impacts of distributed PV on system load for the worst day of the year [IEE]	19
Figure 15 – Proposed frequency behavior requirements of an updated IEEE 1547 [IEE]	20
Figure 16 – Proposed voltage behavior requirements of an updated IEEE 1547 [IEE]	21
Figure 17 – Task breakout showing potential AI prediction component (after [AGR]).....	26
Figure 18 – Detailed energy incentive planning (from [CIF])	27

List of Tables

Title	Page Number
Table 1 – Energy impact of implementing wireless and new revenue generating services	9
Table 2 – Efficiency improvements in power supply technology	10
Table 3 – Simple payback period example calculation	22
Table 4 – Example DC loop resistances of various coax hardlines and power feeder	25

Introduction

The cable outside plant consumes the majority of power of the overall network. Power is consumed not only by the active devices (optical nodes, amplifiers, Wi-Fi hot spots, LoRa gateways, micro cells, cell backhaul, 5G, etc.) but also by the process of moving the power through the outside plant to reach these devices and the associated Joule heating (I^2R) losses in the cabling.

Changes underway in the network architecture and increases in the sophistication and functionality of the active devices are predicted to cause an increase in the power required. Remote physical (PHY) and remote media access control (MAC) and PHY are predicted to increase the power dissipation at the locations of conventional optical nodes. Amplifiers with higher gain and higher radio frequency (RF) bandwidth and the resulting linearity demands may also increase the required dissipation at these locations, however this increase may be offset by gallium nitride technology and new energy efficiency measures. Finally, the addition of new active devices such as Wi-Fi hot spots, micro-cells, and LoRa gateway devices will add new powering burdens to the outside plant (OSP).

The present network architecture transports the electrical power for these active devices through existing coaxial cable or cable that is specifically dedicated to transmission of electrical power. The coaxial cable exhibits a significant electrical resistance at the frequencies used for electrical power, typically 60 Hz or 50 Hz. The resistance of the coaxial cable at powering frequencies is virtually identical to the resistance at DC and is therefore typically referred to as the coaxial direct current (DC) loop resistance. The DC loop resistance can result in dissipation of significant amounts of electrical power in the coaxial cable itself as the electrical power is transported through the OSP to active devices.

These sources of power utilization in the OSP will be enumerated in this document. Several expected changes that will affect the power utilization are then reviewed. Finally, energy conservation measures to reduce or minimize the power required to operate the outside plant are explored. The technology trends and recommendations for conservation measures are part of an overall roadmap (see Figure 1) of the SCTE•ISBE Access Network Efficiency (ANE) working group within the Energy Management Subcommittee (EMS) of the SCTE•ISBE Standards program. This roadmap exists to address energy consumption and conservation in the access or “last mile” portion of modern cable networks.

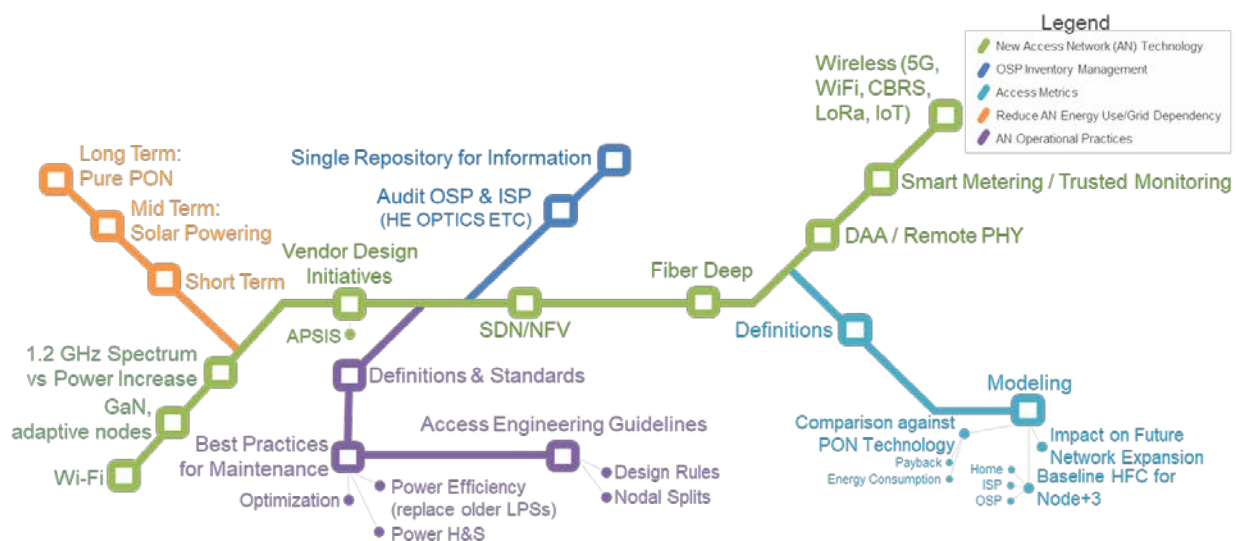


Figure 1 – ANE roadmap of technologies and energy conservation measures

Access Network Energy Consumption, Conservation Measures, and Operational Practices for Improved Efficiency

1. Energy consumption growth with current access network trends

It is well known that the OSP dominates the overall energy bill for a cable operator [ULM]. Figure 2 shows that the access network, as comprised of edge facilities and the OSP represents over 73% of the cable operator's total energy bill, and that the OSP alone can represent over half of the total bill.




Figure 2 – Breakout of energy bill components for cable operators

1.1. Next generation nodes and distributed access architectures

There are many changes underway in access architecture, such as deployment of fiber deep, distributed access architectures (DAAs), including remote PHY and remote MAC-PHY, expansion of the upper RF spectrum limit to 1.2 GHz, full duplex (FDX) DOCSIS™, and even virtualization of the converged cable access platform (CCAP). Unfortunately, *each one* of these are currently predicted to increase the OSP energy consumption. The good news is that DAA, by significantly expanding the capacity of the fiber link from the edge facility to the fiber optic node in the OSP, will significantly improve the energy efficiency of the OSP as measured by the OSP energy efficiency metrics of kilowatt hours per terabyte consumed (kWh/TB). Nonetheless, the total energy bill in the OSP will go up, just not by as much as it would by using conventional hybrid fiber-coax (HFC) architectures to achieve the same capacity increase. A survey of modeling efforts to predict the energy consumption of network architectures is given next.

Ulm [ULM] recently compared several network architectures in terms of energy consumption via the combined edge facility + OSP energy consumption per 100,000 households passed (HP), the results of which are summarized in Figure 3. In this analysis, Ulm neglected the cooling reduction possible in edge facilities from DAAs and pure fiber architectures.




Figure 3 – Power consumed by edge facility and outside plant per 100k HP (from [ULM])

Ulm concluded that for all HFC architectures, the power to drive the outside plant continues to dominate the overall energy consumption of the combined edge facility + OSP, and further that the migration to fiber to the premises (FTTP) will take over a decade. His model predicted a 33% increase in kW/100k HP in adding remote PHY to a fiber deep (FD) architecture. Thus, finding ways to improve the energy efficiency of the OSP to reduce energy consumption will be a critical component to helping cable operators reduce energy costs for the network and the company overall.

Similarly, Loeffelholz [LOE] concluded that the industry migration to FD node plus zero (N+0), shown in Figure 3, makes sense financially in terms of the cost per HP to upgrade the OSP and also in terms of enabling remote PHY/DAA as well as FDX DOCSIS™, and that is indeed the way many cable operators are moving. As will be described below, the desire of many cable operators not to move existing line power supplies (LPSs) so they can avoid permitting costs and delays, is unfortunately resulting in the distributed powering scheme of current OSP to be reverted back into centralized powering schemes, with their concomitant higher Joule heating (I^2R) losses, thereby decreasing energy efficiency.

Holobinko also predicted a 25-30% increase in OSP power consumption and associated energy costs by moving to remote PHY architectures [HOL] (see Figure 4), and he did include the potential reduction in the heating, ventilation, and air-conditioning (HVAC) costs in the edge facilities that are expected from a lower heat load from the telecommunications equipment (Figure 5). Unfortunately, the HVAC energy savings was quite small in comparison to the growth of energy consumption in the OSP, in agreement with Ulm's work. Note that Holobinko used a sample RF design for areas that were selected based on density per mile, node size and aerial/underground ratio, and the sample design consisted of three HFC nodes where the design areas varied in size from ~400 to ~600 HP per fiber node, and the density per mile of coax plant varied from ~95 to 160 HP/mile, while Ulm's work integrated the energy impact across 100k HP. The main point is that a rise in OSP energy consumption of 25-33% is predicted just from moving to remote PHY/DAA by both of these authors.




Figure 4 – OSP annual powering cost (from [HOL])




Figure 5 – Remote PHY annual powering cost (from HOL)

Finally, a Bell Labs Consulting white paper [BEL] came to similar conclusions about remote architectures causing approximately 25% growth in the OSP energy consumption, while dropping the energy consumption of the edge facilities (see Figure 6). In this case, the point made was the virtualization in the

edge facilities and OSP could further and significantly reduce the energy consumption in the edge facilities while only slightly raising the energy consumption in the OSP. However, in light of much smaller energy consumption of edge facilities vs. the OSP (see the pie chart in Figure 2 previously), it is not clear whether the total access network energy consumption (edge facilities plus OSP) will actually decrease from such virtual architectures.

Total energy consumption




Figure 6 – Total access network energy consumption models for conventional and virtualized architectures (from [BEL])

In important point made by Ulm is that while FD architectures theoretically reduce the total OSP energy consumption by reducing the total number of actives and replacing them with modern, more energy efficient technology, if the LPS units are not moved from the original node locations to the new node locations, the OSP energy consumption per HP could be double what it might otherwise have been if LPSs were moved to the new node locations. Unfortunately, due to permitting costs, the LPS locations are seldom moved. This effect is independent of the OSP energy consumption increase due to remote PHY.

1.2. Impact of wireless access on the OSP energy consumption: small cells/5G, Wi-Fi, IoT, CBRS, LoRa, security and surveillance, and more

In a companion SCTE•ISBE Expo 2018 paper by one author [LOE2] it is shown that the impact of deploying a plethora of wireless devices in the OSP is far from insignificant. A summary of the total energy impact of adding full geographic coverage of a variety of wireless technologies to the OSP on line power supplies (LPSs) is shown in Table 1. If all new wireless devices were deployed maximally, the summation of this additional energy load on the OSP is 820 Watts, or approximately 9.2 amps of additional current draw on existing power supplies, which at a nominal level of 15 amps represents a 61% further increase in OSP energy consumption.

Table 1 – Energy impact of implementing wireless and new revenue generating services

Technology	Watts per Unit*	Units per LPS*	Total Power (W)
CBRS	50	8	400
LoRa (32 Ch)	45	1	45
Wi-Fi	45	5	225
5G LTE	75	2	150

There may be hope to reduce this impact of wireless devices on OSP energy consumption by noting that many of these devices do not operate at full power consumption all of the time, and thus have a lower average power than the peak power usage. In an upcoming SCTE Network Operations Subcommittee (NOS) journal paper tentatively titled “Operational practices for the deployment and maintenance of wireless devices on HFC plant,” the Hitachi and Comcast authors note that “understanding real world power consumption by the device and any deviation from the manufacturer’s device specifications can help avoid overestimating power requirements and performing unnecessary power supply upgrades.” They further note that “Services requiring a two-device solution (a separate radio device and HFC-powered cable modem) can place a substantial load on the network, depending on the required power draw of the radio device supplied over Ethernet from the HFC cable modem,” and thus driving to single device solutions could mitigate some of the OSP energy growth shown in Table 1. Finally, they note that “savings can be achieved by understanding the average and peak power draw of the radio device in real world conditions and under normal usage -- rather than conservatively designing to a sustained maximum power over Ethernet (POE) or POE+ output for every device.”

1.3. Impact of fiber deep and not moving power supplies

Fiber deep deployment unfortunately is also raising the OSP energy consumption. An example study done by one of the authors (Spee) is shown in Figure 7, where even after eliminating one LPS post-split for fiber deep, total power goes up by 61%. This is a general result when not moving or adding LPSs to be collocated with fiber optic nodes, resulting in a significant increase in OSP energy consumption. The additional power requirements are due to (a) the additional nodes and (b) the losses in transmitting power from power supply to node, including additional power supply efficiency losses.

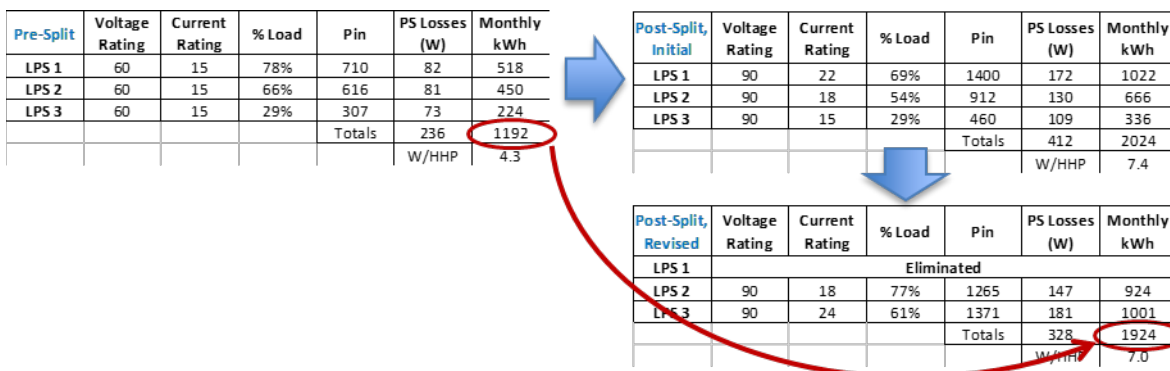


Figure 7 – Energy consumption growth with practical fiber deep deployment

In general, multiple system operators (MSO's) do not want to move power supplies to avoid permitting and cost issues. Likewise the high cost of updating the existing coax to improve direct current (DC) loop performance results in an unattractive return on investment. In some cases, cable replacement also runs into utility pole attachment issues since overlash opportunities are limited.

The results shown in Figure 7 are for going from N+4 to N+0 OSP architecture. The coaxial cable type used was predominantly types 500 and 625. These have loop resistances of 1.7 and 1.1 Ω / kft, respectively, which is typical in many older systems. Many others use type 750 or 875, which drops the loop resistance to 0.76 or 0.56 Ω / kft, and can significantly reduce the additional energy loss from transmitting power to the new, deeper fiber node locations. But optimized coaxial power distribution is still less efficient than a distributed powering architecture where all fiber nodes have their own dedicated and collocated power supplies.

It is possible however to mitigate somewhat the energy inefficiency of not moving LSPs. The initial design increased the projected power requirement by almost 70%, due to both the increased power requirement and also the increased power supply loss. However, the actual power requirement came down a bit in the end: since the LPSs were to be upgraded anyway, this provided an opportunity to increase the power provided by LPS 3 in Figure 7 and eliminate LPS 1. LPS 3 was in a more optimal location for delivering power, so I²R losses and thus overall powering requirements could be reduced. The savings in fixed monthly utility fee and operational expenditures such as batteries, preventive maintenance (PM), and so on, mitigated the increased energy consumption slightly. Again, this project, along with many current projects, had a strict requirement not to change out any cable in order to avoid utility costs and delays. Essentially, fiber deep projects as practically implemented, revert the distributed OSP powering architecture to a central powering scenario.

1.4. Summary of energy consumption trends for the OSP architecture

It has been shown that nextgen remote PHY/DAA technology will likely increase total access network power consumption by 25-30%. Fiber deep deployment without moving LPSs increases OSP energy consumption by another 61%. And adding a full wireless capability across multiple services can add yet another 61% to the OSP energy consumption. Thus, OSP energy consumption has the potential to more than double from the combination of OSP evolution and new wireless service additions.

2. Energy Efficiency Improvements for OSP Equipment

This section will cover energy efficiency improvements in next generation OSP equipment that can offset the OSP energy consumption growth presented in the previous section.

2.1. New LPS technologies with greater efficiency

Table 2 shows how modern 3rd generation LPS units (GEN 3) are more efficient overall in providing the same amount of power to the OSP active devices than previous generations of LPSs.

Table 2 – Efficiency improvements in power supply technology

PS Model	Load (Watts)	Daily Energy Consumption (kWh)	Annual Energy Consumption (kWh)
GEN 1 PS (90V)	600	18.00	6569.37

PS Model	Load (Watts)	Daily Energy Consumption (kWh)	Annual Energy Consumption (kWh)
1985 - 2000	1000	27.94	10196.64
GEN 2 PS (90V)	600	16.90	6168.48
2000 - 2010	1000	26.88	9812.47
	1350	35.82	13074.30
GEN 3 PS (90V)	600	16.87	6157.42
2011 - Today	1000	26.68	9736.79
	1350	35.45	12939.27

Since this will be explored later as a potential way to conserve energy in the OSP, note that the useful life of modern LSP technology, specifically a typical transformer module, is designed for a 15-year design life at an industry standard design temperature profile. The typical inverter module (IM) is designed for 8-year design life at an industry standard design temperature profile. Past history indicates a 15 to 20+ year life cycle for the power supply transformer modules and more than 8-year life cycle for the power supply inverter modules. Finally note that new LPS units also provide technologies like battery balancing, event recording, and accurate DOCSIS™-based monitoring, in addition to higher efficiency. Newer units also maintain their efficiency down to lower loading conditions, perhaps as low as 30% of max load. But while the energy efficiency improvement may be greater at lower loads, the total energy savings are lower due to the smaller energy consumption at low loading conditions.

Figure 8 shows the utility savings from early retirement of older generation LPS units at @ \$0.22 per kWh.




Figure 8 – Utility savings from more efficient LPS technology at \$0.22 per kWh

To determine the utility power from LPS efficiency, I^2R losses and network load to utility power consumption, the following equation is used:

$$Utility\ Power\ (kW) = \left(\frac{P_{Network\ Load} + \sum \left(\frac{P_{@Active}}{V_{@Active}} \right)^2 * L_{Ohmic\ loss\ per\ ft\ of\ cable} * D_{ft\ of\ cable}}{Power\ Supply\ Efficiency} \right)$$

where P and V @Active are the power and voltage respectively at the active device, L is the DC loop resistance per unit length of the OSP cable and D is the length of the OSP cable.

2.2. New plant actives with greater efficiency

2.2.1. GaN technology

While remote PHY technology does consume more power overall, it uses GaN technology which is more energy efficient. The energy savings of newer GaN technology for OSP active devices was highlighted in an ANE working group presentation [DAY] where the combination of GaN technology, active linearization, and digital pre-distortion in next generation OSP active devices could have the potential to save up to 50% of power consumption per device.

A brief description of these technologies is below:

- Digital pre-distortion (DPD)
 - Requires digitized signal (Remote PHY)
 - >15 dB composite triple beat (CTB) reduction in back-off
 - ~25% reduction in power (18W drops to 14W, e.g.)
- Envelope tracking: adjust bias depending on envelope of signal (see Figure 9)
 - Cost & overhead of generating envelope
 - Amplifier bias modulation complexity
 - Voltage and current methods
- Beyond Class A : adjust bias depending on RF signal itself
 - Difficult to control broadband amplitude modulation to amplitude modulation (AM-AM) & amplitude modulation to phase modulation (AM-PM) responses
 - High speed devices are increasingly available to soften challenges
 - Can create actively linearized (AL) amplifiers using radio frequency integrated circuit (RFIC) approach (Figure 10)




Figure 9 – Envelope Tracking Concept (from [DAY])




Figure 10 – Active linearization concept (from [DAY])

In the working group’s discussion, the OSP equipment vendors noted that while there is some marginal improvement due to DPD, there is also a concomitant cost in processing power. Also, along with the improvements in efficiency in the next generation of GaN technology, there is perhaps another 5-10% improvement on the GaN hybrid itself. Current devices have GaAs driving GaN, so if next gen devices are 100% GaN, further improvements may be seen, however note that optimal GaN performance intimately depends on combining the best of a high speed ultra-linear technology e.g., GaAs p-type high electron mobility transistor) with GaN in a cascode topology. All agreed that the adaptive power systems interface specification (APSIS™) use case (see next section) for active bias control (“smart biasing”) would be straightforward to design. The group noted that GaN is a given; without GaN a lot of efficiency improvement would be absent. DPD by itself gives 25% savings. OSP adaptive energy efficiency will mainly come from smart-biasing, but this presumes that DPD is used for all four output ports.

2.2.2. APSIS™ functionality

The other key energy efficiency improvement that may come with new active device technology is APSIS. It’s estimated that 15% of power consumption may be saved in active devices using APSIS compliant technology that attenuates the bias current feeding RF amplifiers during off-peak hours [SAN].

APSIS defines a uniform mechanism to collect energy data and issue power state controls to devices in the network [SAN]:

- CMTS/CCAP
- High density edge QAM devices
- Switches/routers
- Fiber transport platforms
- Remote PHY-edge facilities
- Remote PHY-outside plant (OSP)/MDU
- More general fiber nodes
- RF amplifiers
- OSP power supplies

Diurnal adaptation by one vendor at SCTE Cable-Tec Expo 2016 on a CCAP device demonstrated a 40% power consumption reduction during times of off-peak load, which translates to a % 15 efficiency gain during the OSP daily cycle. Another vendor has estimated a ~% 15 efficiency gain by attenuating bias

current feeding RF amplifiers during off-peak times. And in general, the IT equipment demand response applications have yielded up to 30% cost avoidance [GOV].

2.2.3. Example of new MSO node requirements

The following examples of required and suggested new requirements for next generation node technology from a major US cable operator will drive increased energy efficiency in new OSP active devices [HOW]:

- Enable node ports to have their PAs remotely shut off and on, saving significant power consumption when a port is not used (required)
- Consider power factor-corrected DC power packs to optimize efficiency of the AC network power supplies (suggested, especially for FDX DOCSIS)
- Implement digital pre-distortion (DPD) for the PAs to achieve the required modulation error ratio (MER) at lower DC bias levels (required)
- Support envelope tracking technology to reduce the average bias current in conjunction with the variations in the RF waveform (suggested). This is to modulate the power supply voltage with the RF envelope of the input signal to the power amplifier. Envelope tracking to lower average supply voltage to the power amplifier, thereby lowering the DC power consumption of the power amplifier

It is noted that DAA, software defined networking (SDN), and network functions virtualization (NFV) all provide powerful tools for intelligent power management, and orchestration will be a critical component of managing these energy efficiency measures.

2.2.4. Generic access platform (GAP)

Another opportunity for improving efficiency and visibility in OSP active devices exists in the GAP specification currently in development within the Interface Practices Subcommittee (IPS) of the SCTE•ISBE standards program. A possibility exists for integrating energy monitoring, tracking, and control into the platform specification, or as a later module that adhered to the specification.

2.3. Alternate energy for the OSP

In this section, solar generation in OSP is discussed, along with a review of applicable IEEE standards and implications for powering the OSP via distributed solar energy generation, and finally thoughts on utility provider partnering.

2.3.1. Cable operator explorations of solar powering the OSP

In the early days of the SCTE Energy 2020 program, the group explored solar powering of LPSs via photovoltaic panels on the top of the units. Unfortunately, there simply is not enough area on an LSP unit, especially those mounted on utility poles, to provide enough energy to significantly reduce the grid dependence of the OSP. The concept worked for erbium doped fiber amplifiers (EDFAs) for passive optical networks (PONs), but that is only due to the EDFA being used just to extend the reach of radio frequency over glass (RfOG) networks and thus consuming just under a half a watt of power.

Since then, two key industry events have renewed interest in solar power for the OSP. First, back in 2016, Liberty Global International (LGI) held a Spark Innovation initiative and the winning entry was the idea that solar powered cable customer premises might offer their excess energy during the day to power fiber optic nodes [LGI]. By focusing on using customer-provided (and funded) solar energy merely to reduce the need for batteries in the OSP LPSs, they were able to show a payback in 5 years for the small

investment required to allow the nearby fiber node to accept power from the customer's premises. Second, Comcast recently announced a partnership with a solar energy installer so they might offer solar power generation to their customers.

Solar energy generation appears to be most prevalent in states with higher utility rates, e.g. CA, HI, and even parts of the northeast US, which is exactly where US cable operators would also like to reduce their energy bills.

But there are challenges for this concept, including:

- Adapting plant taps, to accept customer's solar power:
 - One solution is to power nodes directly, but this only works for FD architectures
 - Another is to use customer solar energy generation to reduce battery backup requirements, similarly to the LGI study
- Not repeating the mistakes of IEEE smart grid standards for frequency accuracy

In the LGI study, the customer as a sustainable energy supplier was promoted for energy savings, and the client gets money for spare capacity through a compensation scheme with two options:

- Purchasing conditions (electricity) LGI to the customer (from 22 cents/kWh to 6 cents/kWh)
- The client could enjoy cable services completely free of charge (or at least at a substantial discount)

The goal was to reduce node energy dependency from "traditional" suppliers by 10% and to also find a way to keep nodes online during power outages for at least 3 hours. Their calculations showed the operation should be profitable within 5 years.

Their plan was to install a separate 230V AC powerline between a fiber node and the residential/business to business (B2B) customer premises as depicted in Figure 11. Then, connect this powerline up to the control unit that is already in place.

The feasibility of the approach is based on the fact that household/business energy generation and consumption vary throughout the day (Figure 12). Excess solar capacity at daytime (orange) loads battery capacity (blue), allowing for energy consumption at night and selling excess capacity to the grid (green), or in this case, to the cable operator.

Strengths of the concept include:

- Direct customer involvement: the cable operator relies on client solar generation, striving to reach a financial win-win situation for both parties, plus they are indirectly "subsidizing" and encouraging sustainable energy production
- Versatile: can be expanded to include fuel cell or wind turbine power sources
- No huge up-front investments
- Compensation scheme discourages early client-side termination of arrangement

Weaknesses include:

- Limited control over hardware: the cable operator takes care of the powerline, the client retains control over their own in-house solar infrastructure; conversely this is also a potential strength

- Varying solar panel density: having all our nodes throughout the cable networks powered 100% through solar generation is far from realistic. Luckily, the target 10% seems possible.

Varying up-front costs they identified include the distance from node to client, as well as local governmental/bureaucratic procedures.

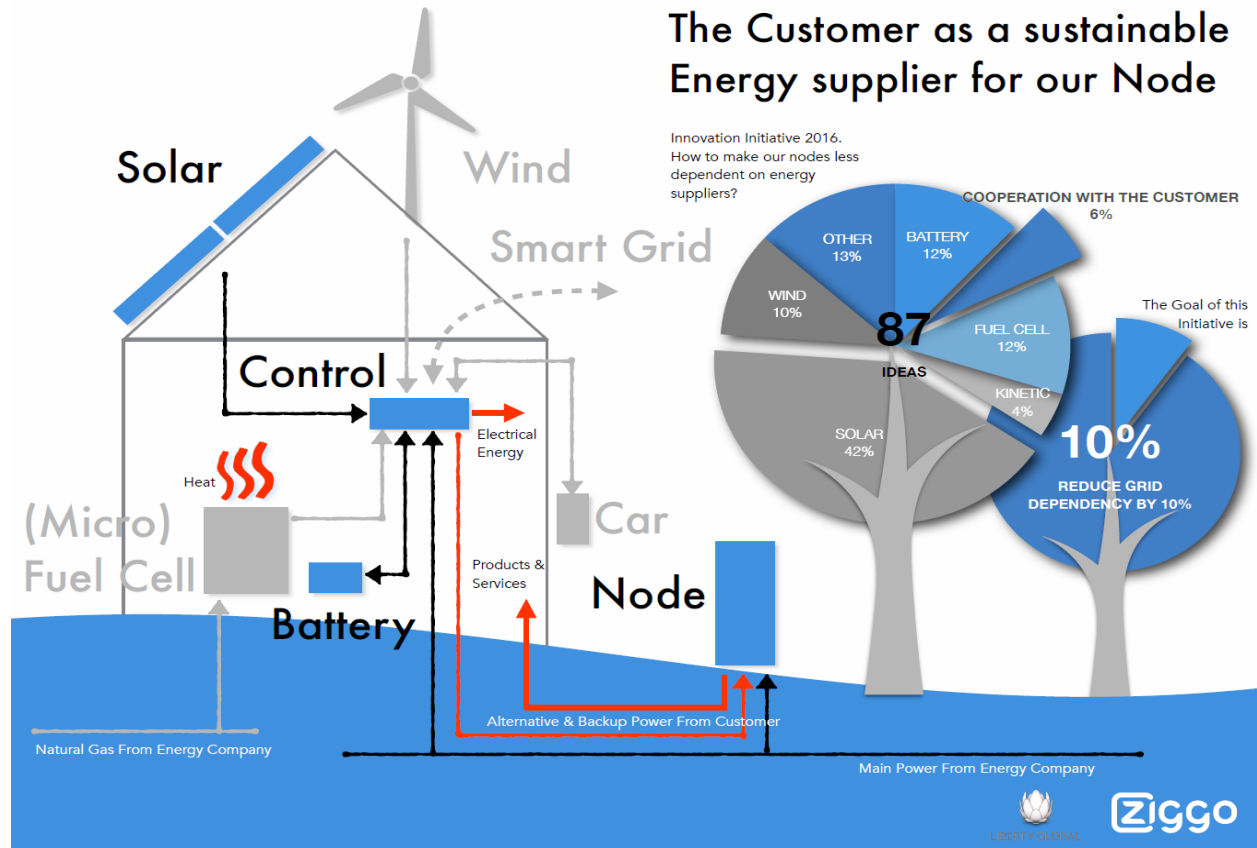


Figure 11 – Customer solar powering of node concept (from [LGI])




Figure 12 – Feasibility of using customer diurnal excess solar capacity (from [LGI])

2.3.2. Implications of feeding power into the 90 V quasi-square wave OSP power system

Whenever multiple power sources fed power into the same AC power network, each of those sources needs to feed power into the network with an appropriate voltage amplitude and waveform and at the proper frequency and phase. This would be relatively easy if the power network had a master controller that knew the instantaneous available power from each source and then dictated the exact instantaneous voltage waveform (along with exact phase) to be fed from each power source. Unfortunately, such a master controller does not exist today, and the implementation of such a controller is likely to be impractical in the short term.

In the absence of a master controller, each power source needs to inject power into the network while maintaining, as much as possible, the network's present voltage waveform and frequency. This is widely done by frequency inverters used to supply power from photovoltaic (PV) panels into the AC mains power grid and it works quite well as long as the energy from such inverters is not a significant percentage of the total power being provided to the network. Once these independent power sources become a significant percentage of the total power, they can interact with each other, yielding undesirable consequences.

IEEE 1547-2003 was created to standardize how independent power inverters should insert power into the AC mains power grid in a safe manner. The standard was created with the assumption that such power inserters would never be a majority contributor of power to the network. Before system operators attempt to insert power from PV sources directly into the 90 V OSP power network, they should learn from the experiences encountered by the electric utilities.

2.3.3. Lessons to learn from IEEE 1547-2003

IEEE 1547 was designed to assure that PV power was never inserted into a power network that was not running as the correct voltage and frequency. The tolerance was quite strict. This was good because it assured that voltage would NEVER be fed into a network that was not being powered by the utility's generation equipment and, thus, PV power would never be fed into a network that was otherwise offline. This is an important safety consideration, which prevents injury or death to people who come in contact with power utility wires that are presumed to be inactive (primarily after storms).

Units compliant with IEEE 1547-2003 were:

- Designed to be sure that power is never backfed to the network if the network is offline
- Required to adhere to very strict voltage and frequency requirements
- Required to trip offline at <59.3 Hz

This standard worked great until PV generation became a significant contributor to the over-all power in the network. An excellent IEEE article describes the effects of distributed energy resources on system reliability and explains the improvements being made to IEEE 1547 to improve system reliability [IEE]. Figure 13 through Figure 16 are from that article and demonstrate the issues to be mitigated.

Figure 13 shows what happened in an actual network when a generating unit went offline. Since insufficient power was being fed into the network, frequency began to lag. Normal protocol for an electric utility in this situation is to shed loads by cutting off portions of the serving area. Unfortunately, much of the power being generated was coming from PV sources, and all those sources went offline at 59.3 Hz as

required by IEEE 1547-2003, exacerbating the problem. When the PV power was needed the most, it was automatically disconnected!




Figure 13 – Actual frequency response to a generating unit trip on Oahu [IEE]

Figure 14 shows that distributed PV (DPV) generation can be a significant percentage of the total power fed into the network during the sunniest part of the day.




Figure 14 – “Duck curve” showing increased impacts of distributed PV on system load for the worst day of the year [IEE]

Figure 15 and Figure 16 show the proposed frequency and voltage behavior requirements of a proposed update to IEEE 1547. These new requirements are designed to assure that the PV power sources continue to feed power into networks that are experiencing minor fluctuations, while also assuring that the PV power sources disconnect when the power network has a critical fault.




Figure 15 – Proposed frequency behavior requirements of an updated IEEE 1547 [IEE]




Figure 16 – Proposed voltage behavior requirements of an updated IEEE 1547 [IEE]

2.3.4. What this means to the OSP

In order to feed PV power directly into the 90 V quasi-square wave power system used in the OSP, inverters will need to be designed and built using a methodology similar to IEEE 1547. The inverters will need to match the existing voltage waveform while feeding a current that is proportional to that voltage waveform in order to assure a high power factor. If the PV power is to be used as utility backup, such inverters will also need to be able to create the quasi-square wave voltage waveform when that inverter is the only power system active. Safety factors must also be considered so that PV generated power does not unexpectedly come on and power a damaged network the morning after a storm.

Solving these issues is beyond the scope of this paper. They are mentioned here so that readers are aware of the issues that need to be considered before PV power can be directly fed into the 90 V OSP power system.

2.3.5. Implications for MSO partnerships with utility providers

One of the ANE members is a public utility, and offers the following considerations for cable operators considering using solar energy to power the outside plant and more generally:

- California the only state that has mandated builders by 2020 to require solar installation on new homes and single-story businesses which is in line with Zero Net Energy goals

- Currently solar power provides 16% of the energy used in California
- Solar power is an alternative driver towards energy production
- There are two choices: builders to provide individual homes with solar panels, or build a shared solar-power system serving a group of homes
- By 2030, California state law will require 50% of the state electricity to be from a non-carbon source.

The acceleration of renewable energy generation in the US will likely cause cable operators to:

- Work collectively with builders, solar and utility companies
- Approach solar companies that can utilize smart inverters
 - Note that there are not a lot out there currently
- Partner with battery storage companies
- Partner with utilities that offer IoT programs

3. Payback framework for access network energy conservation measures

Most access network improvements are driven by performance, capacity, and competitive threats. But could they also make sense as energy conservation measures (ECMs) and payoff in a reasonable timeframe? To answer this question for long term investments usually requires net present value (NPV) type analysis, which takes into account the time value of money and all other dynamic growth and depreciation rates to determine if the investment is acceptable.

But for the engineers and technicians who are considering ECMs for their facilities and networks, a simple payback analysis is often sufficient for initial assessment. Simple payback analysis can compare two options, e.g., continuing to use the existing OSP technology compared to adding new OSP devices, LPSs, or technology upgrades. The rule of thumb for ECM payback is that it is usually acceptable if under 3 years.

Simple payback includes initial costs, rebates, installation costs, cost of power, energy costs, maintenance costs, component replacement costs and end of life costs, as well as how these vary over time, and is calculated as follows:

$$\text{Payback Period} = A + B/C$$

where:

A = The last period with a negative cumulative cash flow.

B = The absolute value of cumulative cash flow at the end of the period A

C = Total cash flow during the period after A.

This will calculate when you ‘break even’ or get your money back. An example is shown in Table 3.

Table 3 – Simple payback period example calculation

	Extra Install	Extra O&M	Savings @ 2% growth	Net Cash Flow	Cum Cash Flow
Yr 1	\$1,000	\$40	\$350	-\$690	-\$690

Yr 2		\$41	\$357	\$316	-\$374
Yr 3		\$42	\$364	\$323	-\$51
Yr 4		\$42	\$371	\$329	\$278
Yr 5		\$43	\$379	\$336	\$613
Cum	\$1,000	\$208	\$1,821	\$613	
Simple payback 3+51/329				3.2	years

There are more detailed and well-known financial analysis that can be considered for large projects: NPV, internal rate of return (IRR) and return on investment (ROI). See the SCTE paper on guidelines for facility cooling technology [DOL] for an explanation of these methods of determining more precisely whether an investment makes sense.

It turns out that for such a short, targeted payback period (3 years), and in the absence of all costs and benefits, an even simpler calculation can be made to quickly determine whether an ECM should be pursued. For the purpose of this operational practice document, the term “coarse payback” will be used to refer to the following simple calculation:

$$\text{Coarse payback} = \frac{\text{extra cost of ECM (beyond nominal costs, cost of BAU or alternative)}}{\text{estimated annual savings of the ECM}}$$

where BAU refers to “business as usual.” Again, if this calculation gives a result that is in the desired range, i.e., roughly 3 years or less, or even slightly more than 3 years, it warrants considering the ECM, collecting all pertinent financial data and doing a more detailed financial analysis to determine its exact payback period or NPV. One way to rule out non-viable ECMs is to make the calculation in the best-case scenario for energy savings, i.e., at the highest utility rate and under the worst energy efficiency conditions in the OSP, such as oldest LPS technology, highest DC loop resistance legacy cabling, and so on. If this calculation gives a coarse payback that exceeds 4 years, it is extremely unlikely that the ECM will ever payoff reasonably for the cable operator. This is the approach used in the remainder of this document for evaluating potential ECMs for application to the access network and identifying those that have the potential to payoff reasonably under some conditions.

Finally, note that another way to compare technologies, and thus ECMs is to compare the cost per home passed, especially when coupling to a planned deployment of new technology. Like NPV analysis, this is more complex and requires far more data, hence the coarse payback will be used below.

4. Access network energy conservation measures and estimated paybacks

In this section, the coarse paybacks of a suite of ECMs will be calculated. Note that all costs used in this section are merely coarse estimates and used only to show how to calculate and compare payback periods of different ECMs. Each cable operator should determine their actual costs from their vendors and perform the calculations for their situations.

4.1. New plant actives/technology as an ECM

It has been estimated that modern OSP active devices that employ GaN technology are at least 25% more energy efficient, and if active linearization and digital pre-distortion technologies are also used in next generation equipment, up to 50% improvement in energy efficiency is possible [DAY]. This begs the question of whether replacing older plant actives with newer devices could save enough energy to pay for

the upgrade over time. However, the upper frequency and RF output power may be extended by the GaN technology and such increases may offset the energy efficiencies, especially if more power is dissipated by the coax itself. Nextgen actives may also include APSIS functionality for further energy savings.

For this ECM, as often happens, cases range from totally infeasible to worth considering, depending on the actual parameters. Here are two examples, with most cases being somewhere in between these:

Case 1: If cost of replacing a plant active (material + installation) is \$20,000, and new node saves 150 W or 1314 kWh/yr, at \$0.16/kWh, total annual savings is \$210 for a coarse payback of 95 yrs.

Case 2: If the cable operator is planning to replace the unit anyway due to maintenance issues/failure or capacity expansion, and newer, more efficient active devices with APSIS functionality only add \$2,000 to the material cost (vs. replacing with a non-APSYS unit), and for high energy cost per kWh of \$0.25 (e.g., Japan or HI), and finally the energy savings are 40% (400 W, 3500 kWh/yr = \$875/yr), then coarse payback would be 2.3 years.

4.2. Early retirement of LPSs: Replacing Gen 1 & 2 power supplies with latest technologies; collocating with nodes

There are energy advantages that result from replacing larger centralized line power supplies with smaller line power supplies that are co-located with each optical node. Such a change reduces the power losses that result from moving power long distances through coaxial cable. The change may also permit the retirement of large, older line power supplies with smaller, newer, more energy efficient line power supplies. However, there is significant initial capital expenditures associated with such a change. There may also be logistical problems such as obtaining access to the commercial power grid. There is also a challenge with properly sizing new line power supplies to support the power requirements of future technologies such as remote PHY, remote MAC PHY, and FDX DOCSIS as was previously described. Finally, if energy incentives are sought for early retirement, as is often done for HVAC systems, the utility provider would require at least one year of useful life left to qualify for an incentive.

Could energy savings from greater efficiency justify replacing older power supplies that are over 15 yrs old? At \$0.22/kWh, roughly \$13-52/yr savings (2nd gen), and \$88-109/yr (1st gen.) units were given in Figure 8. It was seen that the maximum savings are achieved at the highest load even though older generation LPSs are least efficient at lower loads. But even with maximum savings (\$109/yr), the coarse payback period would exceed 20 years for just the equipment costs alone.

Can payback for this ECM be improved? Large utility incentives (at least \$600) with other planned costs accounted for such as addition of an external transponder (\$300) and assuming a near term repair / refurbishment (\$450, every 6-7 years) help considerably. At the same high utility rate (\$0.22/kWh) the coarse payback for the extra cost to retire a 1st generation LPS early and just prior to upcoming maintenance and upgrade costs would only be 3.2 years. For less than ideal circumstances (e.g. lower utility rates, 2nd generation LPS units and lighter loading, early retirement could only be justified if new service offerings or network analytics need to be implemented.

4.3. Running a lower loss line to power deeper fiber nodes

It was seen that by not moving LPSs when deploying fiber deep, the I²R losses increase and take the plant from a distributed powering scheme back to a centralized scheme, which is not as energy efficient. The I²R losses result from the DC loop resistance associated with the coaxial cable. Assume that this DC loop resistance has a value of R ohms. There will be power dissipated in the coaxial cable that is given by

$$Power_{Coax} = I_{LPS}^2 R$$

This power dissipated in the coax is wasted as heat. Note that the power is proportional to the square of the current from the line power supply. Any peaks in the current waveform will result in significant wasted power. Consequently, the peak inrush current that occurs whenever the voltage changes polarity is a significant contributor to the power that is wasted in heating the coax.

One way to reduce I^2R losses is by running a low loss power cable from existing node location to power new node down the line, instead of using the existing coax line surgically to save energy. Table 4 gives example DC loop resistances of various coax types and a relatively new cabling option “PowerFeeder” designed just for power transmission with minimal I^2R losses.

Table 4 – Example DC loop resistances of various coax hardlines and power feeder

Cable Type	DC Loop Resistance
0.75” CommScope P3	0.76 Ω / kft
0.75” CommScope Quantum Reach	1.0 Ω / kft
0.75” CommScope MC2	1.0 Ω / kft
0.5” legacy hardline	1.7 Ω / kft
CommScope PowerFeeder 625 JCAT 3R	0.3 Ω / kft

Consider a best-case ECM for replacing the existing coax with a low loss cable, for example, replacing a 0.5” coax with a PowerFeeder type of 1000 ft. at max current load (18 A). The legacy coax I^2R loss of 550 W drops to 97 W. This would produce \$555 annual savings @ \$0.14/kWh.

Payback: Install, materials, plus permitting/make-ready costs could easily exceed \$2500 for 1000 ft. run, so 4.5 year payback in this case

Therefore, ECM applies only to:

- Legacy hardline with very high DC loop resistance
- High utility cost states/countries
- Minimal permitting/make-ready costs
- Integration with existing deployments (to avoid labor costs associated with the ECM)

Another possible and related ECM would be to minimize the power lost in the coax by controlling the high peak inrush current into the switching regulated power supply in the active devices. This could be done by using power factor correction in each switching power supply in the network. However, this requires modifications at each active device in the network. The feasibility and impact of this ECM is still being investigated by the ANE working group.

4.4. Alternate energy as an ECM for the access network

The LGI study shown earlier gave payback of 5 years when reducing grid dependency by 10%. Would this work in the USA? In some states, the answer is yes, seen in the following:

The cost of industrial electricity in Netherlands (where the LGI study was done) is about \$0.09/kWh (similar to USA average) so payback period is similar. But in high cost states (same ones with lots of solar already deployed), payback would be far faster:

- \$0.14 / kWh (CA), payback is 3.2 yrs
- \$0.25 / kWh (HI, Japan), payback is 1.8 yrs

As a further indication of the viability of this ECM, note that the telcos G.Fast model for next gen digital subscriber loop (DSL) also involves reverse powering. And as it will for cable operators, this DSL reverse powering has to have an orchestration layer that is smart enough to know when power is available, how it's shared, how it depends on the number of homes, etc. Thus, there is some software and management along with the contractual issues to resolve prior to this concept being deployed at scale.

4.5. Machine learning/artificial intelligence as an ECM for the access network

Could we use machine learning (ML) / artificial intelligence (AI) as an ECM? Note first that new technology generally drives the cost of something down [AGR], e.g. semiconductors drove down the cost of arithmetic, the internet drove down the cost of distribution, communication, and search, and so on. Driving down the price of one element often increases demand and expands applicability. AI drives down the cost of prediction, and also decreases the value of substitutes (human prediction), but raises value of complements (data, judgement, etc.).

To evaluate the impact (and ultimately payback) of developing an AI as an OSP ECM, first divide a workflow into tasks as shown in Figure 17, list all the possible places where an AI could be developed, and do an ROI analysis of which tasks would payback well with AI/ML for that task.




Figure 17 – Task breakout showing potential AI prediction component (after [AGR])

As can be seen from companion papers and presentations in SCTE Cable-Tec Expo 2018, AI is already being applied by MSOs for traffic modeling/ analysis, processing proactive network maintenance (PNM) data, optimal orthogonal frequency division multiplexing (OFDM) profile selection, understanding and predicting network ‘behavior,’ and much more. Here are some predictions that might be performed by an AI for the OSP as potential ECMs:

- Benefits of upgrading a node vs. various options vs. not upgrading the node at all (since new technology will increase the energy consumption in the OSP)
- Energy costs system by system
- Costs of ECMs over time
- Energy savings from traffic adaption or, more generally, APSIS functionality
- Network design as a function of competition, energy costs, traffic loads, etc.

After listing the potential prediction and associated tasks per Figure 17 that could be performed by an AI for the OSP, the next step is to do the ROI/payback analysis to see if it pays off in the desired timeframe.

It is often said now that data is the new oil, and this is because with AI/ML lowering the cost of prediction, the value of input data goes up. The ROI for developing an AI should thus include the cost of acquiring data (especially any new sensors to be deployed), analytics, and the required orchestration to implement the AI.

5. Access network operational practices to improve energy awareness and efficiency

5.1. Measurement and verification (M&V) to prove energy savings

To realize cost savings from the ECMs, the cable operator must prove the energy reduction to utilities. Typical incentives / measurement and verification (M&V) models require metering to prove energy savings. The cable operator must go through a detailed process to realize the savings and/or garner incentives from the utilities, as shown in Figure 18.

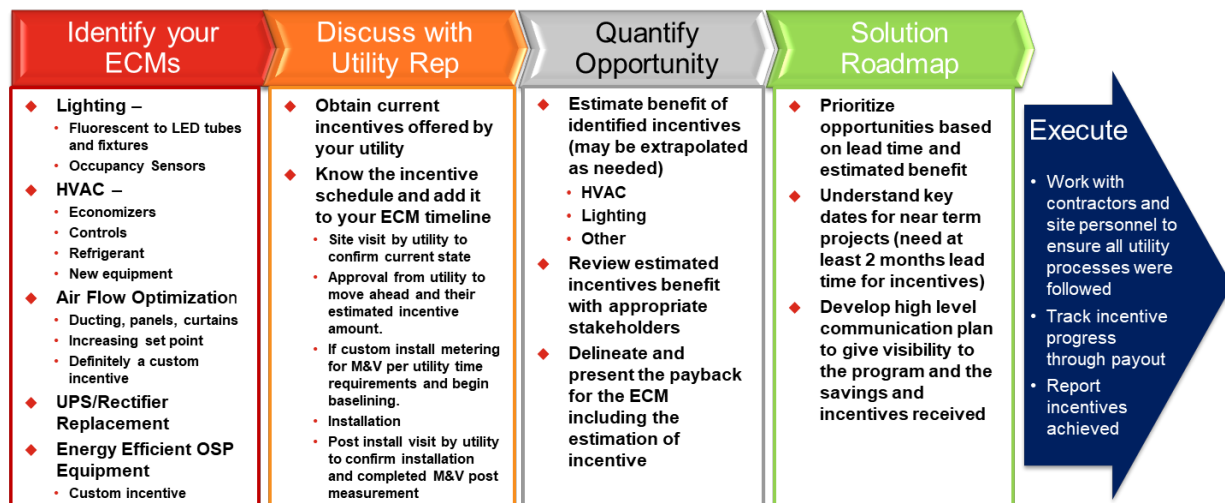


Figure 18 – Detailed energy incentive planning (from [CIF])

The challenge for the OSP in particular is that one must convince each utility on a case-by-case basis. However, if one measures all LPS's in power, there is generally a Gaussian or normal distribution. Then one can approximate the entire population via sampling. Previous efforts of ANE members resulted in agreement with utilities to sample the OSP, measure a subset of LPSs, assume a constant load, and then use utility-approved meters as clamp-on meters to provide data acceptable by the utility companies. Utility personnel usually ride along, measuring a substantial number of LPSs, and then picking a smaller subset of LPSs to come back and check following the energy impacting action taken by the cable

operator. While this process is very laborious and done case-by-case, it does work and can be used to reduce OSP energy costs.

5.2. Longer term planning for dynamic power consumption (e.g. APSIS) and renewable energy usage in the OSP

Unfortunately, the sampling method just described would likely not work for dynamic energy consumption such as would occur with APSIS, energy-proportional wireless devices, and solar powering of the OSP as described previously. Real time monitoring would be required using either built-in capabilities of modern LPSs, add on monitoring modules from the original equipment manufacturer (OEM) or a third party, or via add on IoT sensors that use for example LoRa communications. MachineQ and Leverage are two IoT companies that cable operators have partnered with to provide such IoT OSP-based solutions for both new revenue/services, as well as potentially monitoring the access network of the cable operator.

The challenge of metering line power supply actual power consumption to take advantage of efficiency gained from switching to more efficient line electronics (such as GaN) and actual variations between power supplies runs afoul of a power provider's one size fits all philosophy. It may be possible to qualify some units with a calibrated power metering device that meets power provider standards.

Based on the discussion in the ANE working group and in this document in particular, the time may be approaching for standby power supplies to become more capable and flexible to improve their power consumption footprint. Status monitoring of those devices makes a great deal of sense as a first step. That monitoring can provide a number of benefits for the network and will form the foundation supporting other important functions such as:

- Status monitoring provides the cable operator with a warning if the AC mains power fails. Depending on the charge status of the supplies, batteries, the operator may need provide backup power to that location to support an extended outage.
- With appropriate sensors, status monitoring can inform the operator when the batteries require maintenance. This could reduce the number of regular field inspections saving manpower and vehicle fuel.
- At some point it may be possible to obtain, cost effective, accurate AC line power monitoring modules, with accuracy acceptable to power providers, and incorporate them in the advanced line power supplies. This can provide a means to report the energy consumption for active power supplies which is an idea that is not too far from realization if developers are incentivized. Perhaps the quantities required for outside plant monitoring can drive down cost. For other purposes, power service providers are switching to remotely monitored electric meters. The metering and billing for residential solar systems requires accurate measurement of power flow to and from residences. Many states provide solar renewable energy credits (SRECs) based on residential solar energy generation. Accurately addressing that credit requires power flow measuring and reporting. Those technologies should form the basis of accurate and affordable standby power supply monitoring and reporting devices.
- A first step for operators to become "better partners" on the power grid might be to be able to reduce their power load at times that are critical for power providers. Power companies already offer incentives for actions like disconnecting electric water heaters during critical load periods. The supply status monitoring system can provide the communication control link and assist in the recording of power supply status changes. The cable operator and electric utility will have to negotiate the financial terms of some of these actions. Standby supplies that have disconnected from the power grid will consume some of the standby capacity which will have to be made up by

later charging. Disconnecting the supply in response to a power grid emergency is one possible scenario. Disconnecting during regular daily peak load periods would be another. One interesting opportunity could be to disconnect the supply during periods when “time of day” pricing of electricity is very high. This would require accurate documentation or electric service provider metering of the supply’s time of day consumption. The operator would control disconnect times based on the power supply’s charge status and weather conditions that might predict a need for the supply’s stored energy.

5.3. Measurements to make during normal maintenance

Finally, we can still get visibility into OSP energy consumption for ourselves cost-effectively. The latest generation of LPSs include energy monitoring, but even for the older generation units, it is possible to measure the relatively static energy consumption during general maintenance of LPS and/or deployment of fiber deep, node splits, or DAA. The following measurements might be made without significant additional labor/cost:

- Volt-amperes (VA)
- Volt amperes reactive (VAR)
- True power factor (TPF)
- Displacement power factor (DPF)
- Phase lag angle
- Energy kWh
- Energy cost in \$
- Waveform snapshot

In particular when deploying new OSP devices and architectures, it is recommended to use field power analyzers or similar capabilities to do a baseline before, and M&V after an ECM and/or new node architecture or technology is deployed. Again, APSIS and other dynamic energy measures will require dynamic monitoring/data loggers left on site to prove the energy savings.

Conclusion

While challenging and case-specific, there are indeed OSP energy conservation measures with reasonable paybacks in certain situations. Keeping the payback minimal requires careful planning and logistics to minimize cost components like truck rolls, permitting, make-ready, and some material costs. A key to realizing true savings from these measures is having a process and/or new technology to provide energy monitoring that is acceptable to utility companies

For most ECMs, the plant power is stationary following deployment, so existing methods can be refined, combined and scaled via partnership with the utility companies to keep payback periods low. APSIS and any other dynamic energy control technologies will likely require formal metering or a lengthy partnering process with utility companies to demonstrate average (or even minimum) energy reductions to scale for OSP use.

Abbreviations

5G	5th generation long term evolution cell network
AC	alternating current
AI	artificial intelligence
AL	active linearization
AM-AM	amplitude modulation to amplitude modulation distortion
AM-PM	amplitude modulation to phase modulation distortion
ANE	access network efficiency
APSYS	adaptive power systems interface specification
B2B	business to business
BAU	business as usual
BW	bandwidth
CBRS	citizen's band radio service
CCAP	converged cable access platform
CMTS	cable modem termination system
CTB	composite triple beat distortion
DAA	distributed access architecture
DC	direct current
DPD	digital pre-distortion
DPF	displacement power factor
DPV	distributed photovoltaic
DS	downstream
DSL	digital subscriber loop
ECM	energy conservation measure
EDFA	erbium doped fiber amplifier
EMS	energy management subcommittee
FD	fiber deep
FDX	full duplex DOCSIS
FTTP	fiber to the premises
G.Fast	ITU-G recommendation for fast access to subscriber terminals
GaN	gallium nitride
GAP	generic access platform
GEN	generation
H&S	health and safety maintenance services
HFC	hybrid fiber-coax
HP	households passed
HVAC	heating, ventilation and air-conditioning
I ² R	Joule heating losses in conductors, squared-current times resistance
IEEE	Institute of Electrical and Electronics Engineers
IM	inverter module
IoT	internet of things

IRR	internal rate of return
ISBE	International Society of Broadband Experts
ISP	inside plant wiring
IT	information technology
LGI	Liberty Global International
LoRa	long range wireless communications
LPS	line power supplies
LTE	long term evolution cell phone network
M&V	measurement and verification
MAC	media access control layer
MDU	multiple dwelling unit
MER	modulation error ratio
ML	machine learning
MSO	multiple system operator
N+0	fiber node plus zero amplifiers
NFV	network functions virtualization
NOS	network operations subcommittee
NPV	net present value
OEM	original equipment manufacturer
OFDM	orthogonal frequency division multiplexing
OSP	outside plant
PA	power amplifier
PF	power factor
PHY	physical layer of open systems interconnection model
PM	preventive maintenance
PNM	proactive network maintenance
PoE	power over Ethernet
PON	passive optical network
PS	power supply
PV	photovoltaic
QAM	quadrature amplitude modulation
RF	radio frequency
RFIC	radio frequency integrated circuit
RFoG	radio frequency over glass (SCTE FTTP standard)
ROI	return on investment
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined networking
SREC	solar renewable energy credit
TB	terabytes
TPF	true power factor
US	upstream
VA	volt-amperes

VAR	volt-ampere reactive
Wi-Fi	wireless fidelity network

Bibliography & References

- [ULM] J. Ulm, Z. Maricevic, “Giving HFC a Green Thumb: A Case Study on Access Network and Headend Energy & Space Considerations for Today & Future Architectures,” SCTE Cable-Tec Expo 2016, Philadelphia.
- [HOL] J. Holobinko, “Energy Analysis Remote PHY vs. Integrated CCAP,” presentation to SCTE ANE working group, 2017.
- [LOE] T. Loeffelholz, “Fiber Deep Networks and the Lessons Learned from the Field,” SCTE 2017 Fall Technical Forum, Denver CO.
- [LOE2] T. Loeffelholz, “What small cell wireless wants (and needs) from us,” SCTE Cable-Tec Expo 2018 proceedings, available from www.scte.org.
- [BEL] “An economic analysis of distributed access architectures: The next major cable transformation,” a white paper available from www.bell-labs.com, Dec 2017.
- [DAY] C. Day, “Reducing power in the cable access network,” Presentation to the SCTE Access Network Efficiency working group, August 10, 2017.
- [LGI] S. Kholá and M. Bosman “The customer as a sustainable energy supplier for our node,” winning proposal from LGI Spark Innovation Initiative, August 2016 and presented to SCTE•ISBE EMS Plenary/ANE working group breakout, April 2018.
- [SAN] F. Sandoval, “APSYS Value Proposition for Service Providers,” presentation to IEEE Denver, Feb 2, 2017, available from <http://sites.ieee.org/denver-com/files/2017/02/APSYS-Overview.pdf>.
- [HOW] R. Howald “Energy ramifications of DAA,” presented to the SCTE•ISBE Energy Management Subcommittee plenary meeting, April, 2018.
- [GOV] US Government web site, “Grid Modernization and Smart Grid: Demand Response,” available from <https://energy.gov/oe/services/technology-development/smart-grid/demand-response>
- [HAY] C. Hayes, “Broadband power demands energy efficiency,” from Electronic Specifier online, 18 Jul 2015, with material sourced from Alpha Technologies, available from <https://www.electronicspecifier.com/power/broadband-power-demands-energy-efficiency>
- [DOL] J. Dolan, D. Howard, A. Murphy, K. Nickel, and D. Smargon, “Guidelines for Cable Facility Climate Technology Optimization: Cooling Optimization for Edge Facilities,” SCTE Cable-Tec Expo 2017 proceedings, available from www.scte.org

- [IEE] D. Lew, M. Asano, J. Boemer, C. Ching, U. Focken, R. Hydzik, M. Lange, and A. Motley, “The Power of Small: The Effects of Distributed Energy Resources on System Reliability,” IEEE Power and Energy magazine, Volume 15, Number 6, November/December 2017
- [AGR] A. Agrawal, J. Gans, and A. Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Review Press, 2018, p. 75.
- [CIF] D. Howard, G. Gosko, and J. VanHeynigan, “HVAC Best Practices and Guidelines Webinar,” presented at Comcast Critical Infrastructure Forum, February 2018.

Operational Transformation Via Machine Learning

A Technical Paper prepared for SCTE•ISBE by

Shamil Assylbekov, PhD.

Principle Engineer
Charter Communications
6465 Greenwood Plaza Blvd - Suite 900
720.699.4573
Shamil.assylbekov@charter.com

Devin Levy

Director, Video Operations
Charter Communications
6465 Greenwood Plaza Blvd - Suite 900
3037934199
Devin.p.levy@charter.com

Table of Contents

Title	Page Number
Table of Contents	2
Abstract	3
Content	3
Conclusion	6
Abbreviations	7

List of Figures

Title	Page Number
Figure 1 - Decision Tree.....	4
Figure 2 - Actionable intelligence	4
Figure 3 - Process Flow	5

Abstract

In the current paradigm of Operations acting in a reactive, post hoc manner, something has got to give. MSOs cannot continue to throw bodies at problems in efforts to remediate outages, customer impacting events and impairments. Instead, a data-driven Machine Learning (ML) approach needs to be utilized to change the tide for the better. Cable operators have the ability to merge heterogeneous data from various sources in efforts to qualify and classify problem areas. This will ultimately lead to ML-driven operations which will result in a better sustained customer experience and afford the opportunity for the MSO to work under a lean operations model while employing top engineers to do the work of what would have previously taken dozens of engineers to handle in a post-hoc world. Operational problems go past operations expenditure. Customer impacting events is the name of the game, and by utilizing different ML approaches for different issues, one can classify various events that occurred, provide actionable intelligence and network automation to curtail outages, reduce the mean time to resolution (MTTR) of events in the end providing a better customer experience. Initial efforts of our group have yielded a reduced MTTR for outages and impairments, less escaped defects for feature enhancements, actionable intelligence that affords us the opportunity to make data driven decisions rather than gut reactions or best efforts. This result was mainly achieved via the Deep Learning approach for the purposes of anomaly detection, and a mixture of ML approaches for the purposes of escaped software defect categorizations. At Charter, ML has made Operations a more intelligent organization. Some of the decisions are now made with real time actionable intelligence via our own plethora of analytics; we now have ability to improve the customer experience on a daily basis. And, we encourage the adoption of this new operational transformation, for the betterment of our industry.

Content

Machine Learning has come a long way since the 50s, when it first was conceptualized. In 1952, Arthur Samuel wrote the first learning program to play checkers. The more it played, the better it played. Then in 1997 IBM's deep blue beat Garry Kasparov in game one of chess. Now, we are solving complex multidimensional , operational problems with ML, and in some cases assisting in saving lives too. ML is not just for fun and games anymore, rather ML proves to be a golden opportunity for our industry and our to paths to lessen Service Impacting (SI) events.

In Charter, we use ML to predict and mitigate outages for Spectrum Guide. Spectrum Guide is Charter's flagship video product. It utilizes a plethora of microservices, vendor platforms, internal and external clouds, various databases, and countless application program interfaces (API). Ultimately this is used to deliver an exception experience to the customer where updates and upgrades can be pushed to the box on demand with minimum downtime or interference. With that said though, sometimes things go wrong, and when you have hundreds of different components, you have a lot of opportunity for things to go wrong in a lot of different directions. And that's where ML comes in to play.

One example of ML that Charter uses for Spectrum Guide is for unexpected system crashes. Although one cannot account for the unexpected, one can learn from mistakes. When we first encountered this outage, it took approximately four hours to troubleshoot. Alarms pointed towards areas that were innocuous or ambiguous, and when there are thousands of moving parts, finding the right one as expeditiously as possible is paramount.

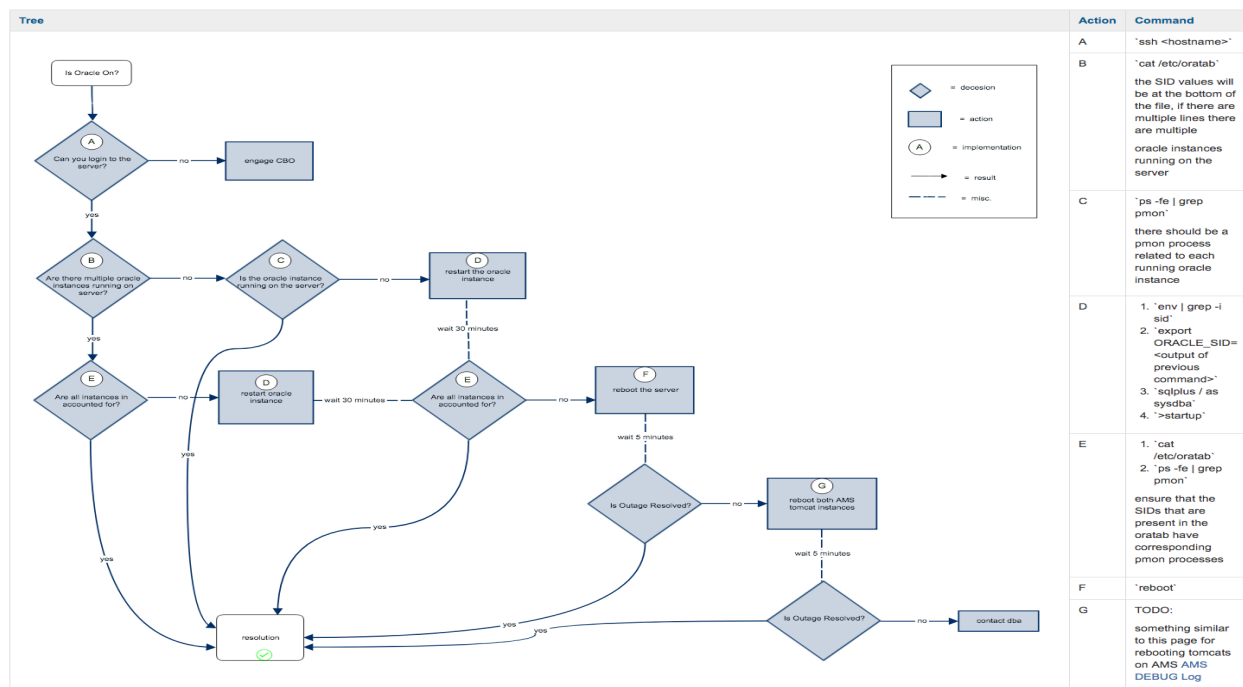


Figure 1 - Decision Tree

When one of the primary the auxiliary databases (DB) experienced a kernel crash at the operating system level, ultimately the database crashed along with it. Any processes that were fed to or pulled from this DB were as a result adversely impacted. After the fault was found, a retrospective took place. We first identified a model for how to identify this type of event, and then put automate in place to follow the model. What we have now is any time this specific issue arises, rather than taking ours to find and resolve we have a reduced MTTR that lasts less than three minutes to resolve instead.

Another example that is more generic but easily adaptable is log analysis. For Spectrum Guide we use time series log monitoring. What we have implemented is anomaly detection. In figure 2 its shown that on August 31st we had a SI event that was immediately followed by automation that reduced the customer impact by 500%.




Figure 2 - Actionable intelligence

By training your data either with human interaction or with a neural network, one can achieve quantifiable improvements in the operational space. But ML is not the end all or holy grail, rather it is another tool in our tool chest. It is not just hammer that looks at all problems in need of being hit on the head. Rather, it is only as good as you make it. For the Video Operations Engineering Spectrum team (VOES) , we have had success in several facets of reducing outages and reducing the MTTR of outages. But, there are still many more SI events that we have not trained our algorithms to address. There's always opportunities to make a tree, create binary classifiers, utilize time series plots and ultimately better the customer experience by incorporating ML.

Machine Learning Initiative for spectrum automated self repair.

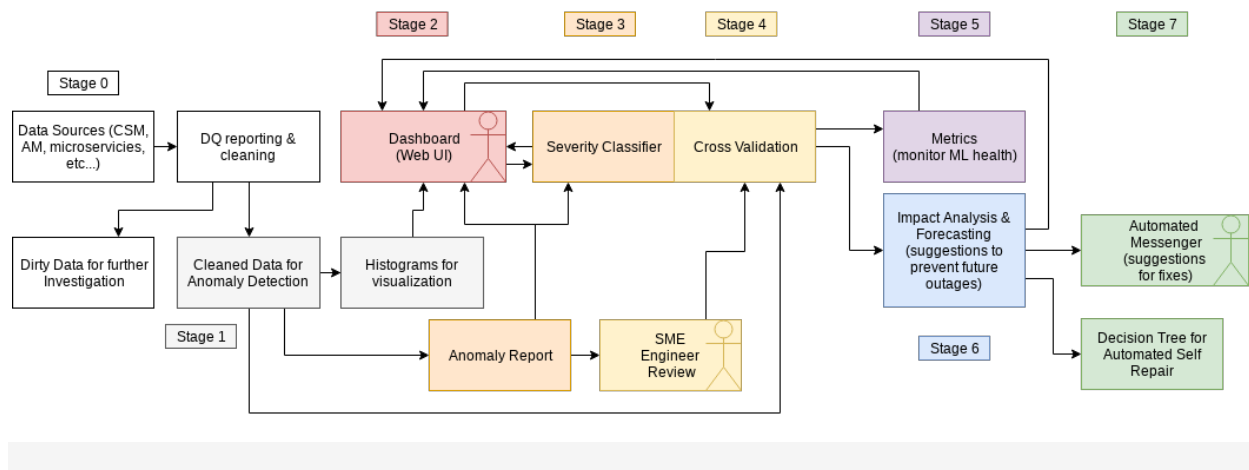


Figure 3 - Process Flow

For more computationally intensive algorithms a distributed computing solution will be needed. Only a subset of one's objectives will be achievable on a small scale and remotely from the actual data files in a batch processing paradigm. One possibility is to set up the clusters or other resources from regional data centers for optimal computational and geographic redundancy. That way the work is done locally to where the data is stored for speed/storage concerns.

In the next phase, one would need to switch investigations from primarily reactive to primarily proactive in order to achieve Anomaly Detection.

This step takes expanded daily reports and generates a meta-report of any anomalous data for investigation. The output of this is a page linked to on the Insights dashboard which lists anomalies with descriptions of the anomaly and relevant images. This will also output a highly curated output with only the anomalous entries for further investigation.

Classification:

As soon as one turns on anomaly detection they will want to classify events based on impact. The meta-report from anomaly detection can then be expanded to include a classification of resulting severity (degree of user impact, outage, this is nothing, engage another team, etc...). Some anomalies are meaningless, while some are indicative of larger problems.

In order to train the severity classification & improve upon it, it would be prudent that the meta-report be reviewed relevant SMEs to see how well the system is assessing the impact of anomalies. The SME

feedback on how meaningful each anomaly is in contributing to solving outages should then be fed back into the severity classifier in order to improve its detection methods.

Note that even after the successful training, machine learning models do not interpret their own results to explain why an anomaly has been flagged. Rather, SMEs will always be necessary not only to train the ML algorithm, but to interpret the needed responses.

Once one has trained on a sufficient dataset of these meta-reports, they can begin operational testing. What is needed at this point is making sure that the classifier is not over-fitting training data by performing k-fold cross validation on our training datasets. The training dataset will be a rolling window of the last X days-worth of data to ensure that the model is based off of, and up to date.

At this point one can consider building/turning on automatic features for detection and different automatic alerts at various levels of severity in an attempt at forecasting problems for repair before they become an outage.

Next, at this point one can begin working on impact analysis to suggest areas for their teams to focus in on in order to reduce outages the most efficient way possible. This analysis would look at the anomaly detection from before and attempt to suggest what anomalous features are consistently contributing to calls.

If the ultimate goal is automated repair and self-healing, then the proposed solution is not sufficient. One approach would be to attempt deep learning for automated repair. It could follow a similar paradigm to what is proposed, but would require SMEs to not only flag the importance of an anomaly but list the affected systems AND the appropriate actions taken (read: calling fully tested, generalized, and automated repair scripts) to resolve the problems in those systems. Unsupervised deep learning algorithms would not be wise for automated self-repair.

Any self-repair system should be implemented upon a well understood decision tree providing a classification of the problem where the conditions to entering the automated repair state are a clear and well understood set of comparisons. IF $x > \text{someValue}$ AND $y < \text{otherValue}$ OR ... Then DoTheThing. The good news is that the classification algorithm described above could be fed into the conditionals of such a decision tree. The classifier shows that the X variable is statistically significantly high and that correlates to outages, so that can be one condition toward calling the automated scripts.

Conclusion

Create a Minimum Viable Product (MVP), it is a good idea and a proof of concept can go a long way.

Picking one data source and implement Insights, Anomaly Detection, and begin building a seed model for Classification. The deep learning feedback loop can be left out at first for the MVP. Object oriented design and proper use of version control will be critical for ensuring long term goals are met while not impairing short medium term goals too much.

Insights and Anomaly detection should prove immediately useful upon completion, and the following stages will need to be trained before their value becomes apparent. If certain details of implementation are handled correctly then the feedback algorithms, meta-report generation, classification, and parts of the anomaly detection should be reusable for other data streams including other sources.

Abbreviations

ML	Machine Learning
MTTR	Mean Time to Resolution
DE	Differential Evolution
NFL	No Free Lunch
SI	Service Impacting
API	Application Program Interface
DB	Database
MVP	Minimum Viable Product

Opportunities And Challenges of Implementing Wireless

Small Cell / WiFi / IoT

A Technical Paper prepared for SCTE•ISBE by

Todd Loeffelholz
VP Product Management
Alpha Technologies Inc
3767 Alpha Way, Bellingham, WA
(360) 392-2172
tloeffelholz@alpha.com

Table of Contents

Title	Page Number
Table of Contents	2
Today's Evolving Network	3
Understanding the Network	3
1. Power Supply Current Draw	4
2. Plant Voltage	5
3. Heat Maps / Geographical Maps	6
3.1. Heat Maps	6
3.1.1. Current Availability	6
3.1.2. Wireless Coverage	7
Planning Ahead	8
4. DOCSIS™ 3.X Backhaul	8
5. 6 Battery Cabinets	9
Powering New Services	10
6. New Services	10
6.1. Small Cell	10
6.2. WiFi	10
6.3. I-IoT	11
6.4. Security and Surveillance	11
Conclusion	12
Abbreviations	12

List of Figures

Title	Page Number
Figure 1 - Typical Current Draw	4
Figure 2 - Power Supply Voltage Profile	5
Figure 3 - Current Available Heat Map	7
Figure 4 - IoT Coverage	7
Figure 5 - D3.0 Transponder	8
Figure 6 - 6 Battery Enclosure	9
Figure 7 - Enclosure with Wireless Equipment	9

Today's Evolving Network

This has been a truly amazing industry to watch over the past twenty years. The pace at which it evolves and recreates itself is amazing. Five years ago, the general thought was that FTTH was the only way to go and if you weren't going FTTH, you were basically going to be extinct. The telcos spearheaded the charge to bring life to their aging assets and "leapfrog" cable. Once they started slowing down, a well-known search engine decided to jump into the fray. Through it all, CATV plants continued to upgrade and push fiber deeper understanding that the key to the whole equation was data. Not necessarily the speed of the data but keeping ahead of the consumer demand for data. It didn't take long for DOCSIS™ to catch up and enter the world of 1Gig by 1Gig but at the end of the day, typical consumer demand is not driving the need for a 1 Gig symmetrical service.

There will always be a need to upgrade the network. Consumer demand will continue to evolve and in general people will continue to drive to faster service. The question is fairly simple – What else can drive revenue and growth in the industry? In the past twelve months, a new growth engine has materialized – wireless. This engine is not new to the industry as there has been a few forays into the cellular space over the past twenty years; however, up until now the value sale hasn't been there. As the cellular networks evolve, the number of small cells will drastically increase per mile to support 5G. The cellular leaders have been discussing the challenges of deploying small cells for the past few years. Where are they going to get the power, real estate, and backhaul to deploy 10X the number of small cells? The answer is fairly simple, the HFC network. FTTH was once a major threat to the industry and now it is turning into an advantage. What is missing with a FTTH network? **Power!** Power is a key factor in building a small cell network and what doesn't FTTH have? **Power!** In the past year, Sprint has signed two deals with MSOs to utilize the HFC network to deploy 4G radios. Expectation is that there will be over 20,000 small cells deployed this year to support Sprint's 4G LTE network. In rough numbers, if the LTE radios are consuming ~75W per radio, this is roughly 1.5MW of load being added to the HFC network. This type of load profile hasn't been contemplated and added to the network since circuit switch telephony was rolled out in the early 2000s.

Small cell is just one of the many wireless technologies which will create many revenue driven opportunities as well as a few operational challenges. 4G LTE is predominantly being utilized by the mobile network operator (MNO) and is on licensed spectrum. Citizen Band Radio Service (CBRS) technology is being investigated by the MSOs to provide small cell coverage in the unlicensed spectrum. WiFi will continue to rollout and either compliment or compete with small cell technology as the technology advances. Last but definitely not least, is Industrial Internet of Things (IIoT) technology, there are a number of Low Power Wide Area Network (LPWAN) technologies which will drive the adoption of machine-to-machine technologies.

It is an exciting time for the industry and with that excitement there is going to be significant opportunities and challenges to deploy.

Understanding the Network

The first step in tackling a new challenge is to understand where you are today. In this case, it is really about understanding the current state of your network specifically with regards to power and coverage. There are hundreds of thousands of miles of coax in existence in the US alone. This vast amount of coverage makes it ideal as a backbone for all of the new wireless networks. With the majority of the network continuously monitored, there is a lot of information available which can be utilized to better understand characteristics such as the load on the network, amount of current available for new services,

plant voltage, and heat mapping of the network. Almost all MSOs utilize an EMS (Element Management System) to manage and monitor their outside plant HFC power network. A proper EMS system, at a minimum should contain the following information:

- power supply location
- power supply output voltage
- power supply current draw
- information on backup battery run time
- outages per year.

For this paper, the focus will be on the initial three items as these will give a good sense for the type of services / load which can be supported by the current power grid.

1. Power Supply Current Draw

In the past twenty years, the typical HFC power supply has 15A of available current. Over this time period, roughly 83% of the power supplies are 15A with 11% supplying 18A of current. In the past five years, this metric is shifting to the 18A power supply. From 2013 to present roughly 67% of the power supplies are 15A and 26% are 18A. The remaining 7% account for the bookends, power supply sizes above and below 15A – 18A. In general, the trend is to deploy larger power supplies - thus increasing the amount of power available for the network. Last year's Expo paper, *'Fiber Deep and Lesson's Learned from the Field'* highlighted the need for additional power to support Remote PHY deployments. The same can be said for supporting new revenue generating services such as small cell, WiFi, IoT, and security and surveillance.

The best place to begin is to pull information from the EMS. A typical profile of the amp draw of a network is shown in Fig. 1. As can be seen from the graph, the typical load on a standard HFC power supply is 6A. The average for a segment of network tends to be between 5A and 7A. Taking a conservative approach to understanding the available power, assume that all of the power supplies are 15A power supplies. This gives roughly 9A of excess power. A good rule of thumb is to maintain 3A of overhead thus there are 6A available for next generation architectures or revenue generating services.




Figure 1 - Typical Current Draw

It is good to maintain a minimum level of overhead for future capacity and to support hot days. This will increase resistance in the coax, thus increasing the load on the power supply. 3A or ~20% of maximum current rating are good rules of thumb.

Once the amount of available current is determined, an operator can calculate the total amount of available power for a given network, section of a network, or a specific region. Assuming that a particular section of the network has 5000 power supplies, a simple calculation to understand the amount of available power is to multiply the # of power supplies (5000) by the amount of current available (6A) by the voltage (89.5V) which equals 2,685kW (Assuming 90V plant).

In the early 2000s, the HFC plant was built out to support circuit switched telephony. This is fundamentally the foundation for why the bell curve shown above is centered around 6A instead of 9A. The network was built to support a terminal at every house. On a positive note, it built a power base which provides the HFC networks a competitive edge to deliver the next generation of services. This is one of the main drivers why the telcos are interested in partnering with the MSOs to deploying the next generation of wireless devices. The HFC networks have a broadband powered network designed to support next generation services such as small cell, WiFi, IoT, security and surveillance.

2. Plant Voltage

The distribution of 60V versus 90V varies significantly between sections of the network. Based on previous research, the amount of 60V plant can vary from less than 5% to as high as 53%. A typical profile can be seen in Figure 2. It is worth investigating what is happening with the categories in the breakout section. Modern power supplies should have output voltages which are either 60V or 90V. If there is something else showing up in the EMS system, it is worth a look to understand what is going on. At first glance, it appears there is a portion of this network which is legacy 75V plant. This can be seen with the 5% in the “medium” category. There are 5.6% of the power supplies contained in the zero, low, high, and other categories which do not have a simple answer and require investigation.




Figure 2 - Power Supply Voltage Profile

In general, it is better to move the plant to 90V as quickly as possible. A plant running at 90V provides longer reach and better efficiencies over the coax. It is all about the I^2R losses in the coax. The higher voltage reduces the current thus reducing the load the coax has on the power supply. There are also a number of small cells which work from 75V to 90V. These small cells will not work on 60V plant without up converters. The primary reason that 60V plant is still in operations today is due to legacy actives in the network. Mostly these legacy actives are being removed over the years but they still exist and need to be accounted for during engineering work on new services. On the positive side, fiber deep / N+0 virtually removes the remaining amplifiers from the network which eliminates the reason for 60V plant. As plant is migrated to fiber deep, it is a general rule of thumb to transition it to 90V.

3. Heat Maps / Geographical Maps

Now that we understand both the amount of power available and the plant voltage, there are ways to apply this knowledge to better understand the capacity to deploy new services.

3.1. Heat Maps

Definition - A representation of data in the form of a map or diagram in which data values are represented as colors. (www.google.com)

A properly implemented and maintained EMS will also contain the location of the assets. Available current, locations, and voltage provide powerful information which can be represented utilizing heat maps. To start though, the key is to understand what question is being answered. Is the interest in understanding the amount of power available in an area or to understand the amount of coverage which can be overlaid with a specific technology? The following examples provide a quick idea of the capabilities of heat maps.

3.1.1. Current Availability

Figure 3 shows the amount of available power available at each power location. Each colored circle represents the excess power available at a site. The colors are represented as follows:

- Blue (>12 Amps)
- Green (9 – 12)
- Yellow (6-9)
- Orange (3 – 6)
- Red (<3)

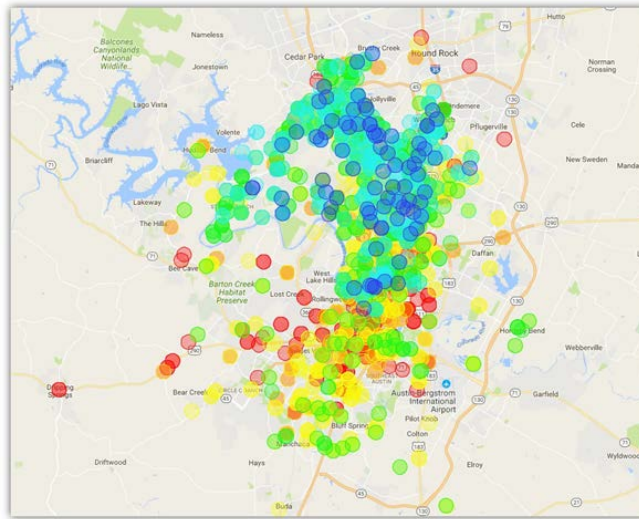


Figure 3 - Current Available Heat Map

Depending on the type of service being deployed, one can get a quick understanding if additional power is required. This heat map can also be used to trend power consumption overtime. As new services are deployed, it is beneficial to revisit the heat map and compare initial to current over regular 3 month / 6 month intervals to better understand how the load on the network is affecting the power system.

3.1.2. Wireless Coverage

The map in Figure 4 is an example of wireless coverage based at a power supply. Locations of the power supplies were utilized and then overlaid with a 2.5km circle to represent a LoRa wireless radio 20 feet in the air. The purpose for this heat map is to get a feel for coverage provided by mounting the LoRa radio at a power supply location. The advantages would be utilizing the existing location without having to cut into the strand as well as utilizing the existing DOCSIS modem at the power supply location. As can be seen in Figure 4, there is significant coverage provided in urban areas with minimal impact to the plant.



Figure 4 - IoT Coverage

The power supply cabinet is a valuable asset in an HFC network as it provides power, backhaul, and real estate in a single location.

Heat maps are fairly simple to utilize and implement. Programs such as Google Maps can upload Excel data files which will then plot the power supply locations over a geographic map. Color-codess, symbol types and sizes can provide a high level of flexibility and information for quick understanding of the information.

Planning Ahead

Every day there are technicians in the field performing some sort of maintenance on the plant, installing new power supplies, or mitigating outages. Are the technicians doing the same thing as they did yesterday or is there a master plan in place to upgrade the network every time a technician is at a site? Labor is one of the largest expenses an operator has; therefore, it is key to utilize that resource in the best way possible. There is no question that changes are coming for the HFC network. New partnerships are announced regularly between mobile operators and MSOs. The mobile operators understand that a key to their success is utilizing the power, backhaul, and real estate capabilities of the HFC network. To make this happen cost effectively, it is better to plan ahead and start upgrading the network every time a technician is in the field and working on the plant. There are a couple of simple items which if planned properly will help prepare the plant for the future demand.

4. DOCSIS™ 3.X Backhaul

Data Over Cable Service Interface Specification (DOCSIS™) has been critical to the evolution of the industry. DOCSIS™ has provided a means of communication which rivals fiber and supports today's and tomorrow's data demands. As with everything though, there is a significant number of legacy devices in the field which will not support tomorrow's demands. DOCSIS™ 1.0 and 2.0 devices have been the prevalent devices utilized in the network for monitoring power supplies. In most respects, this is all that is needed to monitor a power supply. The challenge is that more and more power supply locations are being utilized for new services. Services such as small cell, WiFi, IoT, and security and surveillance require more bandwidth than can be provided by a D1.0 or D2.0 modem. Along with the bandwidth, D3.X devices also provide additional proactive network maintenance (PNM) features which support troubleshooting and plant maintenance. Every time a technician is at a power supply site, the current transponder should be upgraded to a D3.X transponder. This is a simple way to prepare the network for next generation services.



Figure 5 - D3.0 Transponder

5. 6 Battery Cabinets

When deploying new enclosures, the standard methodology is to deploy the appropriate enclosure for the task at hand. In most cases, this would be a 3 battery enclosure to support a 15 Amp power supply with 2 – 4 hours of backup power. As new services are deployed, additional power will be required to support the new loads on the network. There are a few advantages to deploying all six battery cabinets. A six battery cabinet provides room to grow. If a 21 Amp or potentially a 24 Amp power supply is required in the future, the six battery cabinet can house 4 or 6 batteries to manage 48V or 36V battery strings (respectively).



Figure 6 - 6 Battery Enclosure

The larger enclosure also provides additional space for placing equipment for new services. The additional battery shelf can be repurposed for active equipment to support new services. An example of this can be seen in Figure 7. A global quad player has both HFC assets as well as wireless spectrum. Changing local regulations which restricted the placement of new wireless towers created difficulty for the operator to extend their wireless network. Their solution is to place 6 battery cabinets and utilize the 2nd battery shelf to house their wireless equipment. The solution provided multiple benefits. The primary benefit being the ability to extend their network in areas which were previously uncovered. Secondary benefits include: hiding the equipment from potential thieves, providing backup power for increased reliability, utilizing the D3 modem for cellular back and front haul, and placement of the antennas at the bottom of the enclosure. A simple solution to a complex problem which allowed the operator to build out their wireless network. Note: when looking at this type of solution, make sure to run thermal calculations for the specific region to guarantee neither the power supply nor the active equipment overheat.



Figure 7 - Enclosure with Wireless Equipment

As technicians are out in the network every day, make sure that their focus is to build a better and more reliable network. Every time a power supply is maintained, make sure that there is the latest DOCSIS technology installed. Whenever a new enclosure is placed, consider a six or eight battery cabinet along with a 240V service or the ability to upgrade to 240V service. These two simple upgrades will make it easier to support new B2B services.

Powering New Services

Network powering has been a growing point of discussion over the past two years. As the network re-architects itself from a centralized/headend centric platform to a distributed architecture with Remote PHY, the power distribution within the HFC plant is also changing. Initially, operators planned a N+0 strategy but were quickly overwhelmed with the number of CMTSs required within the headend or hub. It was quickly realized that remote PHY is a solution to the growing number of CMTSs within the headend. With remote PHY, the PHY layer of the CMTS is placed in the OSP close to the customer. This technology is great for pushing fiber deeper into the network and reducing the amount of equipment within the headend. Remote PHY also takes advantage of digital optics within the OSP thus reducing the impact of noise on the network. The key is managing the changing electrical load as it moves from the headend to the OSP. As discussed previously, the HFC plant typically has excess capacity today to support the rollout of remote PHY. Even though Remote PHY will increase load on the OSP roughly 40%, most operators only expect a 10% increase in the number of power supplies required to implement remote PHY. If operators were only planning on deploying remote PHY, the network would be relatively well positioned. The challenge being that operators are also actively pursuing new revenue generating services or cost avoidance initiatives. In the past twelve months, the first of these type of services can be seen through recent deployments of new technologies. In general, these type of services can be broken into four segments – small cell, WiFi, IoT, security and surveillance.

6. New Services

6.1. Small Cell

Small cells are by far the most talked about new service over the past twelve months. Operators are typically looking at two types of services. The first type of service is to support an MNO with deployment of either 4G or 5G small cells. This service would be a revenue stream for the operator while also allowing an MNO to quickly deploy additional coverage or new services. The benefit to the MNO is significant reduction in deployment timelines and cost. In 2018, over 20,000 4G small cells were deployed on a MSO network. The business opportunity is real and the MNOs are looking for partners to support their 4G and 5 G rollouts. In international markets, there are many times when an MNO is the operator (quad player) thus they get both advantages. The second type of service is to support an operators mobile virtual network operator (MVNO) agreement. With an MVNO agreement, an operator will deploy their own mobile service and pay the MNO for cellular and data usage. To reduce the fees to the MNO, an operator can deploy a wireless network on top of their HFC network. This is usually done via open spectrum such as the CBRS 3.5GHz – 3.7GHz band. There are a number of operators in trial reviewing the technology and understanding the business case for rolling out their own wireless cellular network. Key to success will be understanding the business model which fits the MSO.

6.2. WiFi

WiFi in the OSP has typically been a “sticky” service. It is a service which is designed to maintain customers but not grow revenue. This has done very well over the years and there are two new business cases which

are starting to emerge which make WiFi relevant again. Both business cases are based on cost avoidance. The first business case is targeted at operators who are managing inflating costs with their mobile provider. Deployments of WiFi at a power supply site provides a technician with quick access to the network for downloading and uploading reports, metrics, and service logs. The site can also be provisioned to support customer access. The second business case is cellular offload. Similar to the CBRS business case, operators are utilizing WiFi to offload data from the cellular network. The cost of the cellular data is significant enough that operators are working to get calls and data onto their own network as efficiently as possible. When planning the power consumption of either business case, the WiFi radios themselves are typically low power - in the range of 15 Watts to 25 Watts plus the DOCSIS modem, the total budget will be under 40 Watts. The challenges of WiFi are in the coverage range. When talking about WiFi, coverage range is talked about feet and not miles. The outdoor environment typically provides much larger coverage than indoor. Estimating a 250' coverage range per radio, a typical power supply would need to support 10 to 20 radios depending on the density of the network.

6.3. I-IoT

Industrial Internet of Things (I-IoT) is a growing market which allows the ability to transmit small bits of information over long distances utilizing very low power. The applications are endless from tracking your pet, to monitoring the flow of a river, the use cases are just starting to be conceived. For an operator, simple internal applications would be adding tamper sensing equipment to products in the outside plant or tracking vehicles via an internal network versus paying for the service. There are also a number of operators utilizing an I-IoT network to build a new business model. The challenge faced by the industry is the readily available Narrow Band IoT (NB-IoT) technology available to 4G operators. 4G operators will be able to deploy NB IoT quicker than MSOs; however, there is already proving to be unique business opportunities for both technologies. Key to success will be building networks quickly and gaining customers. There are a number of technologies being utilized to deploy an IoT network. Long Range Wide Area Network (LoRaWAN) is one of the leading technologies in the US for MSOs. From a powering perspective, this is a low power technology which has fairly significant coverage. On average, a LoRa radio set at 20 feet will cover roughly two miles of space in a rural environment. In a heavy urban environment, it will cover roughly one mile. For a typical HFC network, this would require one to two LoRaWAN radios per power supply. With an average of 80W per radio, an operator should plan on average roughly 120W of power per power supply.

6.4. Security and Surveillance

In today's world, the ability to offer security and surveillance is a significant business opportunity. Traditionally, this opportunity has been owned by the telecom operators. With the invention of new hardened modem gateways, the HFC network is becoming a much more preferred network for the placement of security cameras. In the typical telecom deployment, non-backed-up utility power is utilized - thus, when it is critical, emergency based situations and the power is out, so is the surveillance equipment. The HFC network provides a highly reliable, battery backed-up network which utilizes a single connection to both power and backhaul a camera. Power, backhaul, and real estate are essential elements to every security and surveillance business opportunity. It is hard to predict the number of cameras which could go on a network. For cold environments, a camera can use up to 60 watts. The 60 watts covers the camera power, the heater, and the pan/tilt/zoom function. For tepid environments, a camera will typically utilize about 30 watts.

Conclusion

HFC networks are poised to be the network of choice over the next ten to fifteen years. The key differentiating factor from standard fiber networks is power availability. With the next gen DOCSIS technologies rolling out, the HFC network will compete head to head with fiber networks and in general will provide more bandwidth than the typical customer will use. The key to success is understanding the business opportunities which are available in the market. Focus time on understanding each of the technology families which are emerging and how the network can support these technologies. It is also important to understand where the network is at today and what options does an operator have for supporting these new services. There are a lot of tools available today to help facilitate an understanding of the network today and also many options to plan ahead for the next generation network.

Abbreviations

B2B	Business to Business
CBRS	Citizens Band Radio Service
DOCSIS	Data Over Cable Service Interface Specification
EMS	Element Management System
FTTH	Fiber to the Home
HFC	Hybrid Fiber Coax
I-IoT	Industrial Internet of Things
LoRaWAN	Long Range Wide Area Network (Type of LPWAN)
LPWAN	Low Power Wide Area Network
MNO	Mobile Network Operator
MSO	Multiple System Operator
MVNO	Mobile Virtual Network Operator
PNM	Proactive Network Maintenance
SCTE	Society of Cable Telecommunications Engineers

Orchestration: What Is Really Behind This Overloaded And Overused Term?

An Overview of What Makes An Automation And Orchestration System

A Technical Paper prepared for SCTE•ISBE by

Alon Bernstein
Distinguished Engineer
Cisco Systems

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Automation vs. Orchestration.....	3
Where Do I Start ?	3
Customer facing and Provider facing	5
Management Layering	5
Domains	6
Workflows.....	7
Declarative vs. Imperative.....	8
Never repeat the same error	9
Bill of Materials	9
Introduce failures.....	9
How many automation use cases ?	9
So how did the industry survive without automation and orchestration ?	10
An Opportunity	10
Conclusion.....	11
Abbreviations	11
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1 ONAP closed loop automation.....	4
Figure 2 Devops.....	5
Figure 3 TMN layering.....	6
Figure 4 Cable Modem scan workflow.....	7

Introduction

Automation ! Orchestration ! These two magical words solve everything... with a click of a button OPEX shrinks to nothing. But not so fast...the software that drives automation and orchestration behind the scene is complex and highly customized and it cannot eliminate the inherent complexity of a service provider network, the data center and the outside plant. It can only help manage this complexity. It can't eliminate existing processes, but it can turn them from manually operated ones to a software operated processes.

Furthermore, the transition to the cloud native world, which is even more distributed and large scale (hence complex) than existing systems, makes automation and orchestration become more than OPEX reduction tools. It's impossible to operationalize these highly distributed systems without automation and orchestration.

This paper will attempt to separate fact from fiction when it comes to automation and orchestration. It will outline the steps required to go from a swivel chair process to an automated one and along the way explain the key concepts and layering required for automation.

Automation vs. Orchestration

The terms “automation” and “orchestration” tend to be bundled together, to the point that they seem interchangeable. The distinction this paper offers is as follows:

- Automation is the overall framework for replacing manual processes with software.
- Orchestration is the specific task of coordinating activities across several domains (we discuss the term “domain” later in this paper).

Where Do I Start?

In the following sections we'll account for the processes that need to be automated. That would help focus the discussion on the scope of automation.

Let's pick a very simple example and follow it through in order to help map high-level abstractions onto concrete actions. Say we want to ping all the cable modems in a service group to validate that they respond within a well-defined time range. Most likely there are already vertical applications that can do it, this paper would explore how you would build such an application with an orchestration/automation mindset.

A good place to start from is ONAP (Open Network Automation Platform, see ref 1). Even if you don't plan to use ONAP it's still a good reference model for our discussion.

ONAP itself is a large framework, and this paper will not discuss it. The only part we explore in this paper is the ONAP automated life-cycle described in Figure 1.

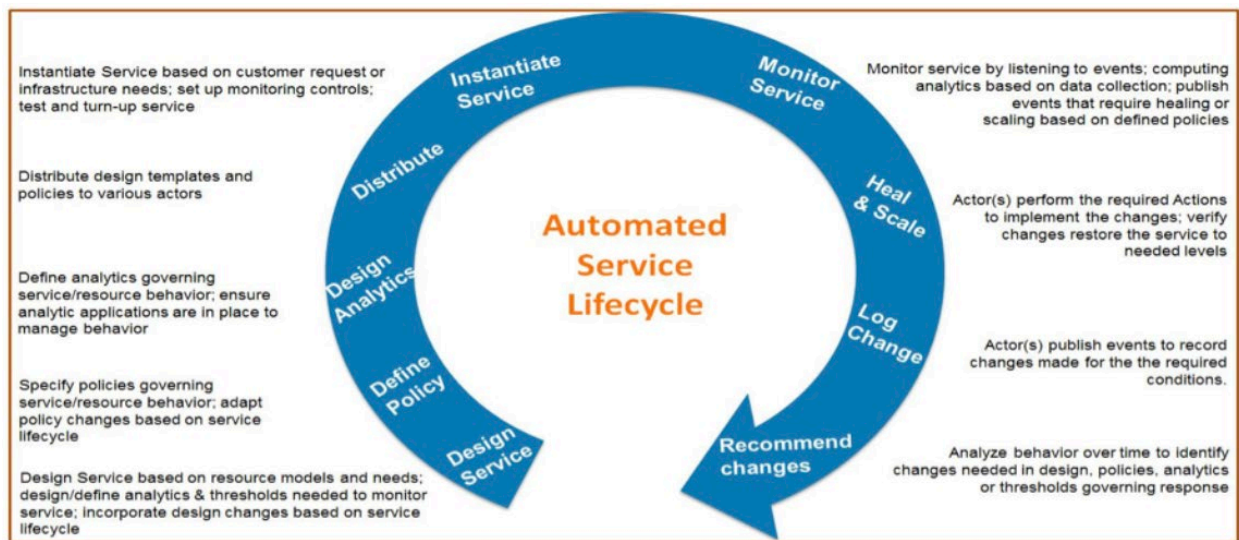


Figure 1 - ONAP closed loop automation

Let's go over the various phases here and in one line describe the key actions and requirements that have to be answered. One confusing thing about Figure 1 is that the service itself (ping) has an automated lifecycle of its own. For example, the ping service requires memory and there needs to be a policy what to do if memory allocation for the service fails. However, the focus of the study here is the user-facing service to be implemented:

1. Design service: define the scale, goal, key modules.
2. Define policy: what to do if a ping fail ? think of the actual pass/fail criteria for the ping, e.g. an average ping time greater then 20ms is a fail.
3. Define analytics: a collection of min/max/average of the ping times to modems could qualify as "analytics data" for our simple example
4. Distribute design template and policies to various actors: we can choose the network control center as the actor for are use case. More complex use cases may have multiple actors, each one with its own set of credentials, authorizations and capabilities.
5. Monitoring: that's when we actually activate the service and start monitoring the ping times
6. Based on the results of the ping take corrective action. This can trigger a whole set of other services, e.g. changing the modulation profile for a cable modem could help ping issue but it's a whole separate service.
7. Log the changes
8. Based on the analytics and observations make a long-term proposition on how to improve service.

Note that there is no orchestration or automation mentioned in each of the steps. The whole cycle is the automation of the service. The intelligence part of the automation is mostly under the analytics and recommended change parts but can be spread out to any component.

It's useful to compare this service life-cycle to the classic devops (development and operations) cycle diagram (see Figure 2):

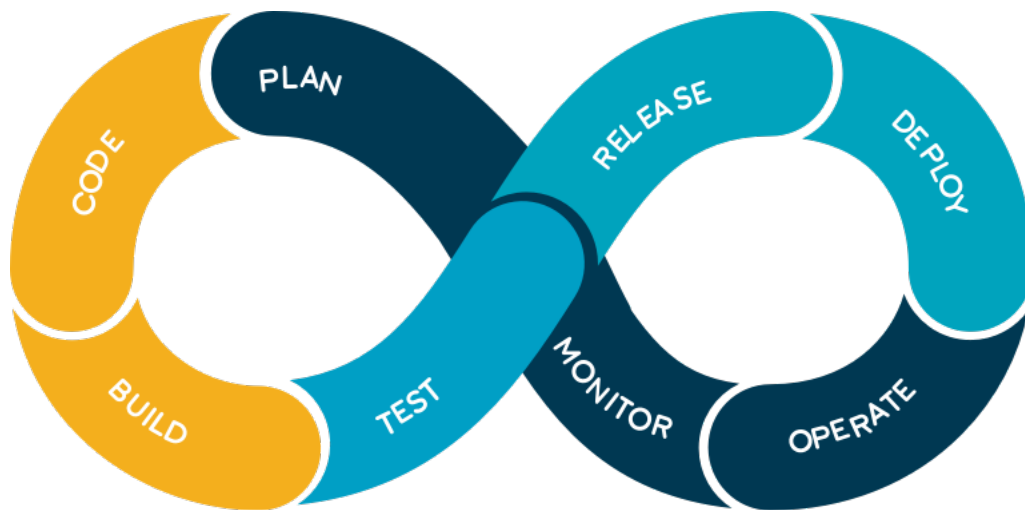


Figure 2 - Devops

As you can see the service life cycle can overlap the service life cycle, where the left-hand side of the figure (the “dev”) correspond to the initial stages of the service life cycle and the right hand side of the service (the “ops”) maps into the later stages of the service lifecycle .

Customer facing and Provider facing

Some of the orchestration and automation services are internal to the service provider and some are for consumers. For example, the process of installing an RPD (Remote Phy Device) is an operator facing service while providing an L2VPN service to a business customer is naturally a customer facing service. In principal the same tools can be used to automate either the provider facing or customer facing services, however there are several key differences:

1. Provider facing services are focused on what is known as “day zero” and “day one” operations where day zero is the initial install of the equipment (in the virtualization case the instantiation of the infrastructure and based containers). Day one is typically the base configuration needed to get the system up. The provider facing services are all CAPX and OPEX because they are concerned with bringing up the service provider “platform”.
2. Customer facing are those that generate revenue for the service provider. These are the Day two type of operation. Note that there are operations that are not configured directly on network devices but rather are signaled e.g. voice calls.

Management Layering

A model that has been trusted and deployed for over 20 years is the TMN (ref [4]):




Figure 3 - TMN layering

Although this is an old model (the words “ATM” and “ISDN” sprinkled as examples throughout the recommendation hint at its age) the key concepts survived the test of time:

1. Network elements – the actual network devices. The big change since the 90’s is that we have a new generation of virtual/cloud-native network elements, but functionally they are still “network elements”.
2. Element management – the scope of managing a single element, for example, software update or configuration change for a specific device. In the age of SDN the concept of “controller” has been introduced and a good part of the device specific management is done by a controller. Note that even if a network element is cloud-native, and even if we use JSON for configuration, it still needs to be managed and this fits into the TMN model nicely.
3. Network management – this layer manages connectivity across a whole network comprising of various devices. This is the layer where a lot of orchestration and automation are required
4. Service management – the elements, element management and network management are all “provider facing” the actual services that the operator sells, for example, high-speed data for cable customers, are at this level. Note that this level extends beyond the network because it includes items such as credit card processing to make sure the customer paid for the service.
5. Business management: the layers on top of the services can include information that is required to run the business, for example, tracking customer satisfaction.

Domains

A complex network can be built out of several domains. For example, a cable network is comprised of the HFC plant, The DOCSIS protocols, access network, core network, back-office etc. To create services, one has to work across all these domains.

What makes the domains complex is that each one of them has its own set of tools and set of experts. The skills and tools used to manage the HFC are very different than those used to manage a data center. As a result, there is little to no sharing of expertise, and a swivel chair process where one organization (aka “domain”) passes tickets to another to perform an action. The ticket typically sits in a queue and the result is that service activation can take days if not weeks.

It’s the orchestration part of automation that takes care of operating end-to-end solutions across domain and does the magic of building consistency across a wire array of tools and expertise that each domain has. This is where workflow automation comes into place.

Workflows

A workflow is similar to a flow chart. It lists a set of actions and their dependencies. A standard called BPMN2.0 (Business Process Model and Notation, see ref [3]) was defined to capture workflows in a uniform way. A simple workflow for the example on cable modem scanning can look as the following:

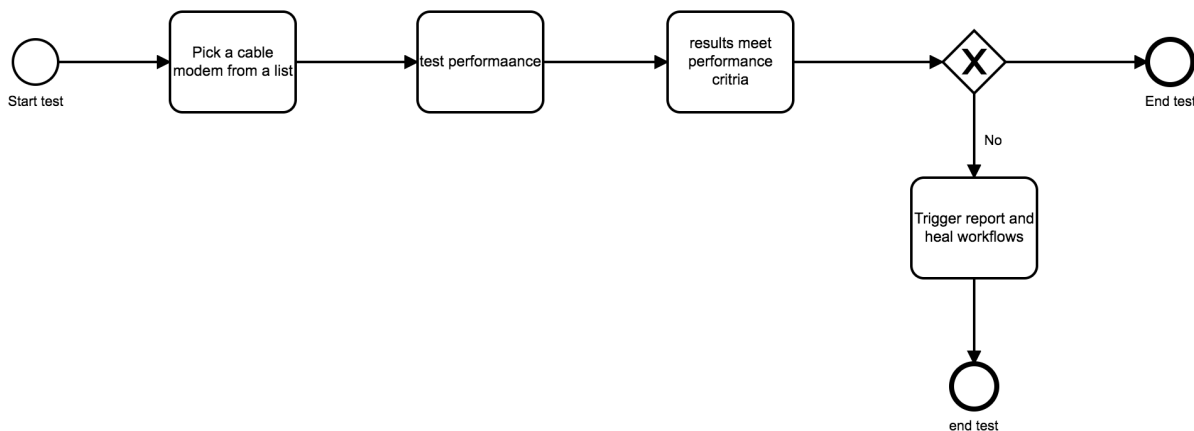


Figure 4 - Cable Modem scan workflow

Even with this simple workflow a key point can be highlighted: The workflow itself is not intelligent. It is essentially a state machine that calls APIs (Application Programming Interface) and takes actions based on the results of these API calls. The actual performance test could be a simple ping or some more complex operation. The workflow just calls the module that performs it. In other words, the workflow defines what needs to be done, not how to do it.

Another way in which workflows are useful is clarifying the use of APIs. In many cases publishing an API list is not sufficient, and a workflow example clarifies how to use the APIs correctly, e.g in what order they should be called.

There are several tools that can take a pictorial representation of a workflow and trigger actual API calls and scripts with it, so that the BPNM2.0 is not only a modelling tool, it can be an actual run-time service.

One question that comes up is the difference between writing a script vs. using a workflow engine. At the end of the day both can get the job done and it's about using the right tool for the right job. Having said that, if the answer to the questions below is positive then a workflow may be a preferred approach:

1. Do you need to communicate the process to non-software-engineers ? If yes then the workflow is your friend because it does not require coding skills. In a devops kind of environment it's a way to bring developers and operation experts closer together. Note that in a pure devops world the operators should be coders as well, but having said that in cable we have physical devices and outside plant components that are not made out of software....
2. Do you need to work across several domains ? If so the workflow provides a neutral environment that is not tied to a specific domain and focuses on API calling only.
3. Do you need an operation at the “what needs to be done” as opposed to “how to do it”. The workflow excels at the former.

How does workflows relate to orchestration ? It is the view of this paper that the workflows are a way to implement orchestration, so they are one and the same. Note that there is no magic here. Orchestration is accomplished only looking at existing operational practices, describing them in workflows and starting the process of replacing each one of the boxes that has a manual operation with a software based one.

Declarative vs. Imperative

There are two ways to achieve a network management goal:

1. Declarative: Detail the end-result, or end-state, and have the system automatically do what is needed to achieve the end goal/state. This method is referred to as “declarative” as the user declares the desired goal/state. Some very successful software solutions are declarative in nature. For example, in Kubernetes a user defines the scaling requirements (e.g how many replicas of a service are needed) and Kubernetes takes care of the actual scaling and placement of the service onto the CPUs in the data center in order to achieve the declared scale figure.
2. Imperative: Detail step by step how to reach a certain goal/state. This is more aligned with the work flow approach and is called “imperative”.

The declarative model seems more attractive – what could be better than simply declaring what you want and let the system magically figure it all out ? There are two key observations about this in practice. The first one is that any service definition is in essence “declarative” because it states what needs to be done, not how to do it. For example, “I want a 20mbps upstream and 20

mbps downstream high-speed data service” is as good as declaration as any, nothing really new here. The second observation is that somehow, somewhere, someone needs to write the software to turn the declaration into reality. In other words, a script or program or workflow has to be written anyway, it’s only a question of having a layer of “declarative” abstraction above it (which is indeed a useful thing to have), but nothing happens by magic.

Never repeat the same error

One of the key benefits of automation is repeatability and consistency. With manual processes there is always the risk that errors are the result of an operator mistake. Even in cases where failures occur intermittently a consistent a repairable process make the debug process and root cause analysis easier.

Automation is not static. Once an error occurs in production, and even once the error is fixed by means of software or configuration change, it’s possible to write a script that validates the error does not happen again. In that sense the automation resembles a set of unit-tests in a software development process and fits into the devops model.

As the list of verification scripts gets longer it becomes a risk-reward question of how many of them needs to be run, and a choice one can made when submitting a change to a “bakeoff” setup vs. in production deployment as to how many validation scripts need to be executed.

Bill of Materials

With automation any change can be treated in a similar way, whether it’s a software change or a configuration change, because any type of change carries a risk and can cause unexpected behavior. A “Bill Of Materials” (BOM) for a change can include the software modules that are changed and/or the configuration change, along with various validation scripts.

Introduce failures

Service provider networks are built for high-availability, but when are these high- availability systems tested in production ? The answer is that they should not be tested only when an unplanned failure occurs. Errors can be injected on propose. This concept has been popularized by Netflix in what’s called “chaos monkeys” (see ref[5]), and can be considered part of the service automation loop.

How many automation use cases?

At this point we get to the hype around orchestration and automation. We can apply the automation lifecycle to both customer facing and provider facing services, and in each one of the layers of the 5 layers of the TMN model and across several domains (let’s assume 5 domains for a cable network). That’s $2 \times 5 \times 5 = 50$ different areas where orchestration automation can apply and each can have the 9 stages of life cycle management for a total of $50 \times 9 = 450$ cases , and of course each one of them can have 10s if not more of use cases. It becomes clear that there is no

“finger snapping” that will make it happen, and if implementation resources are limited it’s a question of figuring out which processes are the ones that benefit automation the most. What automation represents more than reduction in human resources is a shift to devops. The brains to manage the network are still needed, but they invest their brainpower in creating scripts and workflows rather than repeating the same action again and again. The number of automation workflows and scripts that have to be written also enforce the devops vision that everyone, dev or ops, need to be able to code, there is just too much coding work to be done !

So how did the industry survive without automation and orchestration?

Very simple. There was automation and orchestration all along, but with 3 caveats:

1. It was not called “automation/orchestration” it was called network management.
2. It was applied only when absolutely necessary. E.g. from the get go it was clear that deploying cable modems will not scale without automation and as early as the DOCSIS 1.0 spec the groundwork for automating the cable modem registration was put in place.
3. It was done in verticals: let’s say an operator needed to keep track of some condition in the cable plant. In most cases a vertical application with its own set of collectors, data analysis tools, GUI, etc. was built with little sharing of other applications. The cable modem registration mentioned above is another example of a one-off mission specific vertical. Today we might have treated the cable modem registration as part of the IOT framework and use a whole different automation framework which would have been more consistent with other “things” that have to be managed.

What the new automation and orchestration advocates relative to the old way of doing network management is:

1. Build the orchestration and automation as applications on top of a common platform (e.g. ONAP).
2. Use open models such as YANG.
3. Use open source tools and code as a “standardization” framework as opposed to standard bodies paperwork.
4. Automate and orchestrate everything. Anything that has to be done more than once needs to be automated. Anything that requires coordination across multiple entities requires orchestration.

An Opportunity

The next wave of change in the SP networking space is cloud-native based networking devices. In the cloud native world functions are broken into micro-services and there can be tens, hundreds and thousands of those micro-services. At this scale automation and orchestration is no longer a nice-to-have-OPEX-reduction play. It’s a must because it would be impossible to manage this level of complexity and distribution by hand. This presents the opportunity to start

with automation at the cloud native area and as tools and expertise are built, start branching out to other domains.

Conclusion

Orchestration and automation are not new. What is new is the use of open source tools, open model and workflows to support automation and the absolute necessity of having automation in the cloud world. Hopefully this paper helped in giving an overview of the key components and architectural concepts needed to create an automated network.

And last but not least, in the spirit of devops, every developer is an ops person and every ops person can write scripts and workflows. This is the key to automation.

Happy coding !!!

Abbreviations

API	Application Programming Interface
BOM	Bill Of Materials
BPNM	Business Process Model and Notation
Capex	Capital expense
DevOps	Development and Operations
ONAP	Open Network Architecture Platform
OPEX	Operating Expense
RPD	Remote Phy Device
TMN	Telecommunications Management Network
YANG	“Yet Aounter Next Generation” data Modelling language

Bibliography & References

1. <https://www.onap.org>
2. ONAP closed loop Automation: <https://onap.readthedocs.io/en/amsterdam/guides/onap-developer/architecture/onap-architecture.html>
3. <http://www.bpmn.org>
4. TMN reference model : <https://www.itu.int/rec/T-REC-M.3010-200002-I/en>
5. Choas Monkeys : <https://github.com/Netflix/SimianArmy/wiki/Chaos-Monkey>

Practical Deployment Lessons of a Centralized 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

Richard J. Walker

Vice President, Engineering & Info Technology & Services
Shared Services
2700 Oregon Road Northwood, OH 43619
1-419-724-3735
rjwalker@sharedsvcs.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Architecture Description	3
Deployment Details	5
Real-World Considerations	6
1. Compute and Network Resource Location	7
2. Networking	8
2.1. Layer 2 or Layer 3 Converged Interconnect Network	8
2.2. Traffic Prioritization and Capacity Management	9
2.3. Redundancy	9
3. Timing.....	9
3.1. Remote DOCSIS Timing Interface	9
3.2. DOCSIS Latency	10
4. Operations	11
Conclusion.....	11
Abbreviations	12
Bibliography & References.....	12

List of Figures

Title	Page Number
Figure 1 - I-CCAP, PHY Shelf and Remote PHY Node Deployment Comparison	4
Figure 2 - Traditional I-CMTS Deployments Across Multiple Hub Locations.....	4
Figure 3 - vCMTS Deployed in a Centralized HFC Architecture.....	6
Figure 4 - Adding Remote PHY Nodes to a PHY Shelf Deployment	8

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. One possible starting point is virtualizing a CMTS in a centralized deployment model in a headend or hub.

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 centralized architecture. Areas of consideration when enabling virtualization via shifting from a hardware-based CMTS to a vCMTS will include the use of Remote PHY as a protocol/specification between a vCMTS Core and a PHY Shelf in a centralized facility (e.g. headend or hub), and usage of IEEE-1588. The ease of expandability will be highlighted, including node splits with PHY shelves, as well as Remote PHY Nodes (RPN). Lastly, the usage of streaming telemetry will be explored in comparison to legacy monitoring techniques.

Architecture Description

The architecture described in this paper is made possible by recent 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 May 2018 with the tenth revision of the specifications.

“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).”¹ 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 virtual CMTS (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” and “Real-World Deployment of a Virtual Cable Hub”, 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 Distributed Access Architecture (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 two examples where the RPD and RF may be located in the network, in the headend/hub or in an optical node.

Figure 1 shows an I-CCAP deployment architecture, as well as Centralized and DAA deployment architectures, both of which use the Remote PHY signaling to communicate between the vCMTS and the RPD (existing in the PHY Shelf and RPN). Remote PHY signaling includes the Downstream External PHY Interface (DEPI), Upstream External PHY Interface (UEPI) and the Generic Control Plane (GCP).

¹ Remote PHY Specifications, CM-SP-R-PHY-I10-180509, pg. 10

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.

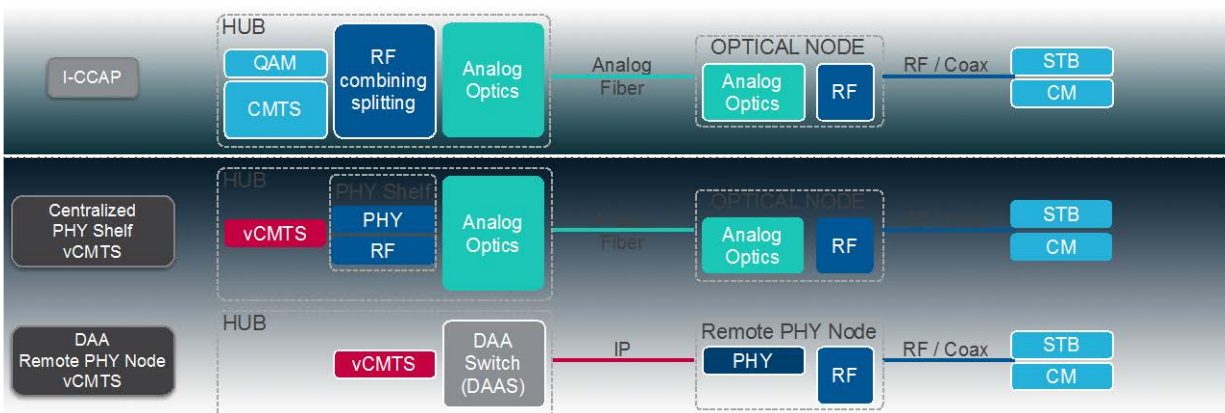


Figure 1 - I-CCAP, PHY Shelf and Remote PHY Node Deployment Comparison

This paper focuses on a Centralized deployment architecture, whereby another variation of the PHY is located in secondary hubs. The term “Centralized” is used to describe this deployment architecture as the existing HFC and analog optics in the field are leveraged. Figure 2 shows a traditional deployment with I-CMTS chassis deployed in each hub location within the operator’s footprint.

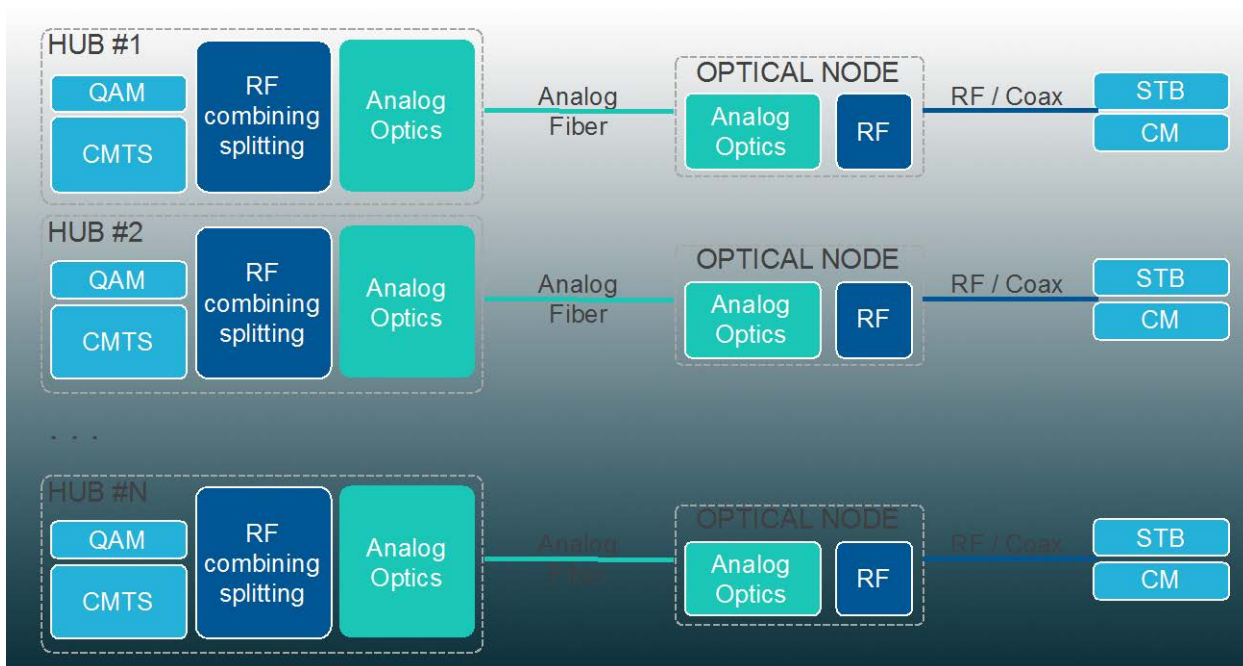


Figure 2 - Traditional I-CMTS Deployments Across Multiple Hub Locations

This traditional I-CMTS deployment was the starting point before the transformation to deploy vCMTS and PHY Shelves. The transformation to a vCMTS deployment was driven by the benefits of virtualization, such as²:

- Reduced space, power and cooling
- More frequent and shorter development cycles
- Sustainable capacity growth, elastic scalability and increase flexibility
- Improved Total Cost of Ownership, including reduced operational (OpEx) and capital expenditure (CapEx)

In particular, with vCMTS running on COTS x86-based 1-RU servers, the entirety of the CMTS Core functionality was consolidated to a single centralized location. The PHY Shelves with the RPD functionality were deployed in secondary hubs where the legacy I-CMTS chassis were previously deployed. These PHY Shelves connected to the existing HFC infrastructure. This transition delivered immediate value while leveraging all other legacy infrastructure, such as broadcast and VOD services being processed and delivered by existing EQAMs.

Further future value may be attained as this is a step towards deploying DAA, with initial small scale and eventual ramp up of RPNs connected to the same centralized vCMTS hub location over an IP network. More servers and capacity can be added, as needed, to scale at the chosen pace over time.

Deployment Details

The deployment details covered will describe the ending point after centralizing and consolidating the vCMTSes into a single hub, deploying PHY Shelves in secondary hubs and leveraging the existing HFC infrastructure.

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. All vCMTS Cores are located in a single, centralized hub and are connected to the Converged Interconnect Network (CIN).
2. Core Routing Engines (CRE): the switch fabric connecting the vCMTS Core servers to the rest of the CIN. All the CREs are located in a single, centralized hub and are connected to the CIN over a Layer 3 network.
3. Core routers: large-scale core routers connecting the access network with the core backbone network, located in a single, centralized hub.
4. Distributed Access Architecture Switches (DAAS): aggregation switches connecting the RPDs in the secondary hubs with the vCMTS Core servers (via the core routers), located in the secondary hubs.
5. PHY Shelves and RPDs: many RPDs share a single highly-available chassis in a PHY Shelf, which connect to the vCMTS Core servers over the CIN, and output RF over the existing HFC infrastructure. The PHY shelves are located in different secondary hubs.
6. 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."³

² Real-World Deployment of a Virtual Cable Hub, pg. 5

³ Remote PHY Specifications, CM-SP-R-PHY-I10-180509, pg. 27

Figure 3 shows the devices and connectivity between the vCMTS Core servers over the CIN, which is “the network between the CCAP Core and the RPD. 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.”⁴ In this deployment type, the CIN traverses between the centralized vCMTS hub and the secondary hubs over multiple switches, leveraging a Layer 3 network.

One of the benefits in this deployment type is that different RPDs and PHY shelves may connect to the same or different vCMTS Core servers in the vCMTS hub, which allows flexibility and scalability in growing capacity. For example, multiple secondary hubs (and the RPDs in those hubs) can be connected to the same vCMTS Core server.

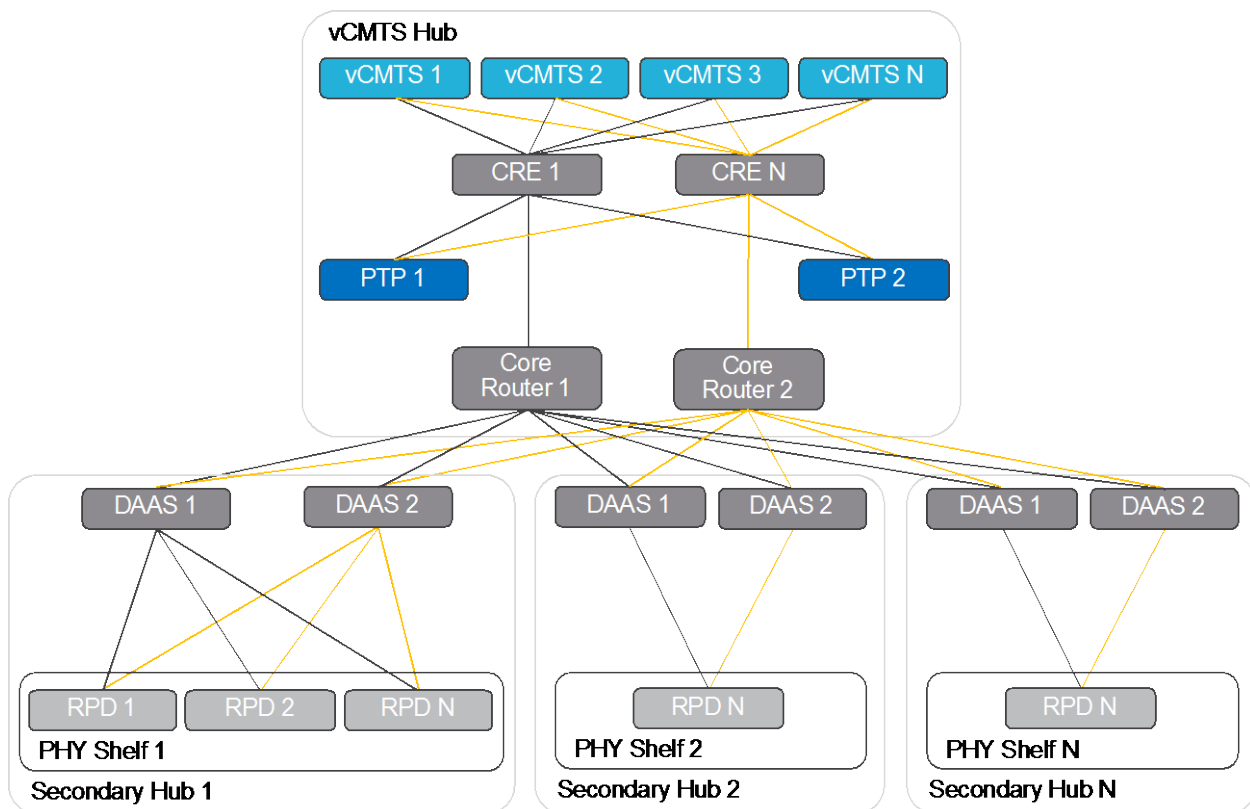


Figure 3 - vCMTS Deployed in a Centralized HFC Architecture

Real-World Considerations

While there are many benefits and opportunities with this new Remote PHY-enabled vCMTS in a centralized HFC deployment architecture, there are real-world considerations which should be considered. In particular, the following topics are described: compute, networking, timing, RPD or PHY type and operations.

⁴ Remote PHY Specifications, CM-SP-R-PHY-I10-180509, pg. 24

1. Compute and Network Resource Location

One of the key considerations when evaluating this architecture is where the CMTS Core is located, for two reasons. The first is that with I-CCAP, the CMTS Core and RF had to be co-located. I-CCAPs are chassis-based products and are scaled between 13 and 16-RU per I-CCAP chassis. There was no choice in the past as I-CCAPs would be deployed at each hub location, 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 provides the operator an opportunity to decide where the CMTS Core should be located.

The second reason for CMTS Core location determination 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 consolidated in an existing hub location, which would be the single hub location for vCMTS Core servers now and in the future. With a single vCMTS location, future expandability of distributed RPNs can be accomplished over the same deployment architecture, connecting each RPN to the appropriate DAAS in a nearby secondary hub, which may already be connected to an existing PHY Shelf. This expandability option is very efficient, as PHY Shelves and RPNs can connect across the CIN to existing vCMTS Core servers and DAAS, providing the operator the option to expand the compute resources of vCMTS Cores and network resources of the DAAS as capacity demands.

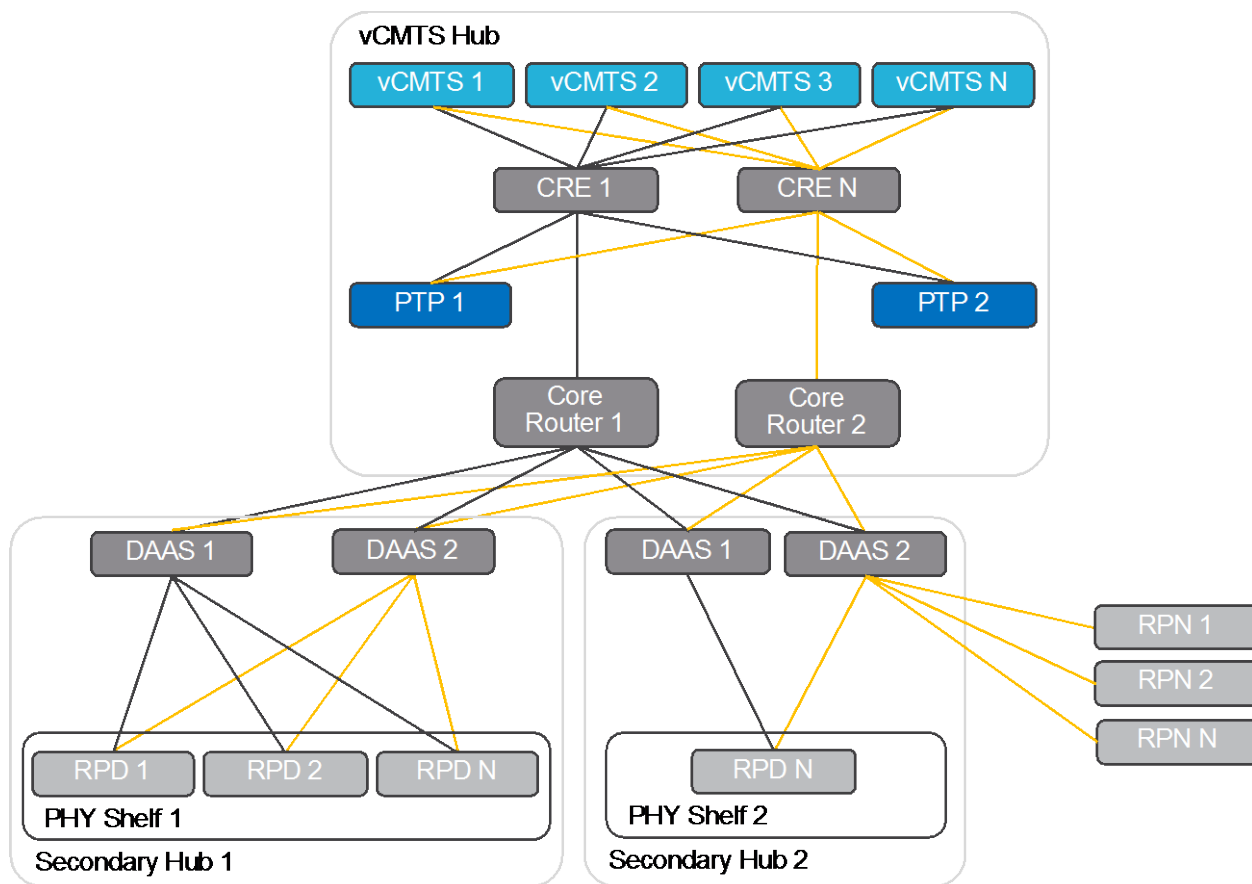


Figure 4 - Adding Remote PHY Nodes to a PHY Shelf Deployment

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.

2. Networking

2.1. Layer 2 or Layer 3 Converged Interconnect Network

There are many real-world lessons learned in the area of networking, as separating the vCMTS Core from the RPD not only spreads the CIN network elements across a variety of locations, and as described before, but it also spreads across different network resource instances and different network configurations. In other words, different deployments need to consider how many hops exist between the vCMTS Core and the RPD, and if the CIN is layer 2 or layer 3.

In this real-world deployment experience, layer 3 networking was selected for its benefits of greater flexibility and protection, as well as it is simpler to manage and scale between the CRE, core router and DAAS with this deployment architecture. Layer 3 networking has embedded control plane capabilities which provide these benefits to the operator, as well as easier configuration across the different network devices.

2.2. Traffic Prioritization and Capacity Management

Another important deployment consideration is that this architecture may expose shortcomings in the existing operator network and force improvements in IP network robustness and traffic prioritization. It is simpler to deploy when it can be guaranteed that the CIN is congestion free. However, if this can't be guaranteed, capacity and congestion management of the CIN network devices are necessary.

Capacity management is important to determine when expanding CIN network resources are necessary and installing these network resources before congestion occurs. If congestion does occur, capacity management and traffic prioritization need to be calculated and evaluated. 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 to the extent that is impactful to critical control packets, such as DOCSIS MAP MAC Management Messages (MMMs), no amount of traffic prioritization will resolve the impact to stable network operations.

2.3. Redundancy

Network redundancy is another deployment decision with many options available, including:

1. Link redundancy
2. Chassis redundancy
3. Line card redundancy within a chassis

The decision criteria for each network device type (CRE, core router and DAAS) is based on failure domain size (how many subscribers are impacted), the built-in physical redundancy within each network device and the cost of the incremental networking equipment. In this deployment example, lower-cost CRE and DAAS network elements are 1-RU devices and are deployed with redundant links and with two redundant instances to the vCMTS Cores and PHY Shelves, so there is protection in case a single CRE or DAAS fails. On the other hand, the core routers are carrier-class chassis-based devices with redundant line cards and don't require link redundancy to each CRE, since a separate network path via another core router is available in the CIN should a core router fail.

3. Timing

3.1. Remote DOCSIS Timing Interface

With all Remote PHY deployments, timing specifications such as R-DTI need to be adhered to, regardless of the location of the RPD. "The MHA v2 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].”⁵

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 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 which is the highest quality or “best” clock within the network, in case the grandmaster clock quality is compromised or fails.

3.2. DOCSIS Latency

Latency is an important parameter in any QoS system, including DOCSIS-based solutions, keeping in mind that end-to-end latency measured between a subscriber’s CPE and the end-point server extends beyond the DOCSIS portion of the network. DOCSIS has different latency in the downstream and upstream directions, which is a common trait of point-to-multipoint access technologies.

I-CCAP latency measurements are well known and are described in the MULPI specification, with a minimum latency associated with the DOCSIS MAC protocol for best effort traffic. The minimum latency budget consists of the worst-case round-trip propagation delay (variable), queuing delays within the CMTS, processing delays within the CMs and downstream delays caused by the PMD-layer framer and FEC interleaver.⁶

The expected and measured behavior is that ping times increase linearly as distance increases between the CMTS Core and CM, attributed to the one-way propagation delay of 0.8 msec. per 100 miles (0.5 msec. per 100 km.). However, maximum bandwidth is not impacted as this distance is increased.

There is no difference in DOCSIS MAC and PHY processing delays between I-CCAP and Remote PHY deployments. However, this topic is relevant with Remote PHY deployments as the CMTS Core may be centralized further in the network than the RF (located with the RPD) and may extend the distance, as well as the number of network hops between the CMTS Core and the CM. The transit time in both directions is comprised of these two contributing factors: propagation through the fiber and coax, and the transport processing delays. Real-world experience has shown that the dominant contributor to increased

⁵ Remote DOCSIS Timing Interface, CM-SP-R-DTI-I07-180509, pg. 6, 16

⁶ MAC and Upper Layer Protocols Interface Specification, section 7.2.1.6

latency within the DOCSIS portion of the network is in fact the propagation delay, with the increased number of network devices in the CIN minimally impacting the transit time and associated latency.

As with an I-CCAP, any Remote PHY-based CMTS Core needs to account for the MAP advance time supporting the maximum transit time between the CMTS Core and the CM.

The impact of end-to-end latency and the contribution of DOCSIS latency within the network to best-effort services is beyond the scope of this paper.

Nevertheless, with Remote PHY-based deployments, the round trip and latency performance are more in the operator's control in comparison to I-CCAP. The operator can maintain and, in some cases, improve latency performance of the system. Conversely, the operator may deliberately choose an installation scenario that will reduce latency performance in order to gain in other aspects, such as centralizing vCMTS Cores in a single hub location.

4. Operations

Technology is great, but it must be supported by real-world operational practices which support deployment of any new technology. In the experience of deploying a vCMTS in a centralized HFC architecture, there is an immediate opportunity to start by consolidating monitoring of the vCMTS Cores in a single consolidated hub or single network operations center (NOC). This consolidation allows monitoring of any service group within the set of vCMTS Core servers. Additionally, with a software-based virtualized CMTS Core, there are many points of inspection in the software which 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.

Conclusion

Cable operators have an existing footprint which they continue to grow and improve, but can't be overhauled overnight. Technologies such as Remote PHY and virtualization continue to extend the tool set which cable operators can use to stay competitive with capacity growth demands and challenging market environments. This paper has described a starting point which takes an existing I-CCAP centralized HFC deployment and changes a few architectural elements in the network with immediate benefits while paving a path to DAA. This approach leverages an installed base of vCMTS Core servers and will ultimately support a hybrid set of centralized and distributed dense PHY shelves, smaller and more remote PHY shelves and remote PHY nodes.

Abbreviations

CapEx	Capital Expenditure
CCAP	Converged Cable Access Platform
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
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

DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-115-, May 9, 2018, Cable Television Laboratories, Inc.

Matatyaou, Asaf. Real-World Deployment of a Virtual Cable Hub. Publication. San Jose: Harmonic, 2017. Web.

Matatyaou, Asaf. Transforming the HFC Access Network with a Software-Based CCAP. Publication. San Jose: Harmonic, 2015. Web.

Modular Headend Architecture v2 Technical Report, CM-TR-MHAv2-V01-150615, June 15, 2015, Cable Television Laboratories, Inc.

Remote DOCSIS Timing Interface, CM-SP-R-DTI-I07-180509, May 9, 2018, Cable Television Laboratories, Inc.

Remote PHY Specification, CM-SP-R-PHY-I10-180509, May 9, 2018, Cable Television Laboratories, Inc.

Predicting Service Impairments from Set-top Box Errors in Near Real-Time and What to Do About It

How Machine Learning Can Preempt Calls and Tickets: Results from a Trial

A Technical Paper prepared for SCTE•ISBE by

Justin Watson

Senior Manager, Product Management
Comcast
Philadelphia, PA
Justin_Watson@comcast.com

Roger Brooks

Chief Scientist
Guavus, Inc.
San Jose, CA
roger.brooks@guavus.com

Andrew Colby, Office of the CTO, Innovation Lead, Guavus, Inc.

Pankaj Kumar, Senior Manager Analytics, Guavus, Inc.

Anant Malhotra, Principal Analytics Engineer, Guavus, Inc.

Mudit Jain, Principal Analytics Engineer, Guavus, Inc.

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
1. Problem Statement and Challenges	3
1.1. Customer Care Constraints.....	3
1.2. Challenges: Data, Calls and Tickets	3
2. Machine Learning.....	4
3. Customers at-Risk and Risk Drivers	5
3.1. Trial Set Up	5
3.2. Outsized Risk Factor Impact.....	6
3.3. Drivers of Risk.....	6
3.4. Model Accuracy.....	7
3.5. Risk Drivers and Incidents	8
3.6. Treating Risk Proactively	8
3.7. X1 Errors and Incidents.....	9
4. Ticket Problem Code Prediction.....	9
4.1. One Suspect Set, Many True Positives	9
4.2. High-Risk Sub and Highly Probable Resolution.....	10
Conclusion.....	11
Bibliography & References.....	11

List of Figures

Figure 1 – Predictive Maintenance Use Case.....	4
Figure 2 – Problem Taxonomy Generated from Ticket Data	5
Figure 3 – Aggregated Actual (red) and Predicted (green) Calls.....	7
Figure 4 – Aggregated Actual (red) and Predicted (green) Tickets	8

List of Tables

Title	Page Number
Table 1 – Risk Buckets: Predicted and Actual Calls vs. Total Subs with Errors (percentages)	6
Table 2 – Risk Buckets: Predicted and Actual Tickets vs. Total Subs with Errors (percentages).....	6
Table 3 – Predicted and Actual Calls Associated with Risk Drivers (percentages).....	7
Table 4 – Predicted and Actual Tickets Associated with Risk Drivers (percentages)	7
Table 5 – Percentage of Incidents Correlating Strongly with Risk Drivers	8
Table 6 – Calls Saved by Addressing Risk Drivers (percentage).....	9
Table 7 – Tickets Saved by Addressing Risk Drivers (percentage).....	9
Table 8 – Problem Code Prediction: 13 Classes, 3 Sets of True-Positive Tickets	10
Table 9 – Proactive Steps: Accurate Predictions on Resolutions.....	10

Introduction

Getting ahead of subscriber problems is a difficult but powerful way to reduce costs and increase customer satisfaction. This paper describes a proof of concept (POC) trial that harnessed machine learning and Comcast X1 service impairment data to identify at-risk subscribers and risk drivers, and to further indicate next best-actions to take in response to the predicted issues.

Machine learning is well-positioned to address the blind spots of customer support teams. It can be architected to scale to tens of millions of simultaneous event streams and handle real-time, complex predictive analytics. The highly accurate analytics used in this trial enabled us to identify subscribers potentially affected by impairments responsible for generating 36 percent of calls and 46 percent of tickets. Only 5 percent of the device population drove these care events. To enable action, our analytics also identified the associated risk drivers. In another exercise, we predicted a large quantity of true positive tickets per year related to 13 newly clustered ticket classes, with known resolution paths, and associated 57 percent of those tickets with single problem codes. (Note: All tickets referenced in this paper are technical tickets.)

Shared among internal stakeholders, these kinds of insights can drive numerous benefits. They can enable service providers to proactively address technical problems of subscribers; reduce the number of calls, tickets and truck rolls as a result; and more quickly resolve impairment events that do arise. The net result is reduced costs and improved customer experience.

1. Problem Statement and Challenges

1.1. Customer Care Constraints

Facilities-based service providers know that providing excellent customer care takes tremendous effort. From NOCs to support personnel to maintenance technicians to software, equipment, and fleets of vehicles, it requires an extensive combination of resources to take care of customers. These costs add up. Scaled to millions of subscribers, customer care becomes a big number, one that even gets the attention of financial analysts.

Yet despite the investment and efforts, visibility remains limited. “We’re still reactive” - that’s what one operator admitted to the author of a paper on customer experience delivered at this conference last year [Cunha.] The assessment still applies widely across the industry. Past patterns help us schedule resources, network monitoring and telemetry provide device-level insight, and integrated ticketing systems gather what we know onto one screen, but it remains difficult to get ahead of the customer.

How to address the technical problems of subscribers who are likely to call, before they call, was the guiding question of a two-part data analytics project, undertaken in late 2017 and early 2018, whose results we share in this paper. Before getting to the study and results, let us first point to some challenges posed by the data and discuss how machine learning is well suited to meet them.

1.2. Challenges: Data, Calls and Tickets

The problem today is not too little data. In addition to information drawn from customer database and network telemetry common to most MSOs, Comcast can leverage the X1 set-top box. This next-generation IP video platform, which was launched five years ago, not only proved popular among subscribers, it has also generated tremendous amounts of data. Those include the large numbers of error streams that provide the foundation of this analytic exercise. (See Figure 1.)

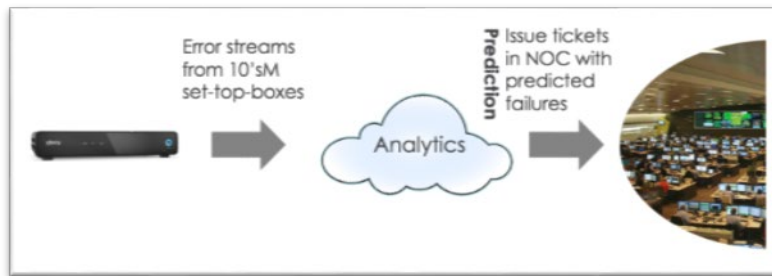


Figure 1 – Predictive Maintenance Use Case

To anticipate customer calls, you need to not only process massive quantities of data, but also do so at near real-time speed. Legacy approaches associated with manual correlation, centralized processing and storage bottlenecks are poorly equipped to deliver the desired results.

Handling the data in this case also means prepping them for analysis. The challenge is common to other industries: raw data rarely arrive in neat taxonomies. Streaming off the X1 set-tops are 1-2 thousand types of errors codes that must be grouped, hierarchically ordered and contextualized within meta information linked to particular boxes.

The calls and ticket have other limits. Several tickets may arise from a single customer, each of which may have an incorrect or correct assessment. Because tickets bear subscriber account numbers, not MAC-level addresses, there is the further problem of not knowing which box in a multi-box home is impacted, forcing the machine to learn only from single-box homes. Given customer delays, the time stamp on tickets also may not convey when the impairments actually occurred.

Then there are the tickets that never arrive. Between 10 and 40 percent of subscribers who have some reason to call in to complain, do not do so. Sometimes called “silent sufferers,” they may just be apathetic or distracted or waiting for things to improve. In this study, however, they are also conservatively tagged as ‘incorrect care event predictions’, even though the errors codes did correlate to incidents. Consequently, this has the adverse effect of negatively impacting our accuracy score and should be kept in mind when reviewing the final accuracy scores.

A final issue with tickets is that some percentage are non-technical. As noted at the outset, for present purposes, drawing largely from X1 error codes, along with outage and reconnect data, this analysis is concerned with technical tickets, not those associated with billing or other types of problems.

2. Machine Learning

Machine learning is a good fit for these kinds of data. Rightly applied, this field of AI has numerous use cases. In a paper last year, several Comcast colleagues discussed using machine learning to simplify field operations, in particular, to detect spectral impairment [Dorairaj, et al.] In another paper, a Guavus colleague pointed to how a combination of machine intelligence and operational analytics is an effective way to assure virtualized networks and services [Sundelin].

Machine learning uses various tools across the entire process, from exploration and feature engineering to training, evaluation, tuning and deployment. The solution used in this exercise is a modular data ingest and analytics platform, containerized for the cloud, and highly scalable over a distributed architecture. Among the algorithms employed are the following:

- Quantile Transformer - enables features to follow a more uniform distribution
- Linear Support Vector Classification - allows selection of features based on weights
- Spectral Clustering - reduces spectrum of the similarity matrix before clustering data
- Probability Chunking - provides probability-based segregated chunks of predictions
- Hierarchy of Models - combines multiple models, all trained at each level of Tree-based taxonomy

One aspect worth underscoring is machine-learning's orientation toward probabilities. Unlike more deterministic techniques that provide yes/no answers, machine learning generates probabilistic results. The tradeoff between precision (correctness of the prediction) and recall (breadth of coverage of the predictions) is a common way to both measure the accuracy of the predictions and align the models to the business value sought. A lot depends on how averse one is to false positives.

3. Customers at-Risk and Risk Drivers

3.1. Trial Set Up

To preemptively address technical problems of individual subscribers, the first step is to identify who is at risk of calling or ticketing. Doing so also involves looking at the risks that are driving those incidents.

Who are the at-risk customers? Given that X1 set-top errors reflect actual impairments in service, we can assume those data are associated with some percentage of customers who do call. The initial data selected for this project followed from that premise. Our population of data came from all subscribers who had any X1 set-top errors over a 7-day period.

In this same exploratory phase, we assessed the problem portions of the ticket data, or problem codes. Using several machine-learning algorithms, including Spectral Clustering, we organized these codes into a taxonomy with the leaves of the taxonomy tree representing "classes" containing similar error data. (See Figure 2.)

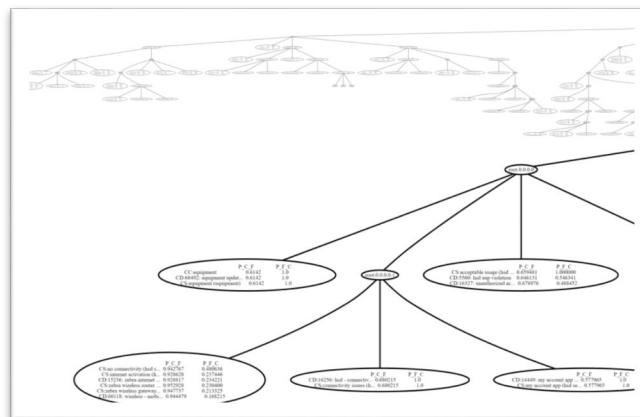


Figure 2 – Problem Taxonomy Generated from Ticket Data

The next phase was identifying those in our subscriber pool most at risk of calling or ticketing. We combined the stream of X1 errors with reconnect, historical outage and other contextual information about individual subscribers and boxes. Feature engineering was then done to reformulate the properties of the data to best align with the use case and the machine's ability to interpret the data. We seeded the

model with all calls and tickets and those aggregated features from three rolling windows of 1hr, 24hrs and 168 hrs. The output was an hourly prediction, aggregated for 7 days.

To assess greater or lesser risk, we used risk bucketing, a machine-learning method for correlating variables, in which similarity is calculated by rank score comparison and then displayed, largest group to smallest, with the smallest bearing the greatest probabilities.

3.2. Outsized Risk Factor Impact

The outcome of these exercises reveals striking results. The multi-bucket model, indeed, shows increasing number of predicted calls and tickets as group-sizes diminish. Bucket 4 warrants the most attention. It represents only 5 percent of total subs with errors, but drives 36 percent of actual calls and 46 percent of actual tickets. (See Tables 1 and 2.) This set offers the best target for preemptive action and cost reduction. Knowing that a certain percentage of issues can be handled proactively, we can envision reducing a number of these calls, tickets and other support on a recurring basis.

Table 1 – Risk Buckets: Predicted and Actual Calls vs. Total Subs with Errors (percentages)

Risk Bucket	Predicted	Actual	Subs w/Errors
0	8%	6%	35%
1	15%	14%	25%
2	19%	18%	20%
3	24%	26%	15%
4	34%	36%	5%
total	100%	100%	100%

Subscribers in Risk Bucket 4, while accounting for only 5% of the total subscribers with set-top errors, drive 36% of actual technical calls.

Table 2 – Risk Buckets: Predicted and Actual Tickets vs. Total Subs with Errors (percentages)

Risk Bucket	Predicted	Actual	Subs w/Errors
0	6%	5%	35%
1	15%	13%	25%
2	18%	16%	20%
3	19%	20%	15%
4	42%	46%	5%
total	100%	100%	100%

Subscribers in Risk Bucket 4, while accounting for only 5% of the total subscribers with set-top errors, drive 46% of actual technical tickets.

3.3. Drivers of Risk

To dive deeper into high-risk Bucket 4, we built another ranking that correlated risk drivers, i.e. our newly organized problem codes, with tickets and calls.

In descending predictive power are X1 errors, previous outages, previous calls and device model. At the top for calls are X1 errors and previous outages, which drive a roughly equal number of actual calls, and together account for three-fourths of the total. (See Table 3.) For predicted Bucket 4 tickets, X1 errors drive an even larger number of actual calls, and twice the events as previous outages. (See Table 4.) In both cases, previous calls outrank device model as explanatory features.

Table 3 – Predicted and Actual Calls Associated with Risk Drivers (percentages)

Risk Drivers	Predicted	Actual
X1 Errors	41%	37%
Previous outage	34%	37%
Previous calls	19%	20%
Device model	6%	6%
Total	100%	100%

Table 4 – Predicted and Actual Tickets Associated with Risk Drivers (percentages)

Risk Drivers	Predicted	Actual
X1 Errors	51%	48%
Previous outage	23%	21%
Previous tickets	18%	22%
Device model	8%	9%
Total	100%	100%

3.4. Model Accuracy

As one can see from Tables 1–4, predicted calls and tickets track closely with the actual ones. Figures 3 and 4 provide additional evidence for the high accuracy of this model. They also indicate the cyclical, time-series nature of these data and the 7-day length of both training and evaluations periods.




Figure 3 – Aggregated Actual (red) and Predicted (green) Calls




Figure 4 – Aggregated Actual (red) and Predicted (green) Tickets

3.5. Risk Drivers and Incidents

Turning from prediction accuracy to causal correlation, we look again at the drivers identified in Tables 3 and 4, here from a slightly different angle. The X1 Errors remain worthy of attention. For tickets, 20 percent of high-risk Bucket 4 subs correlate strongly with X1 Errors, higher than the other three drivers combined. For calls, 14 percent correlated strongly with X1 Errors, still higher than the next highest driver. (See Table 5.)

Table 5 – Percentage of Incidents Correlating Strongly with Risk Drivers

Risk Drivers	Tickets	Calls
X1 Errors	20%	14%
Previous outage	9%	11%
Previous calls	7%	9%
Device model	3%	2%

3.6. Treating Risk Proactively

The end game of this exercise is proactive measures. The first area we examined was what percentage of at-risk calls and tickets might be saved if any given risk drivers were removed. With the X1 placing high in the previous risk analyses, drivers related to X1’s Cross-Runtime Environment (XRE) or possibly the Reference Design Kit (RDK) software stack were logical candidates. Moreover, the focus of this study was not to address outages that occur beyond the control of the operator, rather issues that could be resolved remotely. There will always be previous calls; and device failure is, in part, simply related to device lifecycle.

For both tickets and calls, this learning exercise revealed some top candidates for feature removal. We have anonymized the actual error codes, but if risk drivers A and E were successfully addressed, then 42 percent of calls and 53 percent of tickets would be saved, respectively. (See Tables 6 and 7.)

Table 6 – Calls Saved by Addressing Risk Drivers (percentage)

Risk Drivers	Call saved if feature is removed	Fraction of subs impacted	Call propensity impacted subs
A	42%	91%	0.092
B	31%	95%	0.063
C	14%	100%	0.038
D	12%	100%	0.048

Table 7 – Tickets Saved by Addressing Risk Drivers (percentage)

Risk Drivers	Ticket saved if feature is removed	Fraction of subs impacted	Ticket propensity impacted subs
E	53%	92%	0.128
F	17%	100%	0.039
G	14%	100%	0.043
H	13%	100%	0.039

3.7. X1 Errors and Incidents

Another look at X1 data disclosed more relationships. By extracting the at-risk subscribers for given risk drivers, we discovered that about 20 percent of tickets are highly correlated with XRE errors, and 8 percent highly associated with particular X1 errors, impacting thousands of predicted tickets. We also found that 12 percent of calls were highly correlated with XRE errors, and 5 percent highly associated with particular X1 errors, also impacting thousands of predicted calls.

Removing a high-risk feature affects both those who call or register tickets and those who do not. In the first case, it eliminates the time and effort of contacting customer care; but for everyone it removes a service impediment, arguably improving service. The benefit redounds to the service provider, as well, reducing the number of calls received and a percentage of truck rolls sent to subs who might have simply needed a software patch.

4. Ticket Problem Code Prediction

4.1. One Suspect Set, Many True Positives

A separate exercise in addressing technical problems of subscribers who are likely to call, before they call, involved ticket problem code prediction. Extracting distinct and homogenous ticket problem code classes, we used a clustering algorithm to improve classification. Then over a four-day period, we trained the model, with an evaluation period on the fifth day. Hourly predicted results, aggregated for 1 day, were based on a correlation of X1 errors, outages, and reconnects with ticket problem codes.

The result, once again, reveals that certain risk factors have outsized influence. The overall prediction was 13 ticket classes associated with a large population of true positive tickets per year. From that large sample of tickets emerged three class-related sets (See Table 8):

- Set 1: Eight classes with a single problem code; total tickets covered, eight; true positive tickets predicted per year, 57 percent of total.
- Set 2: Four classes with two problem codes; total ticket problem codes covered, eight; true positive ticket predicted per year, 33 percent of total.
- Set 3: One class with three problem codes; total ticket problem codes covered, 3; true positive tickets predicted per year, 10 percent of total.

In Table 8, we report those classes for which our prediction model achieved the highest precision, irrespective of the associated recall. The tradeoff between precision and recall is adjustable and can be calibrated separately for each of the problem classes. The key takeaway, however, is that a majority (57.4 percent) of the total true-positive tickets are associated with eight ticket classes that have a single code each. Conveyed to the right stakeholders, that kind of insight can help drive both quicker and deeper resolution of issues, reducing average care handling time and costs.

Table 8 – Problem Code Prediction: 13 Classes, 3 Sets of True-Positive Tickets

Predicted ticket classes, problem codes	Precision	Recall	True positives/year
Class 1, code A	98%	21%	Set 1: 57.4%
Class 2, code B	91%	38%	
Class 3, code C	100%	74%	
Class 4, code D	100%	28%	
Class 5, code E	98%	50%	
Class 6, code F	90%	27%	
Class 7, code G	81%	32%	
Class 8, code H	86%	5%	
Class 9, code J, K	99%	20%	Set 2: 33.0%
Class 10, code L, M	83%	37%	
Class 11, code N, P	92%	37%	
Class 12, code Q, R	73%	4%	
Class 13, codes S, T, U	96%	35%	Set 3: 9.6%

4.2. High-Risk Sub and Highly Probable Resolution

This true-positive exercise yielded further insight into proactive problem-solving. The method is to identify those ticket classes for which accurate predictions can be made regarding an appropriate resolution, and then link the fix to an individual subscriber. One example is a provisioned modem for an incorrect boot file, for which our classification analytics found the most likely solution among several for this problem. (See Table 9.)

Table 9 – Proactive Steps: Accurate Predictions on Resolutions

Sub Id	Risk score	Potential ticket problem code	Possible resolution codes
XXX	0.44	Incorrect Boot file	Provisioned modem: 80.96% Customer equipment: 9.21% SIK to customer: 2.32% Excluding Voicemail: 1.71% Reconfigured: 0.12%

Conclusion

The cable industry has invested heavily in customer care and frontline technical support. But so far, we have yet to get very far ahead of the subscriber – and know who is going to call before they call, and why. One development capable of changing the game is the combination of machine learning and rich operational data, such as error data from the IP-enabled X1 set-top box.

As our results from the POC exercise show, there is potential for considerable insight and follow-up actions. This falls into three domains:

- A. We found that 5 percent of subscribers, in an identifiable high-risk category, are driving 36 percent of all technical calls and 46 percent of all technical tickets. We were also able to associate 57 percent of total true-positive, predicted tickets with one large set of ticket classes notable for having a single problem code.
- B. For a significant number of the devices exhibiting errors, we were able to identify the likely diagnosis that would have been made by the agent, as reflected in the ticket if that person sought support. This information can be passed to agents to reduce their call times and better inform the actions taken, thus improving customer experience and making more efficient use of internal resources.
- C. It is not a stretch to say that these findings – if shared among leaders in customer care, finance, quality engineering and other teams – could lead to coordinated strikes against the leading risk drivers and preemptive actions. The outcome, again, could be notable reductions in operating costs (calls, truck rolls, customer care, etc.) and increased customer satisfaction.

Bibliography & References

Cunha, G. Approaches for Proactively Managing Customer Experience and Reducing OPEX in a Cable Operations Environment. SCTE-ISBE 2017.

Dorairaj, S., and Chris Bastian, Bernard Burg, and Nicholas Pinkernell. Simplifying Field Operations using Machine Learning. SCTE-ISBE 2017.

Sundaresan, K., and J. Zhu. Access Network Data Analytics. SCTE-ISBE, 2017.

Sundelin, A. Leveraging Machine Intelligence and Operational Analytics to Assure Virtualized Networks and Services, SCTE-ISBE, 2017.

Preventing Unwelcome Guests in Your Home

Securing IoT Devices Within the Home Network

A Technical Paper prepared for SCTE•ISBE by

Dave Belt
Technology Evangelist
Irdeto
Conifer, Co.
(303) 653-7647
dave.belt@irdeto.com

Table of Contents

Title	Page Number
Table of Contents	2
1. Introduction.....	3
2. Devices and Their Management	3
2.1. Managed Devices (CPE).....	3
2.2. Unmanaged Devices (COAM).....	3
2.3. Operator Influenced COAM (Hybrid).....	4
3. Device Security Features and Their Accessibility	4
3.1. Application Code Signing	4
3.2. Secure Boot.....	5
3.3. Secure Micro	5
3.4. Hardware Root of Trust and Device Identity	5
3.5. Debug Detection	5
4. Device Threats and Their Impact	6
4.1. Data Privacy	6
4.2. Identity Theft.....	6
4.3. Access to Backend Services.....	7
4.4. Device Disablement	7
4.5. DDoS Attacks	7
4.6. Ransomware	7
4.7. Theft of Service	8
5. Conclusions.....	8
5.1. Know Your Devices	8
5.2. Understand Device Security Features	8
5.3. Provide Access Based on Trust.....	8
Abbreviations	9
Bibliography & References.....	9

1. Introduction

With the explosion of IoT devices within the home network, new challenges are emerging for MVPDs from a device management perspective. In the traditional cable model, Operators had full control of the system from plant to device providing the ability to not only manage system security, but also the end user experience. With the emergence of the TVE experience, operators have slowly been accepting and eventually embracing the presence of COAM devices within the ecosystem.

Herein we will look at the devices and threats within the Home Network with the end goal of understanding these devices and their capabilities and enabling the creation of practical policies for their management. First, we will look at device management models and how this affects our control over their behaviour. Next, we will walk through the type of security capabilities that we expect within trusted devices. Finally, we will look at the types of threats that can occur within these device ecosystems and lay some guidelines for managing them.

2. Devices and Their Management

Home networks are built from devices. These devices collect data, present information and services as well as provide the underlying infrastructure over which the home network operates.

The nature of the threat a device poses is dependent upon its functionality, physical implementation, as well as its management model. A key distinction of the Operator's control over a device is influenced primarily by its management model, in particular Consumer Premise Equipment (CPE) vs Consumer Owned and Managed (COAM) devices.

Ultimately the device is the primary entry point into any ecosystem and the most accessible point to the hacker. The securing of these endpoints is pertinent to ensuring the integrity of the ecosystem as a whole.

2.1. Managed Devices (CPE)

Consumer Premise Equipment is the device model most familiar to the MVPD and has been the backbone for most of their history. By providing the customer with their service consumption devices the operator maintains comprehensive control over the device's user experience, but more important to this discussion, the device security.

Operators have been driving new security requirements into these devices over the last several decades and have the ability to adopt new security technologies quickly in response to emerging threats.

In general, when considering the network as a whole, CPE devices present a lower threat risk simply due to the control the Operator has over them. This of course puts the burden on the Operator to implement the security functionality however the only way to have full control is to own the system.

2.2. Unmanaged Devices (COAM)

Unmanaged devices are produced by CE Manufacturers presumably unrelated to the Operator. These devices are produced for general public consumption but due to an intersection of services, the device needs to interact with the Operator's ecosystem in some manner.

The first encounter between unmanaged devices and Operators was the emergence of OTT video. Consumers had a desire for the TV Everywhere experience and as such the market demands the consumption of Operators' video on COAM devices.

The lack of relationship between the Operator and COAM Manufacturer creates an inherent lack of control over the device capabilities and security. In order to get their service on the device, the Operator must create an app which is frequently subject to the Manufacturer's submission rules. Simultaneously, the Operator may have little visibility into the security controls built into the device at manufacture, leaving the Operator at risk to the limitations of the device.

2.3. Operator Influenced COAM (Hybrid)

A third management model is rapidly emerging with Operators, especially with the deployment of IoT based ecosystems. Operator influenced COAM devices fall somewhere in between CPE and independent COAM.

Increasingly Operators are creating business relationships with existing CE Manufacturers. The Manufacturer gets increased volume due to inclusion and promotion within the Operator's offerings, and the Operator gets some level of influence over the device features.

While not as controlled as CPE, this model provides greater control for the Operator, but is still subject to the whims of the Manufacturer's broader roadmap. Due to the appeal of the CE devices to the consumer, this model is becoming increasingly common.

3. Device Security Features and Their Accessibility

MVPDs through their traditional business model of video delivery have come to expect, and indeed have pioneered many device level security features. These features have since made the jump from STBs to common consumer devices.

While security features may be present on a device, access to those features may vary greatly by platform. Frequently these features are only available at the hardware or OS level. This is done by the manufacturer not to reduce 3rd party integration capabilities, but rather to secure their own device ecosystem.

It is in this context that the CPE vs COAM management models make the difference. In a CPE model, security features are fully accessible to Operator integration whereas in the COAM model, the security features may not even be known to the Operator.

Knowledge and access to available security features is critical to building secure home network ecosystems.

3.1. Application Code Signing

Application code signing consists of applying a secure digital signature to a binary software image. Prior to execution, that signature is verified and on success the application is permitted to run.

Code signing is pertinent to verifying that the intended code is running on the device as opposed to rogue software potentially injected by a hacker. This is a key tool to preventing malware attacks on devices. More advanced architectures implement dynamic code signing, where signatures are checked while the application is running. This prevents runtime code injection attacks on the device.

Many COAM devices have code signing implemented, however it is mostly for the system software running on the device. If an operator is integrating at the application level, they will likely not have access to this security functionality.

3.2. Secure Boot

The secure boot is built on the code signing technology previously mentioned. When a secure device boots up, this boot process first performs a signature check of the OS and firmware image prior to boot.

In order for this boot to be truly secure, it must reside within a secure hardware processor, including the asymmetric public key used for code verification. Exposure of this to the software level allows a hacker to circumvent the signing process and potentially inject rogue code. A secure boot is considered to be a bare minimum for developing secure devices in contemporary embedded devices.

The boot processes for COAM and CPE devices are similar, however on CPE devices the Operator has the ability to have their integrated code signed with the system code.

3.3. Secure Micro

The secure micro consists of an area of the device processor that is very tightly protected. Operations within the secure micro are prevented from being exposed at the software level preventing a hacker from access. All cryptographic and keying operations are then performed within this space. A secure micro is a basic requirement in order to implement a secure boot described previously.

White Box Cryptography is a solution that allows the creation of a secure environment in software, but is best used in environments with no or inaccessible secure micro.

This is the main feature of COAM devices that will likely not be accessible to the application developer. The secure micro is frequently used by CE manufacturers to secure their platform, but access to it requires proprietary APIs from the Si manufacturer.

3.4. Hardware Root of Trust and Device Identity

Utilizing the secure micro described previously, a Hardware Root of Trust consists of a unique key or set of keys programmed into the device. This provides the ability to target a unique device with a secure non-tamper identity.

Within any device management ecosystem, device identity is crucial to the system management. By having a secure unique identity, the ecosystem operator can be assured that the devices attaching to the ecosystem truly belong there. Simple device IDs are easily spoof-able by a hacker but with the use of an Authentication Key Exchange (AKE), the identity can be securely identified.

Using this identity, the device can also be uniquely field targeted with firmware or credential upgrades. A payload package encrypted based on the root is uniquely encrypted for that device and cannot be extracted by others.

3.5. Debug Detection

Embedded devices intended to run a known set of firmware and software, frequently have debug detection implemented. Upon attaching a debugger to the device, the firmware can take various evasive actions from reporting it to completely disabling the device. Debug detection is one of the main tools used to keep the hacker out of the device to begin with.

Penetration of the device itself gives the hacker access to the device's data, its operation and potentially to the ecosystem's back end systems.

4. Device Threats and Their Impact

4.1. Data Privacy

As mentioned early on, home networks are made of devices that collect information, present information and potentially take action on that information. The protection of the consumer's data is paramount to protecting a device-based ecosystem. Failure to do so destroys the consumers trust and subsequently that of the ecosystem.

Likely the greatest concern from a customer perspective is that of in home video. The proliferation of video based IoT devices of late, from security cameras to nursery monitors to even the Ring doorbell which is an outside device, provide the hacker with unprecedented access to the household. On a certain level the customer is more violated by a stranger seeing the inside of their home than if their credit card was stolen.

Access to video content provides a distinct physical threat to the consumer in general. This data is easily used for a potential home attack due to the ability to monitor the comings and goings of the resident of the home as well as being able to map out a home and identify items of interest for theft. Monitoring of minors coming and going within the household provides a useful tool for potential predatory behaviour, as a potential predator can easily establish the day to day patterns of the household.

Digital thermostats provide similar data due to their intended usage. When one leaves the home, the thermostat is turned down and then turned back up when coming home. Analysis of this "Big Data" provides a detailed record of the comings and goings of the household.

Encryption technology over the wire is the obvious way to protect content between devices and servers and TLS/SSL technology has become ubiquitous for maintaining these leaks. Failure to do so is considered a Noob development error in contemporary devices.

Protection of data on the device itself is more complicated and requires a layering of the security technologies discussed herein. The implementation of code signing, and a secure boot ensures the integrity of the device itself. Implementation of debug detection prevents hackers from gaining access and encryption of sensitive data on the device prevents access assuming a hacker has gained control of the device.

4.2. Identity Theft

Not dis-similar to data theft, identity theft consists of the acquisition of personally identifiable information (PII) for the purposes of impersonating that individual. PII has historically been very sensitive for Operators due to regulatory controls built around it.

Increasingly personal devices have one or more login credentials which are subject to compromise. Access to any of these can compromise one or more of the many accounts we all have online including financial, social and digital communication. Simultaneously these devices are caching credit card data within them to allow service transaction with a smooth user experience.

Due to our increasingly connected habits, identity theft is on the rise and will likely continue this trend. Media hype around the topic adds awareness to the issues for consumers, but rarely provides them with the tools to actively combat identity theft.

4.3. Access to Backend Services

If a hacker wants to gain access to an ecosystem's backend services, the first place he will target is one of the devices on the ecosystem. By scraping and disassembling the firmware image, the hacker now has an entire blueprint of how the system operates. Any backend API calls within the system will easily be located within the code along with the access credentials required of them.

Once the calls and credentials are obtained, the hacker can now impersonate the device itself gaining access to the backend systems. Clearly the device has controlled access to the backend systems but at this point the hacker begins to look for vulnerabilities within the server software to dig further into the system.

Access at this level may compromise not only one customer's data, but the entire store of the customer base, effectively magnifying the data privacy and identity theft issues described above.

4.4. Device Disablement

ZigBee devices, which are quite prevalent in the low cost IoT space, have been shown to be quite vulnerable to worms and malware. The nature of their AdHoc network allows one ZigBee device, which is under a hacker's control, to affect other ZigBee devices connected to it.

While the interference with the actual functionality of a device frequently results in a bad user experience, many attacks have more nefarious goals. Recent ZigBee attacks have been demonstrated to block communications with home door locks, leaving the home unsecured. Clearly this can be leveraged by brick & mortar hackers to gain entry to the home. The use of these devices to control home lighting systems is a low risk implementation, however they must be used with care when designing safety critical solutions.

4.5. DDoS Attacks

The Distributed Denial of Service attack has been gaining increased visibility in the press of late. In this attack, a hacker targets a specific model of device that occurs in large numbers on the Internet. Malware is installed on all of these devices and they are used as a massive cluster to perform a DDoS attack on some web presence.

While this attack may not, and likely will not be on the Operator themselves, the PR from an attack of this nature is bound to be damaging to the corporate brand. Simultaneously, a DDoS is designed to send out continuous packet streams to take down a web presence. If enough devices are on the Operator's network, this could possibly cause network disruptions.

This is a significant threat in COAM devices due to the lack of transparency of the devices. While many of the security solutions discussed herein will prevent this type of device compromise, understanding which are implemented within devices is crucial to managing a secure ecosystem.

4.6. Ransomware

Ransomware has also been gaining press of late, but mainly in PC type environments. In this type of attack, a device's software or data is encrypted by a hacker, preventing access to the actual user. A ransom is charged by the hacker to provide a key for unlocking the device.

This type of attack is now gaining more prevalence within embedded devices but occurring at the ecosystem level. Once one device is hacked within the ecosystem, the hacker now has the ability to hack

all of them if they are identical devices. All devices are disabled with rogue firmware upgrades and the entire ecosystem is held hostage.

Attacks of this nature are interesting to hackers as an individual is no longer being extorted but instead a large organization is. This naturally leads to bigger ransoms and a perpetuation of the business model.

4.7. Theft of Service

Operators are in the business of selling digital media services. Theft of these services has been and continues to be a major undercut to the Operator revenue model.

Theft of service frequently focuses around the edge device as this is where the service is delivered and the hacker has access to it. Types of attacks in the category are quite broad but can consist of everything from password sharing to device cloning in the interest of gaining free access.

Again, a multi-layered security approach is key here to hardening the device and preventing the hacker access to it.

5. Conclusions

Herein we've looked at home security specifically from the device level, in particular management models, security features and finally threats to the home ecosystem. With respect to deployment of these systems, there are some key takeaways that one wants to consider during implementation.

5.1. Know Your Devices

In the traditional CPE model, the Operator has a high degree of control and understanding of the devices deployed within the ecosystem. With COAM devices this transparency is reduced, however not eliminated. It's important for the operator to understand these devices and assign levels of trust based on their capabilities. This categorization can vary from CPE, to devices with strong authentication down to unknown devices with a low level of trust. Once these devices are successfully categorized, permission can be granted based on their level of trust.

5.2. Understand Device Security Features

In order to classify devices into trust categories, it's pertinent to understand the security features implemented within a particular device. When the devices are managed by the Operator, this information is readily available however for other COAM devices this may require additional data sources. The OCF provides a device qualification leading to a strong device authentication certification. Certified devices of this type provide a level of trust that can be used for security profiling. Additionally, databases such as Shodan can be utilized for profiling unknown devices. Clearly, this takes more effort than a CPE based model however it is necessary to provide a secure yet open environment where all devices can play safely together.

5.3. Provide Access Based on Trust

Similar to human relationships, with devices we grant access based on the level of trust within the ecosystem. By categorizing these devices based on what we know about their security capabilities, we have the ability to grant access accordingly. Highly trusted devices may be granted access to Operator resources whereas devices with a lower trust level may only get network access and may even be quarantined based on rogue behaviour.

Abbreviations

CPE	Consumer Provisioned Equipment
COAM	Consumer Owned and Managed
MVP	Minimum Viable Product
MVPD	Multichannel Video Programming Distributor
OCF	Open Connectivity Foundation
TVE	TV Everywhere

Bibliography & References

Open Connectivity Foundation (OCF) - <https://openconnectivity.org/>

Shodan - <https://www.shodan.io/>

3 Embedded Hardware Security Features Your Smartphone Needs – Joel Snyder, <https://insights.samsung.com/2018/06/22/3-embedded-hardware-security-features-your-smartphone-needs/>

How hardware-based technology keeps mobile devices secure – Ben Cade, <https://gcn.com/Articles/2018/02/13/hardware-based-mobile-security.aspx>

The 7 Craziest IoT Device Hacks – Mike O'Malley, <https://blog.radware.com/security/2018/05/7-craziest-iot-device-hacks/>

5 Infamous IoT Hacks and Vulnerabilities – Isabel Harner, <https://www.iotforall.com/infamous-iot-hacks/>

Proactive Network Maintenance Evolution to the Optical Domain in Coherent Optics

A Technical Paper prepared for SCTE•ISBE by

L. Alberto Campos, Ph.D.

Fellow

CableLabs®

858 Coal Creek Circle, Louisville CO 80027

303 661 3377

a.campos@cablelabs.com

Zhensheng (Steve) Jia, Ph.D. CableLabs

Distinguished Technologist

CableLabs®

858 Coal Creek Circle, Louisville CO 80027

303 661 3364

s.jia@cablelabs.com

Larry Wolcott

Fellow

Comcast

1401 Wynkoop, Suite 300, Denver CO 80202

303 726 1596

Larry_Wolcott@cable.comcast.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
1. Optical Network Topology and Network Elements.....	4
1.1. Optical Network Topology	4
1.2. Optical Access Distribution Environment	5
1.3. Optical Network Elements.....	10
1.3.1. Optical Sources.....	11
2. Optical Signals, Transmission Environment and Impairments	14
2.1. Point-to-Point (PTP) and Point-to-Multipoint (PTM)Transmission Environment.....	14
2.2. Fiber Characteristics and Impairments	14
2.2.1. Attenuation	14
2.2.2. Chromatic Dispersion.....	14
2.2.3. Polarization Mode Dispersion	15
2.2.4. Nonlinear Effects.....	15
2.3. Optical Connectors and Splices	15
2.4. Optical Amplifiers	16
2.5. Multiplexers and Demultiplexers	17
2.6. Optical Splitters and Couplers.....	18
2.7. Isolators and Circulators	18
2.8. Optical Fiber Switches and Wavelength Switches	18
2.9. Coherent Receiver	19
2.9.1. Digital Coherent Receiver Types	19
2.9.2. Coherent Receiver Architecture.....	20
2.10. Impairments Impacting Coherent Systems	21
3. Optical Link Metrics and Link Characterization Tools	22
3.1. Optical Link Metrics	22
3.2. Optical Link Characterization Tools	23
4. Coherent Optical Transceiver Intelligence	25
4.1. Coherent transmission system.....	26
4.2. Coherent Optical Performance Monitoring.....	27
4.3. Basic Operation Principle.....	29
4.4. Comparison Between Direct Detection and Coherent Detection	31
4.5. Flexible resource allocation.....	32
5. Operational Strategy	33
Conclusion.....	38
Abbreviations	39
Bibliography & References.....	41

List of Figures

Title	Page Number
Figure 1 – Regional and Access Networks Connecting to Backbone	5
Figure 2 – Schematic Representation of Fiber Sheaths in Conduit.....	5
Figure 3 – Cable Access Network Topology Example.....	6
Figure 4 – Traditional HFC Fiber Node Topology	7
Figure 5 – N+0 Fiber Deep Network Topology	8
Figure 6 – Fiber Segment Sample Description with Wavelength Map	10
Figure 7 – Optical Signal Descriptors	10
Figure 8 – Optical Network Elements along Hub and Endpoint Transmission Path.....	11
Figure 9 – Laser structures, a-Fabry Perot (FP), b-Distributed Feedback (DFB), and c-External Cavity (ECL)	12
Figure 10 – Coherent Transmission using Amplitude, Phase, and Polarization.....	13
Figure 11 – Dual Polarization IQ Modulator.....	14
Figure 12 – Components of Typical Erbium-doped Fiber Amplifier	17
Figure 13 – Optical Wavelength Multiplexer and Demultiplexer	18
Figure 14 – NXN Optical Switch and a 1xN Wavelength Switch	19
Figure 15 – Three Coherent Detection Schemes: (a) Homodyne, (b) Intradyne, and (c) Heterodyne	20
Figure 16 – Phase and Polarization Diversity in Coherent Receiver Architecture.....	21
Figure 17 – Optical Spectrum Analysis	24
Figure 18 – Optical Time Domain Reflectometer (OTDR) Analysis.....	24
Figure 19 – Optical Modulation Analyzer (OMA) Metrics.....	25
Figure 20 – Coherent System with a Transmitter, a Transmission Fiber, and a Coherent Optical Receiver with Digital Signal Processing Flow Blocks	26
Figure 21 – Typical Parameters at Optical Layer.....	28
Figure 22 – Butterfly-structured Equalizer for Coherent Optical System	29
Figure 23 – Simplified Fiber Model including Major Transmission Elements	29
Figure 24 – Given versus Estimated CD Example	30
Figure 25 – Adaptive Coherent Transceiver to Support Different Scenarios.....	32
Figure 26 – Approaches for Flexible Data Rates and Software-defined Optics	33
Figure 27 – Average Time to Fusion Splice.....	35
Figure 28 – Example of PNM Software Stack.....	36

List of Tables

Title	Page Number
Table 1 – Electrical Impairments.....	21
Table 2 – Optical Impairments in Coherent Link.....	22
Table 3 – Optical Link Metrics.....	23
Table 4: Comparison of OPM Functions in Non-coherent and Coherent Systems	31

Introduction

Proactive network maintenance (PNM) in the HFC environment has taken advantage of the intelligence available in cable network elements such as the CMTS and CM, as well as plant information to determine type, severity, and location of the impairment. As fiber penetrates deeper in cable networks, the portion of transmission that takes places over coaxial cable is reduced and the resulting fiber networks become more elaborate. In these fiber networks, a fiber bundle branches into more paths to reach these deeper points. Next generation optical systems will have to be deployed in this new optical transport environment. This new optical transport environment will have a greater number of short optical segments that will likely be subjected to more handling as numerous optical drops to customers are installed. This requires enhanced troubleshooting tools as well as very granular data from the optical distribution plant in order to extract the valuable information needed to perform PNM troubleshooting.

Luckily, as the cable industry prepares to introduce coherent optics into its access networks, we have a transport mechanism that enables rich intelligence through the numerous processes that take place within the transceivers. These processes, combined with information gathered with other instruments, plant topology, and device configuration knowledge, can lead to detailed information regarding location, nature, severity, and duration of the problem.

Drawing on similarities from the coaxial PNM predecessor, maintaining high-order optical modulation profiles will require more scrutiny and maintenance than traditional analog and digital optical systems. When operating significantly higher data rates and service level agreements (SLA), many of the impairments that are commonplace within the optical domain will need to be maintained to a higher standard. This is especially important when considering some critical business services such as medical and mobile/cellular backhaul. These PNM capabilities will provide continuous reporting about the availability and quality of the optical links to support the SLA agreements of these services. Most importantly, operators can have a full awareness of problems before they impact the services, and provide an opportunity to proactively mitigate them.

1. Optical Network Topology and Network Elements

1.1. Optical Network Topology

When cable fiber networks were initially implemented in the mid 1990s, they used a tree and branch architecture, both in the fiber and the coaxial parts of the network. Figure 1 shows the fiber portion of the distribution network, extending from hub to fiber nodes, in addition to the interconnecting regional or metropolitan fiber networks typically in a ring configuration.

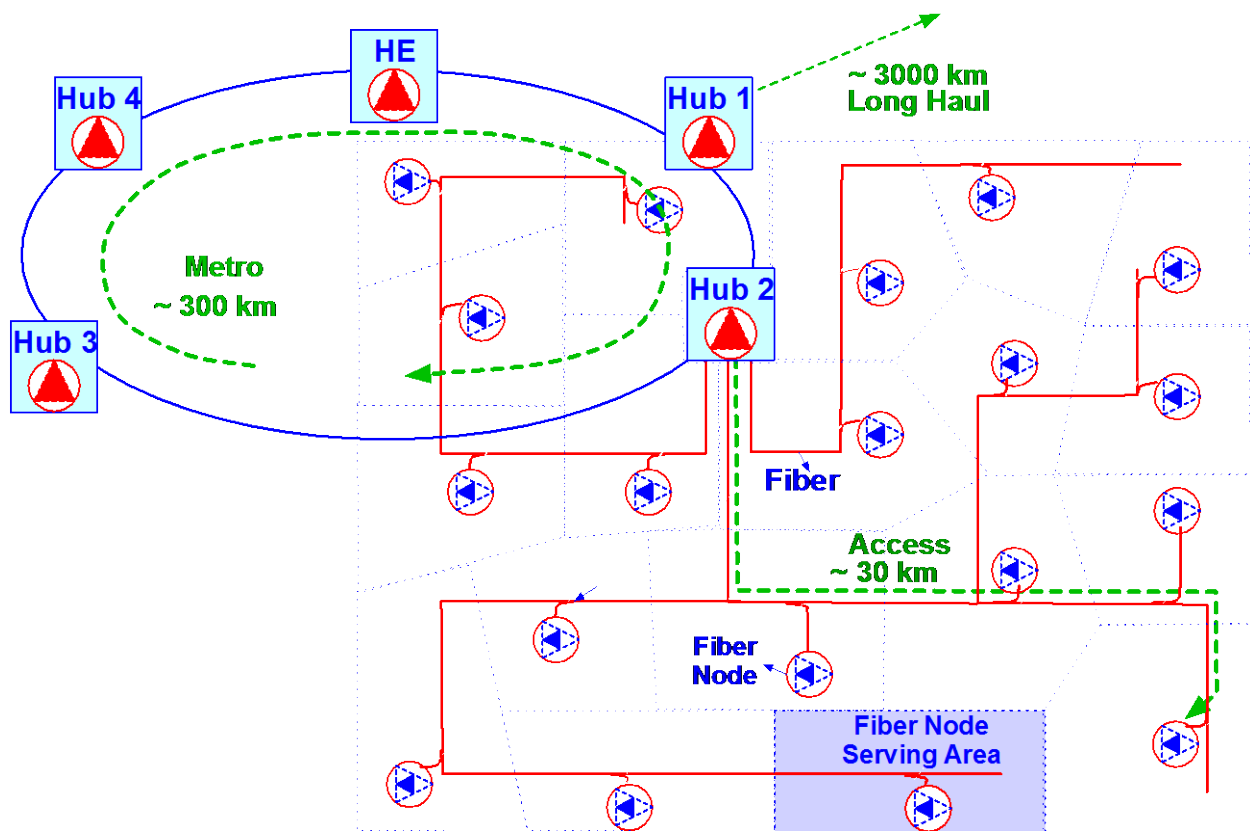


Figure 1 – Regional and Access Networks Connecting to Backbone

1.2. Optical Access Distribution Environment

Prior to embarking in a discussion of cable fiber topology, it is worthwhile to become familiar with the terminology related to how fiber strands are aggregated and carried through cable networks. Fiber strands are grouped in bundles or tubes, and these bundles or tubes are grouped in sheaths. In the case of underground infrastructure, operators deploy conduits through which the sheaths of fiber are blown. The bundles of fiber typically consist of either 12 or 24 fibers. Each fiber and bundle is color-coded to facilitate their management and manipulation (Figure 2).



Figure 2 – Schematic Representation of Fiber Sheaths in Conduit

The fiber access networks extend from the hub or headend to the fiber node. These fiber links are typically laid out by running fiber sheaths with fiber bundles that pass different nodes. From a fiber splice point near a fiber node, a fiber jumper cable with fewer fiber strands is trenched or strung to the node. In the initial HFC buildout, six to eight fiber strands were typically dedicated to a node. Most fiber distances from node to hub are less than 40 kilometers, although in a few areas (where hubs may have been consolidated) distances may reach 120 kilometers. Figure 3 shows in greater detail a representation of the fiber access network extending from hub to nodes.

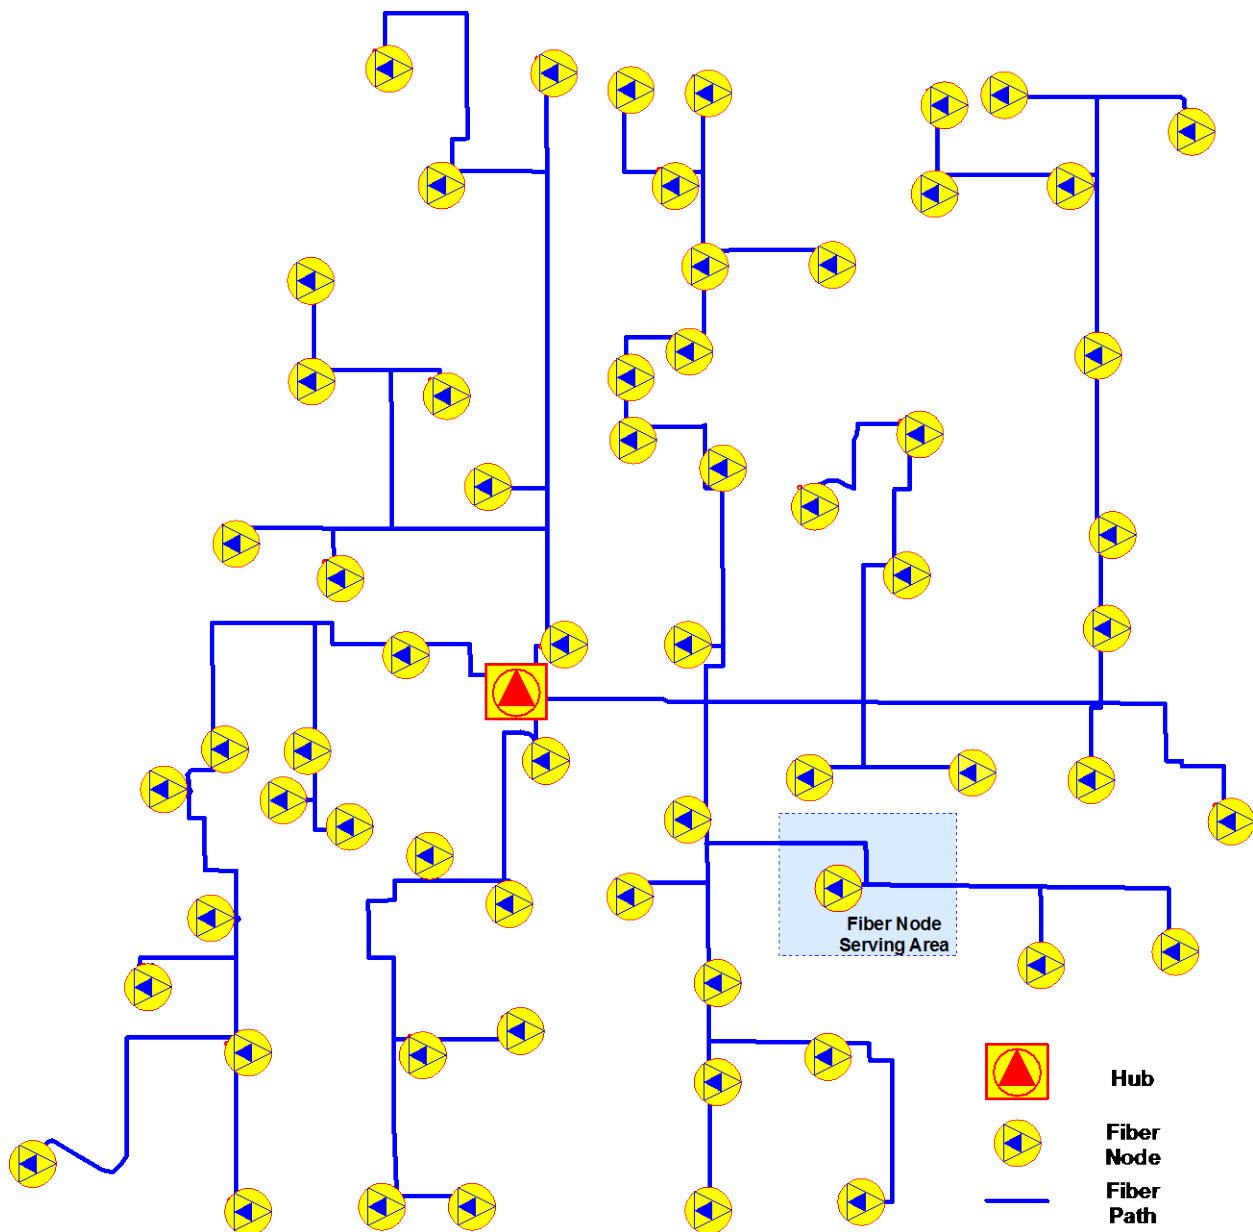


Figure 3 – Cable Access Network Topology Example

The shaded region in Figure 3 represents a traditional fiber node serving area. This fiber node serving area, shown in more detail in Figure 4, extends coaxial cable segments from the fiber node with a few amplifiers in cascade before reaching the subscriber. In addition to the amplifiers that maintain the RF signals at suitable levels, the coaxial cable segment uses taps (green squares in Figure 4), which couple RF energy to the drop cable that connects to the customer premises. It is important to note that while the fiber distribution network exhibits a tree and branch topology, from a connectivity perspective, the optical link between the hub and fiber node is a point-to-point link. It is important to be aware of, and to be able to determine, all the points in the fiber distribution network where the fiber paths bifurcate, as well as any changes in the number of fiber strands within consecutive sheaths of fiber along a transmission path. Transitions from one sheath to another at these points could potentially become problem areas in the future, requiring troubleshooting and maintenance.

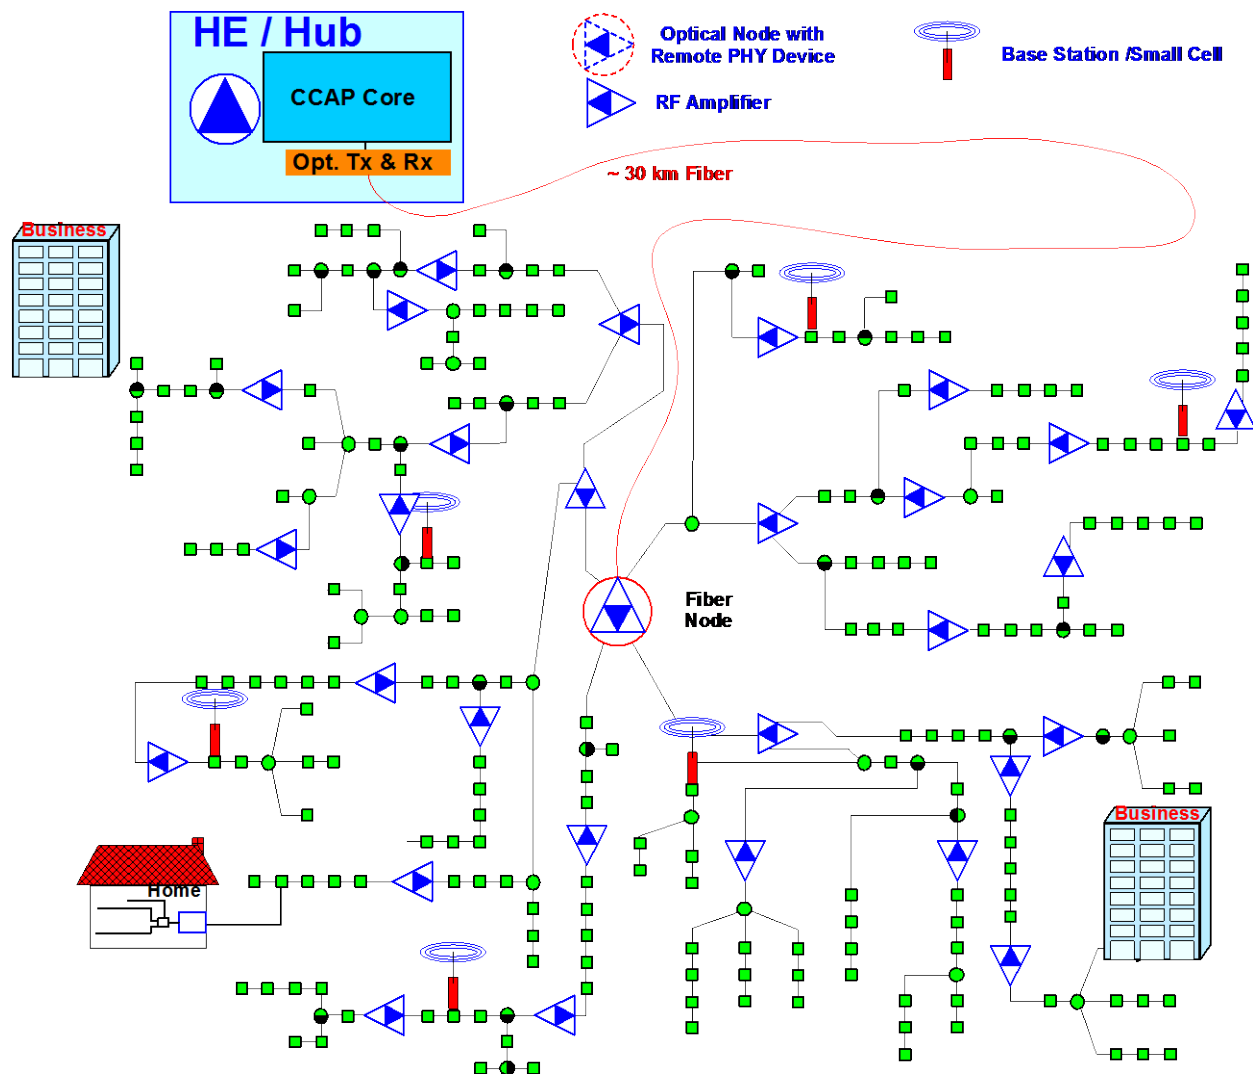


Figure 4 – Traditional HFC Fiber Node Topology

An architecture evolution approach to address the increasing demand in capacity is achieved by segmenting or splitting the fiber node serving area into smaller sections. The evolution of the same legacy node, shown in Figure 4, into an N+0 architecture, a fiber node followed by zero amplifiers, is shown in Figure 5.

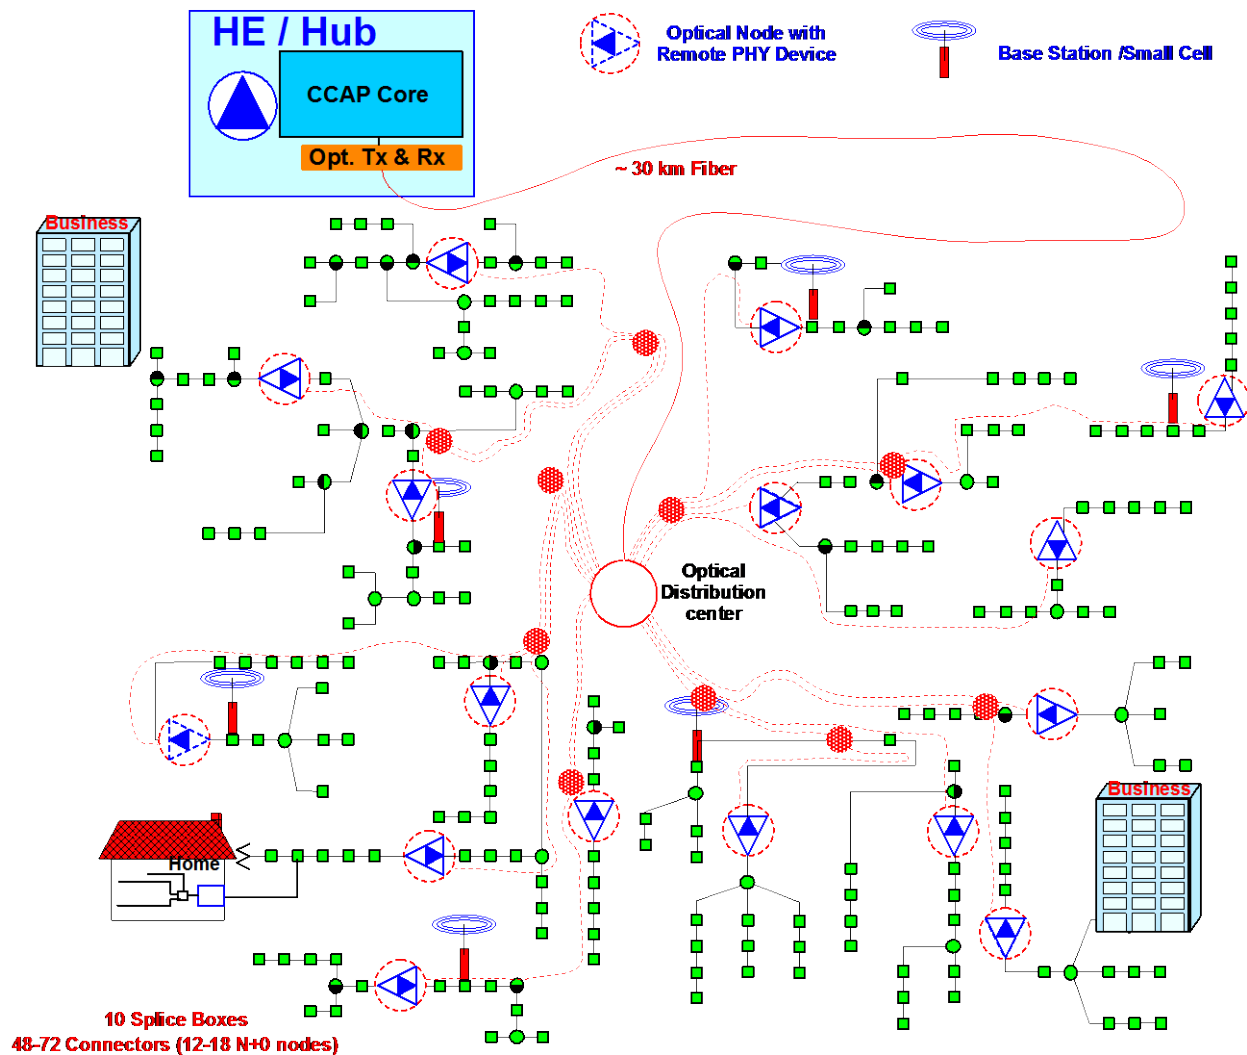


Figure 5 – N+0 Fiber Deep Network Topology

The upgrade of a traditional node into an N+0 architecture typically results in 12 to 18 deeper nodes. The motivation to push fiber deeper is twofold: To better serve residences with increased capacity, and to provide connectivity to businesses, cellular base stations, and wireless access points.

As depicted in Figure 5, the upgraded or “legacy” fiber node no longer serves as a location to transition to RF signals from an optical signal and vice versa, and instead becomes an optical distribution center, or ODC.

Deeper fiber penetration results in a more intricate distribution network that generally terminates within 1,000 feet of residences. The intricacy comes from relatively short fiber sheaths that subdivide into lower count fiber sheaths, and those subdivide again into even lower fiber count sheaths. In fiber-to-the-home (FTTH) or fiber to the end user scenario, it is from these lower fiber count sheaths that fiber drop cables are laid to connect to the endpoints.

While longer, uninterrupted fiber segments connect hub-to-node or hub-to-hub, as fiber reaches closer to homes, the fiber segments are much shorter. There is also an abundance of points where the fiber has been physically manipulated and spliced, therefore there are also many points with a higher likelihood of failure. As fiber penetrates deeper, the fiber topology begins to resemble the original cable topology, from the perspective of splits and segment lengths. That's one reason why it's vital to have granular location information about fiber plant and transceivers. With granular location information, it's easier to determine where problems exist, and their nature, so as to be better prepared in solving them.

A fiber distribution network generally runs a few dedicated fibers from the hub to the "legacy" fiber node, but an evolved network will likely have many more fiber strands from the "legacy" fiber node to deeper points in the network. At this legacy fiber node or optical distribution center (ODC), some optical signal manipulation will likely be conducted. The signal at the legacy node could be translated into the electrical domain by performing detection and retransmission, or it could remain in the optical domain and perform signal routing based on wavelength. Cost considerations and demand for capacity at the endpoints will determine what type of transition takes place at the ODC.

In the access environment, approximately 50% of the fiber is deployed underground and the rest is aerial. Aerial fiber is subject to wind movement, which can cause changes in state of polarization of a coherent signal. Other events that can generate changes in state of polarization (SOP) include lightning and arcing. The capability of adjusting and recovering after a sudden change in SOP is quantified in units of kiloradians per second. Indications about whether the transceiver is compensating or adjusting for changes in state of polarization can provide insight into environmental conditions.

Link length is an important parameter when assessing coherent link budgets and the need of amplification. In addition to accurately determining the link lengths through topology maps, tools exist within the transceiver that accurately estimate the length of the optical link. Several impairments that need to be compensated for are dependent on the optical link length, and accurate estimation of these impairments can facilitate link length estimation.

In a limited fiber strand environment and with a diversity of services carried over cable networks, coexistence of a variety of optical signal types over the same fiber is necessary. The signal types in cable fiber networks include analog optics, non-coherent digital optics, and coherent digital optics. Analog optical signals are transmitted at very high optical output powers, which can drive fiber into non-linear behavior and impact the transmission of other signals on the same fiber. It is important to have an accurate wavelength map of all the optical carriers that reside within each fiber segment, including transmit optical power, signal path traversed and launched location. Accurate optical carriers' characteristics and wavelength information allow us to estimate the impact of non-linear distortion.

An example of the channel map on a fiber strand within a bundle and a sheath is shown in Figure 6.

Signal Wavelength Map for Fiber34/Sheath ID-45									
		Bundle/Fiber		BR/BL		Length		14,000'	
Signal Type	Frequency	Wavelength	Sheath	Channel	Launch	Baud Rate/	Format	Signal Fiber Segments Association	
	Terahertz	nm	Signal Type	Width GHz	Power dBm	Bandwidth GHz	Modulation		
O-Band									
Analog US	229.0	1310				1	SCM	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 23/SheathID-15, Fbr2/ImprID-12	
IM-DD	222.2	1350				10	OOK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 17/SheathID-67, Fbr2/ImprID-4	
C-Band									
Coherent	193.1	1553.60	CO_31	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 22/SheathID-15, Fbr2/ImprID-11	
Analog DS	193.2	1552.80	AN_32	100	12	1.2	SCM	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 23/SheathID-15, Fbr2/ImprID-12	
Coherent	193.3	1551.99	CO_33	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 24/SheathID-15, Fbr2/ImprID-13	
Coherent	193.4	1551.19	CO_34	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 25/SheathID-15, Fbr2/ImprID-14	
Analog DS	193.5	1550.39	AN_35	100	12	1.2	SCM	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 26/SheathID-15, Fbr 27/SheathID-25, Fbr2/ImprID-21	
Coherent	193.6	1549.59	CO_36	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 27/SheathID-15, Fbr 27/SheathID-25, Fbr2/ImprID-22	
Coherent	193.7	1548.79	CO_37	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 28/SheathID-15, Fbr 27/SheathID-25, Fbr2/ImprID-23	
Coherent	193.8	1547.99	CO_38	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 29/SheathID-15, Fbr 27/SheathID-25, Fbr2/ImprID-24	
Analog DS	193.9	1547.19	AN_39	100	12	1.2	SCM	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 14/SheathID-67, Fbr2/ImprID-1	
Coherent	194	1546.39	CO_40	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 15/SheathID-67, Fbr2/ImprID-2	
Coherent	194.1	1545.60	CO_41	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 16/SheathID-67, Fbr2/ImprID-3	
IM-DD	194.2	1544.80	IM_42	100	0	10	OOK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 17/SheathID-67, Fbr2/ImprID-4	
Coherent	194.3	1544.00	CO_43	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 18/SheathID-67, Fbr2/ImprID-1	
Analog DS	194.4	1543.21	AN_44	100	10	1.2	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 19/SheathID-67, Fbr2/ImprID-2	
Coherent	194.5	1542.42	CO_45	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 20/SheathID-67, Fbr2/ImprID-3	
Coherent	194.6	1541.62	CO_46	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 21/SheathID-67, Fbr2/ImprID-4	
Coherent	194.7	1540.83	CO_47	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 22/SheathID-67, Fbr2/ImprID-5	
Coherent	194.8	1540.04	CO_48	100	0	30	DP-QPSK	Fbr 34/SheathID-45, 1x40MuxID-33, Fbr 23/SheathID-67, Fbr2/ImprID-6	

Figure 6 – Fiber Segment Sample Description with Wavelength Map

In order to fully identify the path of an optical signal in an environment that includes detailed topology information and component configuration, the information required is the signal location (i.e., hub), the fiber and sheath identification, and the signal's frequency or wavelength. If the wavelength and the components that the signal traverses are known, as well as how those components have been configured to manipulate or route wavelengths, then the path traversed by the signal can be determined. Figure 7 provides a signal identification example using a name convention that uniquely identifies the signal path, which is also included with other parameters.

Signal ID	Fiber34/SheathID45_193.3_HubA
Launched fiber	Fiber34/SheathID45
Frequency/Wavelength	F_193.3 Terahertz /W_1551.99 nm
Signal Type	CO/IM-DD/Analog
Baud rate/ bandwidth	CO
Launched power	30 GHz
Modulation /Bit rate	0 dBm
Signal path traversed	DP-QPSK
	Fbr 34/SheathsID-45, 1x40MuxID-33, Fbr 24/SheathID-15, Fbr2/ImprID-13

Figure 7 – Optical Signal Descriptors

1.3. Optical Network Elements

Along the optical connection paths and at the end points, there exists a number of optical network elements that aid in the transmission from hub to an optical transmission end point. Figure 8 depicts the most prevalent components in the optical access network.

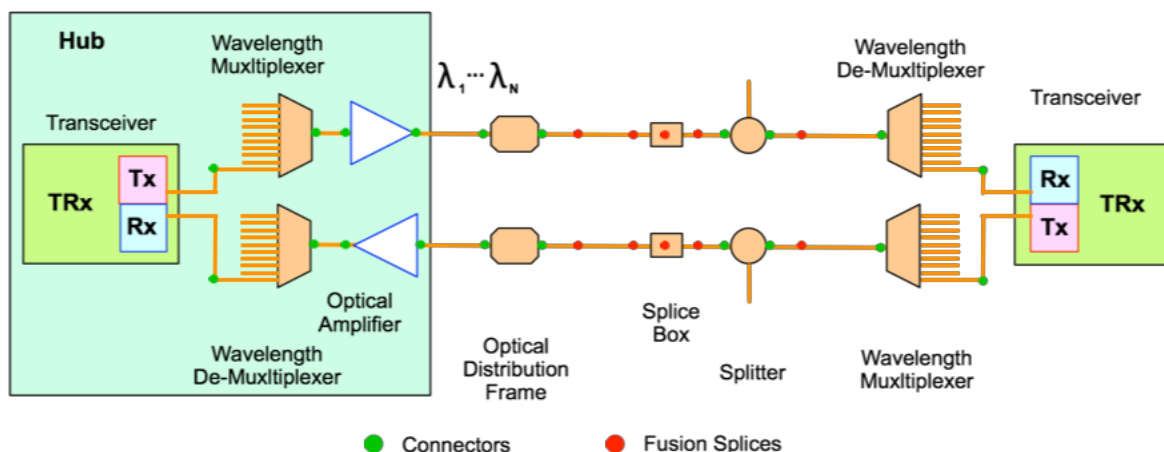


Figure 8 – Optical Network Elements along Hub and Endpoint Transmission Path

Depending on the type of optical transmission system, the transmitters and receivers may not be integrated in a transceiver as shown in Figure 8, but would be separate from each other, as is typically the case for analog optics.

Figure 8 does not include less frequently used components designed to facilitate monitoring or redundancy in the optical network. These include ROADMs, optical switches, wavelength switches, optical filters, attenuators, fiber drops, etc. In addition to the optical transport devices, the optical terminal devices are key elements of the optical distribution networks. Like their coaxial CM and CMTS equivalents, significant information can be extracted from these optical terminal devices regarding the health and characteristics of the optical network. Some key components of this optical distribution network are described next.

1.3.1. Optical Sources

The types of optical components used depend on the signal type in the optical link. The dominant optical links in cable access are the analog optical links. These are intensity-modulated links where the RF signal modulates the intensity of light to convey information through the RF modulated optical carrier. A second type of link used is a non-coherent digital optics link. They are also known as intensity modulation direct detection links (IM-DD), which are used in Gigabit Ethernet and GPON or EPON links. The third type of link is the coherent optical link. These are also digital optical links, but rely on a local oscillator as a reference at the receiver, and are able to distinguish phase and polarization information of the optical receive signal.

An analog optical link and an IM-DD link consist typically of a directly-modulated laser as the source, and a photodetector as the direct detection receiver. In longer links, the analog optics links may also be implemented using external modulation.

The coherent link transceivers are typically implemented using a complex external modulator called an IQ modulator. The key ingredient in a coherent link is the local oscillator laser at the receiver. At the receiver's photodetector, the incoming signal and the local oscillator signal beat together, generating a signal proportional to their product. This distinguishes the amplitude and phase information of incoming signals.

In analog optics, the linearity of the link is very important. When directly modulating the laser diode, the laser diode is operated at specific current and intensity levels such that it can take full advantage of its

linear region. For maximum signal to noise ratio (SNR), the amplitude swing of the modulating signal is as large as possible without exceeding the linear range.

The laser diode biasing could become misaligned as the laser ages, or with temperature shifts that are not compensated. This sub-optimal biasing of the laser diode introduces non-linear distortion and degrades the signal. Sometimes the signal modulating the lasers exceeds the normal amplitude swing. Impulse and burst noise can cause these unwanted laser current amplitude swings, which also introduce distortion. This is called laser clipping and is typically seen in upstream transmissions. Interleaving incorporated in transmission systems is designed to overcome this type of impairments.

In order to achieve the high signal to noise ratio that is required to support DOCSIS 3.1 higher order modulations (16384-QAM and 4096-QAM), a very high optical transmit power is typically required. This high optical power level could drive fiber into nonlinear behavior and introduce distortion, particularly when multiple optical carriers are used.

Most of these problems are associated with an increase in codeword errors and a decrease in SNR and MER. Some of these distortion metrics can be assessed through optical and electrical spectrum analysis that show the generation of non-linear components. Laser clipping, for example, is detected when energy is generated above the maximum upstream frequency.

Diode lasers are resonant cavities that, depending on their design, can generate one or multiple modes of light. A simpler laser cavity like the Fabry-Perot laser has multiple resonant modes across the gain spectral region of the laser. Distributed feedback lasers significantly inhibit the generation of multiple modes through a periodic internal structure that is tuned to a specific resonant frequency. Other structures, like external cavity lasers, are more restrictive in the energy they generate. In later sections, the performance degradation attributable to the laser emission characteristics is discussed. Figure 9 provides a schematic representation of laser structures used in cable.

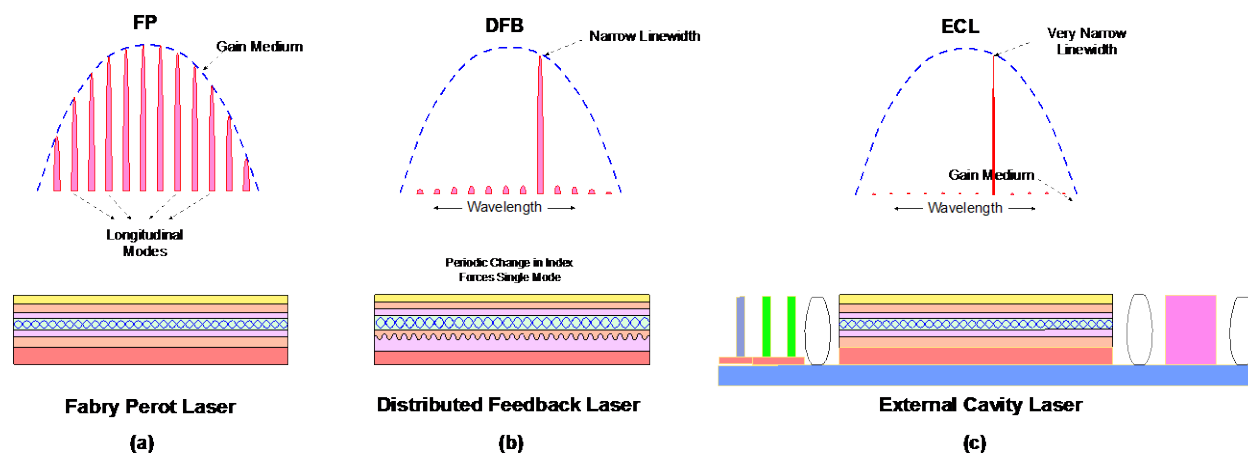


Figure 9 – Laser structures, a-Fabry Perot (FP), b-Distributed Feedback (DFB), and c-External Cavity (ECL)

The high dynamic range required in cable's multichannel environments prompted the use of Distributed Feedback (DFB) lasers in analog optics. DFB lasers are also used in non-coherent digital optics implementations, although they don't demand as high a dynamic range as do analog optics. Early in cable, the upstream signal path required only the transmission of a few channels, at low modulation orders. That allowed the use of the simpler Fabry-Perot (FP) laser, although in many instances FPs suffered from laser clipping, and as a result, the industry migrated away from them.

In certain cable upstream implementations, the RF spectrum was digitized and transmitted on IM-DD links, to overcome the distance limitations in analog links.

As part of the management information of the optical link types -- in addition to specifying whether the signals are analog, IM-DD or coherent -- it is worthwhile to go to a deeper description level and indicate in analog optics if the signals are externally modulated and in IM-DD if the signal is digitized RF.

Intensity Modulation Direct Detection (IM-DD)

Non-coherent digital optical links or intensity modulated direct detect (IM-DD) links typically use On-Off-Keying (OOK), that turns light intensity On and Off to encode information. In this type of modulation, linearity considerations are not critical. The receive sensitivity of the link and the transmit power level have to be such that after all the system losses, the signal to noise ratio is still appropriate for OOK transmission. Non-coherent transmission typically operates at lower power levels than analog optics, so fiber non-linearity is not an issue. In the future, non-coherent digital optical transmission could include multi-level signal transport using, for example, pulse amplitude modulation (PAM). PAM modulation, such as four level PAM-4, is sensitive to component linearity.

Coherent Links

External modulation is typically used in coherent links. The spectral purity of the laser signal enables the encoding and detection of both amplitude and phase information of the optical carrier. This spectral purity is quantified by measuring the laser linewidth. This is the 3 dB spectral width occupied by the unmodulated optical carrier. In coherent optical links the polarization can also be discriminated. All this allows for the encoding of information on an optical carrier in three dimensions; amplitude, phase, and polarization.




Figure 10 – Coherent Transmission using Amplitude, Phase, and Polarization

The modulator to encode information in these three dimensions onto light is the IQ modulator, shown in Figure 11.




Figure 11 – Dual Polarization IQ Modulator

2. Optical Signals, Transmission Environment and Impairments

2.1. Point-to-Point (PTP) and Point-to-Multipoint (PTM) Transmission Environment

In each fiber strand there can be a multitude of optical signals that operating on different wavelengths share a portion or the whole optical path with other signals. Understanding the characteristics of this fiber access environment and its signals is critical in determining the performance of the optical links that reside and coexist within the fiber strand. There are different fiber-related impairments that can impact performance. Some of these impairments are fiber-length dependent, and some are dependent on fiber geometry, material, wavelength, bandwidth, and optical power level. In cable, even though the fiber cable paths follow a tree and branch topology, the actual fiber connectivity is point-to-point and no optical splitting or coupling takes place. There are, however, passive optical networks in residential green field scenarios, where RFoG technology is deployed, and PON networks predominantly deployed in business access. In both cases, these P-to-MP fiber networks are implemented using a 32- or a 64-way split.

2.2. Fiber Characteristics and Impairments

2.2.1. Attenuation

Attenuation in fiber is dependent on the wavelength or frequency. For the particular type of single-mode fiber typically used in cable access, the attenuation is 0.22 dB/km for 1550 nm transmission and 0.3 dB/km for 1310 nm transmission.

2.2.2. Chromatic Dispersion

Dispersion is one impairment associated with fiber length. Dispersion occurs when different portions of the signal travel at different speeds. As a consequence, there is a spreading of the signal over time. There are different types of dispersion: Chromatic, waveguide, modal, and polarization mode dispersion. Chromatic or material dispersion is caused by the change of refractive index with optical frequency. Waveguide dispersion relates to how well the index of refraction represents an ideal waveguide

throughout the fiber length. The differences from an ideal waveguide cause dispersion. Modal dispersion occurs when different propagating modes are present in fiber. In the cable access environment, the predominantly deployed fiber is single mode fiber (SMF), so fiber modal dispersion is not present and waveguide dispersion is negligible compared to chromatic dispersion.

2.2.3. Polarization Mode Dispersion

Polarization mode dispersion (PMD) occurs when two orthogonal polarizations travel at different speeds, which causes pulse spreading. This is caused by random glass imperfections, such as circular asymmetry. PMD is compensated in the DSP of coherent receivers. Since PMD is distance-dependent, compensation for it provides another indicator for link distance estimation. Non-coherent receivers typically have no PMD compensation mechanisms and have to deal with PMD as part of their error correction techniques. This limits performance and link distance. PMD is not an issue in analog optical links, as the modulation bandwidth is about 1 GHz.

2.2.4. Nonlinear Effects

Nonlinear effects in fiber are attributable to intensity dependence of the refractive index fiber medium, and due to inelastic-scattering present at very high optical intensity levels.

These non-linear effects include self-phase modulation (SPM), cross-phase modulation (XPM) and four-wave mixing (FWM). The dominant non-linear effect in fiber is four-wave mixing. In four-wave mixing, if three fields are propagating at frequencies ω_1 , ω_2 and ω_3 , a fourth frequency ω_4 is generated such that $\omega_4 = \omega_1 \pm \omega_2 \pm \omega_3$. FWM is independent of modulation bandwidth and is instead dependent on frequency spacing and fiber dispersion. This effect is critical when multiple high power optical analog carriers are present on the same fiber. Maintaining a detailed fiber wavelength occupancy map --, including optical carrier types such as analog, IM-DD or coherent, optical carrier bandwidth, center frequency, modulation and optical power -- is important to estimate aggregate optical power and potential impacts to any carriers.

2.3. Optical Connectors and Splices

Optical connector and splices are incorporated through cable's fiber network to enable connectivity to the desired endpoints. Connectors are used at endpoint or at mid-point where reconfiguration is not expected. In cable analog optics, the stringent requirements to avoid optical reflection prompted the use of angle-polished or angle-faceted connectors. These connectors have a mating surface that is not perpendicular to the axis of the fiber, but is instead at an angle, so that any reflection leaves the core and vanishes. Other potential sources for reflections are suboptimal fusion splices. Even though these may rarely occur because of the large number of splices in the network, it is possible to find a few bad splices. Connector problems or splice problems can be detected using an Optical Time Domain Reflectometer (OTDR). Alternatively leveraging the coherent transceiver DSP capabilities, optical reflections, and other distortions can be detected, measured, compensated, and located in a similar fashion that equalization coefficient analysis is used to determine reflections in cable's coaxial environment. In the field, fiber patch panels and splices are housed in enclosures or splice boxes in cabinets or pedestals. Again, detailed records of the fiber-to-fiber mapping are important to troubleshoot and to follow the fiber connectivity paths throughout the cable network. These records should point to the wavelength occupancy map described above.

Many times there are sheaths with dissimilar fiber counts entering a splice box or cabinet. As an example, you may have a 312-fiber sheath, a 216-fiber sheath and a 48-fiber sheath in a splice box. This mismatch in number of fiber strands ($312-216-48 \neq 0$) leaves fiber strands available for future use. Good accounting

and management of these available fibers and reclaimed fibers through the effective use of wavelength multiplexing provides the operators with significant long term fiber resources.

2.4. Optical Amplifiers

In cable, a majority of optical links are short enough to not require optical amplification. Nonetheless, a number of scenarios can exist, where optical amplification is necessary. In particular, as cable is moving to a WDM environment, in order to make more efficient use of fiber's wavelength spectrum, the additional loss of multiplexing, demultiplexing and other wavelength manipulation functions may require the introduction of optical amplifiers.

Some optical equipment incorporates an optical amplifier, but in other cases optical amplification is needed to overcome loss in the fiber path and to compensate for splitting, coupling losses, losses in wavelength multiplexing and other losses. An optical amplifier located immediately following the transmitter is called booster amplifier; one located in the optical distribution network is called an in-line amplifier; and one located just before the receiver is called pre-amplifier.

The most common optical amplifier used in optical transport networks is the Erbium-doped fiber amplifier (EDFA), while the semiconductor optical amplifier (SOA) and the Raman amplifiers are less common.

The EDFA functions by using two optical signals. One is the optical carrier to be amplified, typically carrying information, and the second one, which operates at a different wavelength, is used to excite the Erbium atoms to a high-energy state. When these excited atoms return to a lower energy state, they release photons at the same wavelength, phase and direction as the information-carrying signal, therefore amplifying the signal. The optical signal used to excite the atoms is called the pump signal.

The gain in Erbium-doped fiber (EDF) is not flat, so in many cases amplifiers include an optical filter that equalizes the signal by flattening the frequency response. Not all the photons generated are stimulated by the signal to be amplified -- some photons are spontaneously emitted and they contribute to noise. This noise associated with the optical amplification process is called Amplified Spontaneous Emission noise or ASE. Depending on the implementation, the noise figure may range from 4 dB to 7 dB.

EDFAs operate mostly in the C-band (1525 nm -1565 nm) but they can also be designed to operate in the L-band (1570nm-1610nm). Erbium can be excited using 980 nm and 1490 nm pump wavelengths. In low noise applications, 980 nm pump wavelengths are used. Non-linearities are also present in EDFAs, mostly in the form of gain saturation. The EDFA's capabilities, configuration, wavelength occupancy and aggregate power within the fiber have to be taken into account for proper operation. A basic design of an Erbium-doped fiber amplifier is shown in Figure 12.

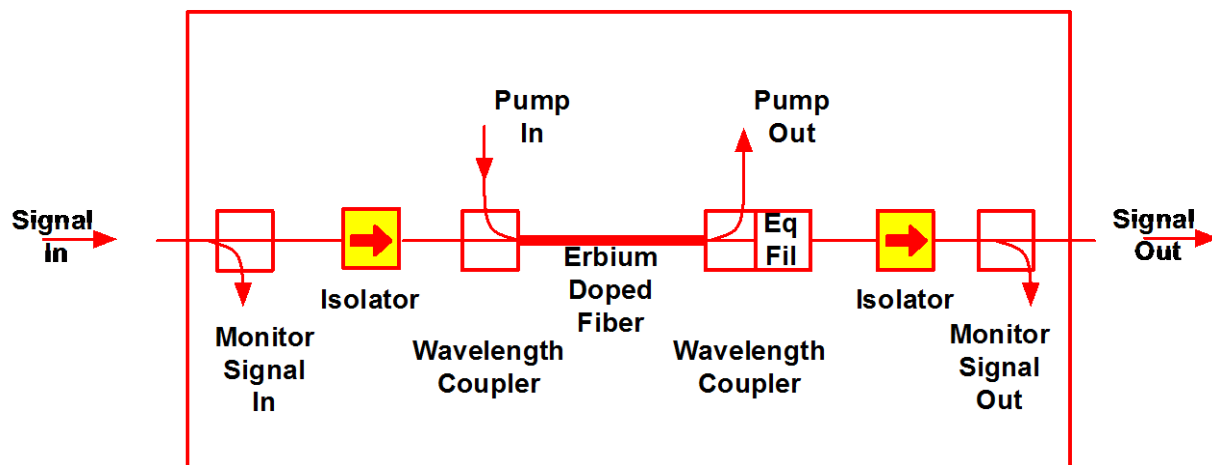


Figure 12 – Components of Typical Erbium-doped Fiber Amplifier

Raman optical amplifiers use stimulated Raman scattering, where light is scattered from a lower wavelength to a higher wavelength. Raman optical amplifiers are not practical by themselves, because they need extremely high optical pump powers (~30 dBm). They can, however, be used with EDFAs to implement ultralow noise amplification. An advantage is that the amplified wavelength is related to the pump wavelength, which provides more flexibility to amplify in different optical bands where other amplification methods may not be practical.

Semiconductor optical amplifiers (SOAs) work in a similar fashion as EDFAs except that, instead of an optical pump that brings the atoms into an excited state, it uses an electrical pump. With it, the electrons are brought to an excited state through the biasing of a semiconductor junction. These excited electrons release photons that are stimulated by the optical signal carrying information. SOAs have a similar structure as a laser diode except that SOAs do not have reflecting facets at both sides of the cavity. SOAs typically have medium gain (<20 dB) and low saturated optical power (<10 dBm). Since SOAs share a similar structure and operation as semiconductor lasers, SOAs are typically incorporated within transmitters or receivers.

2.5. Multiplexers and Demultiplexers

In the topology discussion of Section 1.1, we see that cable uses fiber that penetrates very deep, although there are not many fiber strands available at the endpoints. This prompts the industry to use fiber's wavelength spectrum very efficiently, so that a single fiber can carry a multitude of optical signals. This is accomplished by multiplexing optical carriers on the same fiber. In cable we deal with a diversity of signals. Analog optical signals, non-coherent, or intensity modulated-direct detect (IM-DD) signals and coherent signals -- must all coexist within the same fiber. A separate paper [1] assessing the coexistence of different optical signals in cable is also presented. A key component in the multi-optical carrier future is the wavelength multiplexer and demultiplexer.

A wavelength multiplexer filters specific wavelengths and routes and aggregates them to specific output ports. A wavelength demultiplexer distributes different wavelengths into different ports. Wavelength multiplexers and demultiplexers are bidirectional. Depending on how a device is configured and driven, it may function as a multiplexer, as a demultiplexer, or both.




Figure 13 – Optical Wavelength Multiplexer and Demultiplexer

2.6. Optical Splitters and Couplers

In basic point-to-point optical links there is no need to split the optical signal. Where optical splitting becomes advantageous is when networks are point-to-multipoint. Cable systems tend to leverage PTM networks, PON networks, and RFoG networks. In the cases of PTM and PON, these P-to-MP fiber networks are implemented using a 32- or a 64-way split. In the case of RFoG, the signal that is typically shared among the 32 subscribers is the same RF signal over an optical carrier that an optical node would receive. Some of these 32-way split optical networks may be combined at the hub in order to have a suitable DOCSIS® and video serving group size.

Digital E-PON and G-PON networks are also used by operators, where fiber is again split in 32- or 64-ways, to provide connectivity to business or residential subscribers.

2.7. Isolators and Circulators

Isolators leverage polarization to allow only one direction of transmission, while circulators leverage polarization to force a signal to traverse the three-port circulator, following input-to-output rules where the signal's direction in the circulator depends on the direction of entry.

2.8. Optical Fiber Switches and Wavelength Switches

While optical switches are used today primarily for redundancy applications and automatic configuration, a greater need is expected for conducting wavelength manipulation in the access environment. As mentioned before, the condition of deep but sparse fiber penetration prompts the industry to look into very efficient usage of their wavelength spectrum. The way networks are evolving in cable is happening coincidentally with the advent of distributed access architectures (DAA). In the evolved fiber distribution network, there is a transition -- with dissimilar fiber counts coming into the legacy node or optical distribution center, and going out of the node, to deeper points in the network. In order to retain the flexibility in that portion of the network, wavelength manipulation may be needed. This could be static manipulation, through wavelength multiplexers or demultiplexers, or it could be flexible and agile, through wavelength switches. This wavelength manipulation functionality is used today in fiber backbones through reconfigurable optical add-drop multiplexers (ROADMs). In the access plant, conventional ROADMs provide much more functionality and capacity than what would be required at the node -- but in future and simplified implementations, with a subset of the functionality, they could play a role in the industry's access networks.

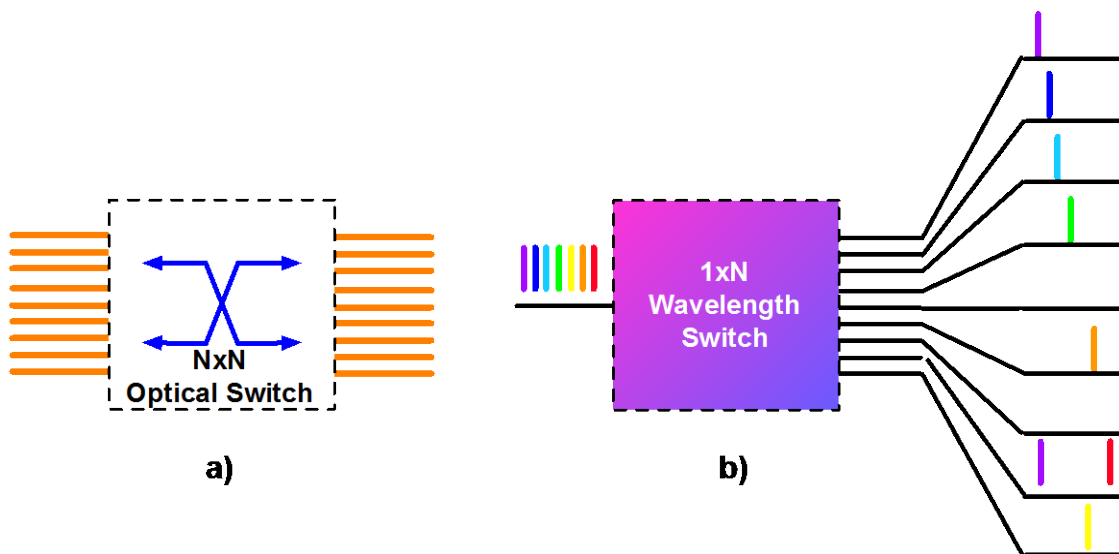


Figure 14 – NXN Optical Switch and a 1xN Wavelength Switch

2.9. Coherent Receiver

2.9.1. Digital Coherent Receiver Types

In a coherent receiver, a local oscillator (LO) is used to down-convert the electrical field of the incoming optical signal to a baseband intermediate frequency (f_{IF}). This coherent detection maps an entire optical field into the digital domain, therefore allowing the detection of the signal's amplitude, phase, and state of polarization. Depending on the intermediate frequency, defined as $f_{IF} = f_s - f_{LO}$, coherent receivers fall into three classes: Homodyne, intradyne and heterodyne, as illustrated in Figure 15, where $Bandwidth_s$ is the optical signal bandwidth.

Intradyne receivers are the de facto choice for contemporary 100G coherent systems. In an intradyne receiver, the f_{IF} is chosen to fall within the signal band by roughly aligning the f_{LO} with f_s . Intradyne detection allows the detection of both the in-phase and quadrature component of the received signal. For that reason, the intradyne receiver is also referred to as a “phase-diversity receiver.” Digital phase locking algorithms are needed to recover the modulation signal from its sampled I and Q components; this requires high-speed analog-to-digital conversion and DSPs.

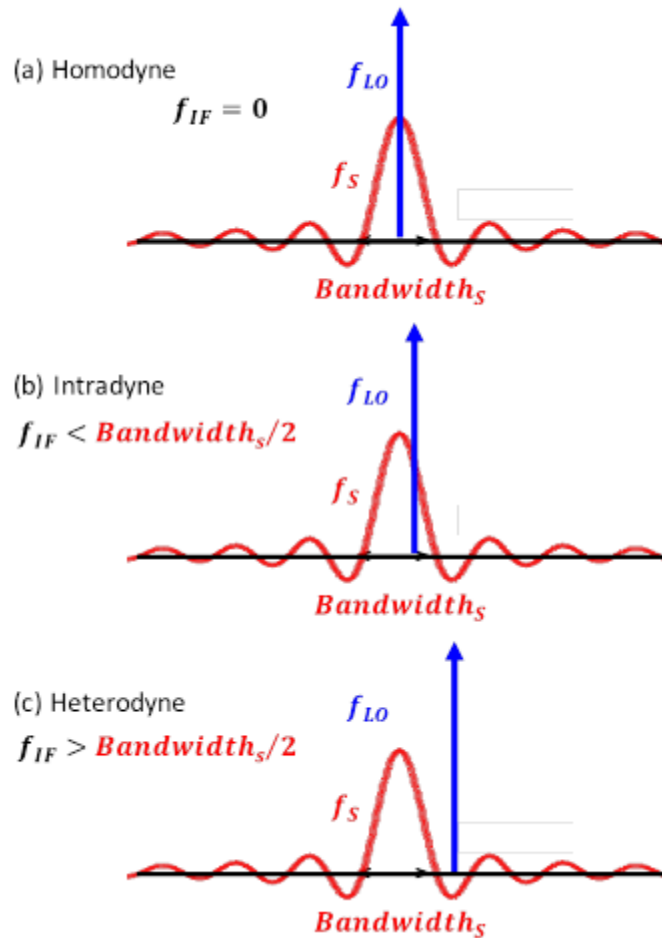


Figure 15 – Three Coherent Detection Schemes: (a) Homodyne, (b) Intradyne, and (c) Heterodyne

2.9.2. Coherent Receiver Architecture

In a coherent receiver, the modulated optical signal and a continuous wave LO beat together in the photo detector, generating a component proportional to the product of their electric fields which can be processed electrically. To detect both IQ components of the signal light, a 90° optical hybrid is utilized to provide a 90° phase shift between its direct-pass and cross-coupling outputs, which is used to discriminate between real and imaginary components of the optical signal. This is done for both polarizations. Balanced detection is usually introduced into the coherent receiver as a means to suppress the DC component and maximize the signal photocurrent.

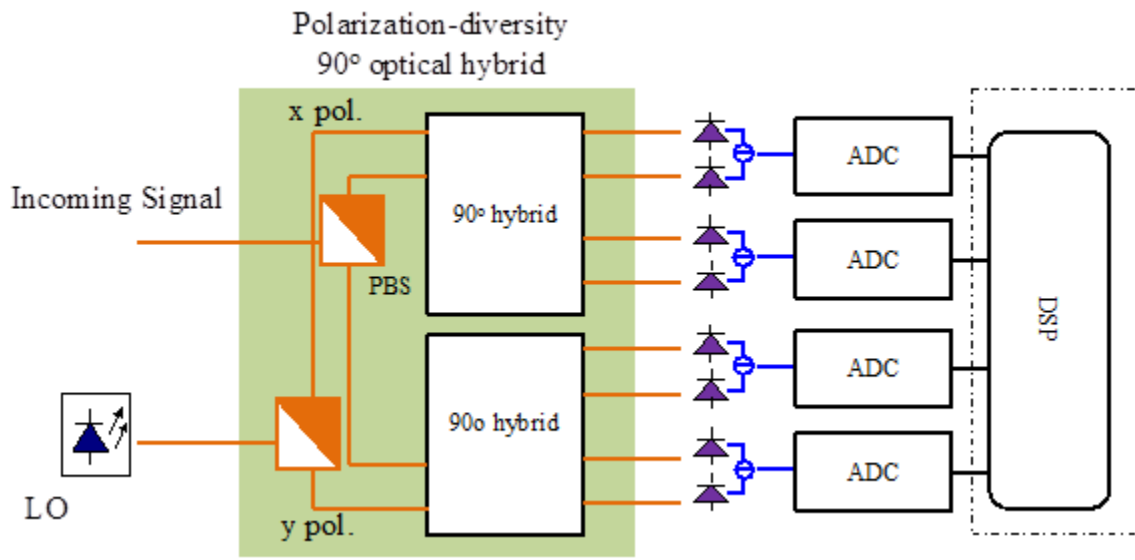


Figure 16 – Phase and Polarization Diversity in Coherent Receiver Architecture

The schematic diagram of a polarization multiplexed coherent receiver is shown in Figure 16. Both the incoming PM signal and LO are split into two orthogonal polarizations using a polarization beam splitter (PBS), after which the copolarized signal and the local oscillator are mixed in two 90° optical hybrids to produce an in-phase and quadrature component for each polarization. The four signals are then digitized by four analog-to-digital converters (ADC), after which DSP can be performed for signal demodulation.

2.10. Impairments Impacting Coherent Systems

There are impairments that impact the signal while it is in the optical domain, and there are impairments that are generated and impact the signal while it is in the electrical domain. Most of the electrical domain problems are related to the implementation of the system's transmitter and the receiver of the system, while many of the optical domain impairments are dependent on the fiber and related optical components along the optical connectivity path. Since the electrical impairments provide insight about the design but not about the plant, this paper places emphasis in the assessment and compensation of the optical impairments. Electrical and optical impairments are shown in Table 1 and in Table 2.

Table 1 – Electrical Impairments

Electrical frequency response
Impedance mismatches
Polarization imbalance and skew
In-phase and Quadrature (IQ) imbalance and skew
Transimpedance amplifier (TIA) and noise

Table 2 – Optical Impairments in Coherent Link

Linear distortion	Chromatic dispersion
	Polarization mode dispersion
	Optical reflections/multipath
	Group delay distortion
	Optical back reflections
	Filter narrowing effect
	Optical components frequency response
Non-linear distortion	Fiber non-linearities (i.e. four-wave mixing)
	Optical amplifier non-linearities
	Modulator non-linearities
Loss	Attenuation
	Insertion loss
	Polarization dependent loss
	Thermal noise
Noise	Shot noise
	Relative intensity noise
	Amplified spontaneous emission noise
	Optical back scattering

3. Optical Link Metrics and Link Characterization Tools

3.1. Optical Link Metrics

Different optical signal types have some common, as well as some unique, metrics that help to assess the quality and health of an optical link. Since many impairments do not change with wavelength, coherent links can be used to measure the health of IM-DD and analog optical links. The optical path may not be common from end to end between an analog, a non-coherent and a coherent optical signal, so when troubleshooting non-coherent and analog links, multiple coherent links may be needed to evaluate the entire non-coherent signal path. A list of useful optical link metrics is included in Table 3.

Table 3 – Optical Link Metrics

Optical Link Budget
Optical Signal to Noise Ratio (OSNR)/ Error Vector Magnitude (EVM)
Bit Errors/Symbol Errors
Pre-FEC BER
Post-FEC BER /Codeword Errors
State of Polarization Variation
Noise Figure
Linear Distortion - Reflections
Linear Distortion - Chromatic and Polarization Mode Dispersion
Link Length
Return Loss
Insertion Loss
Gain Compression (Including non-linear distortion)

3.2. Optical Link Characterization Tools

A diverse set of instruments exists that can be used in the analysis and troubleshooting of optical networks. Some popular instruments include:

Optical Power Meter

Optical Spectrum Analyzer (OSA)

Optical Time Domain Reflectometer (OTDR)

Optical Modulation Analyzer (OMA)

Optical Vector Analyzer (OVA)

The optical power meter consists of a calibrated photodetector, which, by knowing the responsivity of the photodiode versus wavelength, can accurately estimate the optical power level. Because the responsivity doesn't change drastically with frequency, it is sufficient to know the band of the optical signal in order to have a good estimate of the power level. The optical spectrum analyzer, on the other hand provides a detailed behavior of power versus wavelength or frequency (Figure 17).




Figure 17 – Optical Spectrum Analysis

In a similar fashion as the electrical time domain reflectometer, the optical time domain reflectometer (OTDR), operates by sending a narrow pulse, and in the same channel, detects the reflected pulse energy. This allows for the identification of optical transmission discontinuities as well as loss assessments along the optical transmission path (Figure 18).




Figure 18 – Optical Time Domain Reflectometer (OTDR) Analysis

The optical modulation analyzer (OMA) is a very powerful tool that is typically used more in the laboratory than in the field, but it is worth describing because of the informational richness it provides about link performance. An OMA enables measurement of the constellation quality of each polarization,

as well as error vector magnitude (EVM) (which is equivalent to modulation error ratio [MER].) Constellation analysis can also provide IQ and polarization amplitude imbalance and skew. OMAs can independently characterize each information lane (XI, XQ, YI, YQ) by measuring the eye diagrams, in addition to providing BER and channel frequency response. The distortion compensation that the OMA applies to correct the channel provides great insights into channel characteristics (Figure 19).

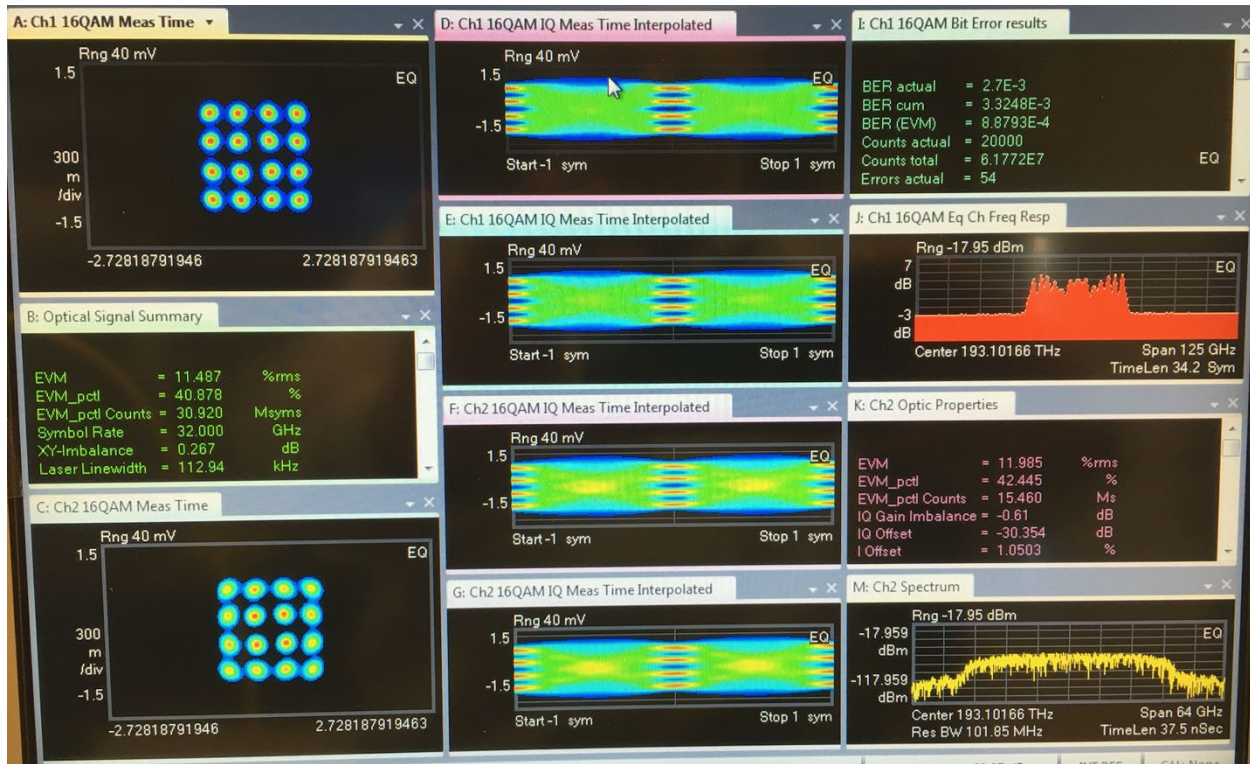


Figure 19 – Optical Modulation Analyzer (OMA) Metrics

The least common instrument is the optical vector analyzer (OVA), which provides equivalent measurements as an electrical vector network analyzer -- insertion loss, return loss as well as system transfer functions.

Some of these instruments are quite complex and expensive. However, in coherent optical networks, there is an opportunity to leverage information generated by the coherent optical transceiver while compensating the different impairments and conditions in the optical link. These parameters can provide useful distortion, reflection, and loss information.

4. Coherent Optical Transceiver Intelligence

Coherent Optics System – The optical coherent transmission system undergoes a variety of processes before it sends its signal to the optical transmission medium.

4.1. Coherent transmission system

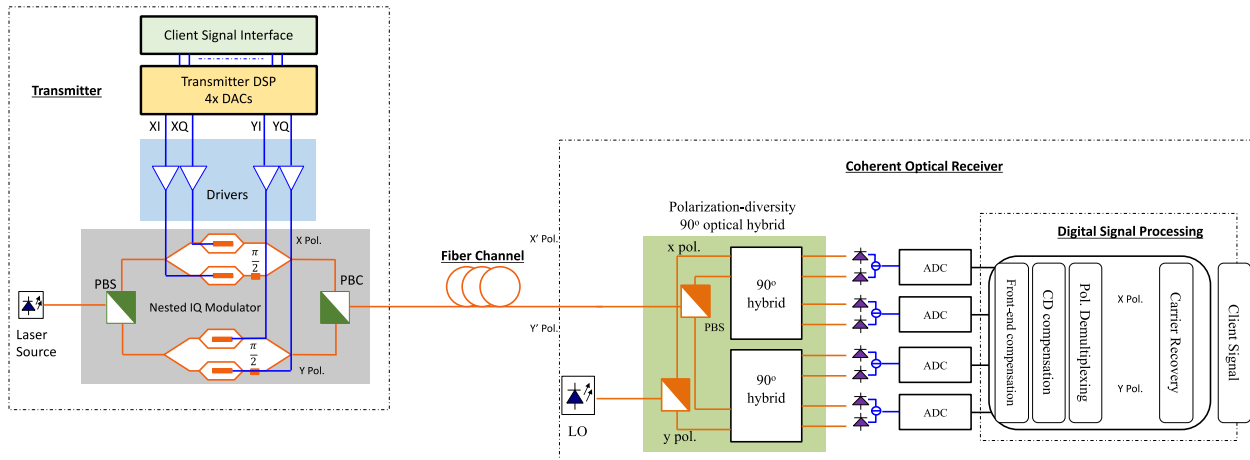


Figure 20 – Coherent System with a Transmitter, a Transmission Fiber, and a Coherent Optical Receiver with Digital Signal Processing Flow Blocks

As an example, Figure 20 presents a coherent optical system with a transmitter, a transmission fiber, and a coherent receiver [2] [3]. In the transmitter there is an optical nested IQ modulator with dual Mach-Zehnder modulators (MZM) for QPSK or higher modulation QAM formats. It can be composed of a phase modulator and two MZMs, and is commercially available in an integrated form. The incoming light is equally split into two arms: The in-phase (I) and the quadrature (Q) arm. In both paths, a field modulation is performed by operating the MZMs in a push-pull mode, at the minimum transmission point. Moreover, a relative phase shift of $\pi/2$ is adjusted in one arm, for instance by an additional phase modulator. This way, any constellation point can be reached in the complex IQ-plane after recombining the light of both branches. One of the most important parameters of the coherent QAM signal modulation in the transmitter is the modulation loss, which depends on the following factors:

- Insertion loss and bias points of modulator
- Driver swing and driver rise/fall times
- Modulation format
- Linearity of modulator
- Spectral shaping and pre-compensation

Operational optimization is needed to maximize the output power of the modulation while balancing linearity.

For generating dual-polarization modulation formats, typically two triple MZMs are used in parallel, each modulating an orthogonal polarization. The two unmodulated carriers come from the same laser and are split into orthogonal linear polarizations with a polarization beam splitter (PBS), then the two independent polarization modulated signals are multiplexed together with a polarization beam combiner (PBC).

The PM signal demodulation at the receiver was described in Section 2.9.2. As a result of the demodulation process four electrical signals or data lanes corresponding to the in-phase and quadrature components of the X and Y polarizations are generated. These four signals are then digitized by four ADCs after which DSP can be performed for signal demodulation.

Coherent optical transceivers now utilize DSP, with the transmitter being responsible for modulation, pulse shaping, and pre-equalization, and the receiver responsible for equalization, synchronization, and

demodulation. At the transmitter, the DSP, in conjunction with the DACs and FEC, convert the incoming data bits into a set of analog signals. In correspondence with the operation of the transmitter, the major advantage of receiver-side DSP stems from its ability to arbitrarily manipulate the electrical field, after the ADC enables the sampling of the signal into the digital domain.

First, the four digitized signals after an ADC are passed through the block to compensate for front-end imperfections. The imperfections may include a timing skew between the four channels, attributable to the differences in both the optical and electrical path lengths within a coherent receiver. Other types of front-end imperfections can manifest in the difference between the four channels' output powers, due to different responses of PINs and TIAs in the receiver, and quadrature imbalance because the optical hybrid may not exactly introduce a perfect 90-degree phase shift.

Second, the major channel transmission impairments -- in particular, CD and PMD -- are compensated through digital filters. The static equalization for CD estimation and compensation is performed first, because of its independence of SOP and modulation format, plus, the impact on the subsequent blocks before the CD estimation is needed to achieve accurate compensation. Then the clock recovery can be processed to track the timing information of incoming samples. Note that it is possible to perform joint processing between the blocks of clock recovery and polarization demultiplexing for achieving the symbol synchronization. A fast-adaptive equalization is carried out jointly for two polarizations through a butterfly structure. Then the frequency offset between the source laser and the LO is estimated and removed to prevent constellation rotation at the intradyne frequency.

Finally, the carrier phase noise is estimated and removed from the modulated signal, which is followed by symbol estimation and hard- or soft-decision FEC for channel decoding. Note that for a particular digital coherent receiver, the ordering of DSP flow may differ slightly because of different design choices. Besides the feed-forward process, it is possible to perform joint processing and feedback among different process blocks, such as clock recovery and butterfly structured polarization demultiplexing.

On top of these typical demodulation processes, the huge amounts of parameters can be estimated and monitored using coherent DSP. In that case, the DSP would be dedicated to recycle data from the coherent demodulation process, so as to turn the coherent transponder into a multi-purpose measuring instrument for network management purposes.

4.2. Coherent Optical Performance Monitoring

Various optical performance monitoring (OPM) techniques using coherent detection have been proposed in scientific literature to monitor one or multiple parameters independently or jointly [4] [5]. Highly desirable features of a coherent optical system include higher robustness, reconfigurability, and flexibility. To enable robust and flexible operation, the coherent optical system should be able to:

- (1) measure its physical state and the quality of the propagating data signals;
- (2) automatically diagnose and repair the failures;
- (3) take actions before data loss and failure occur; and
- (4) allocate resources, including signal wavelength/power, tunable compensation, channel coding, and channel bandwidth.

Figure 21 shows the typical list of monitoring parameters for signal quality supervision at the optical layer.

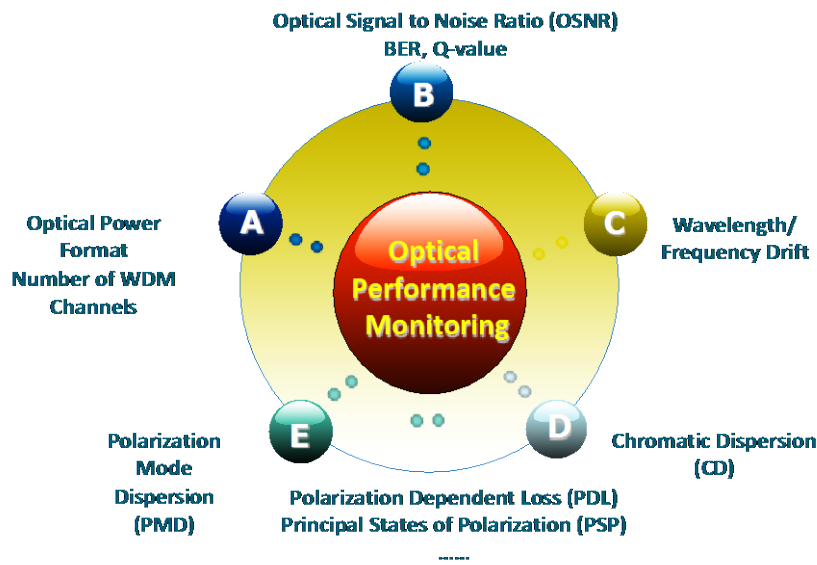


Figure 21 – Typical Parameters at Optical Layer

Monitoring at the digital and optical layers is an important element of PNM. It has shown competitive advantages in that it simplifies system design, optimizes system performance, shortens system installation, and lowers operations costs. Proactive network maintenance relies on advanced trend analysis of characteristic performance parameters that are observed at regular intervals over a long period of time. Repair is initiated as soon as negative trend is visible, and normally long before the client layer is affected. A lot of parameters shown in Figure 21, such as power, OSNR/BER/Q-value, and polarization tracking speed, can be used for proactive maintenance purposes.

It is worth mentioning that the equalizers that are used to compensate PMD and CD also provide great insight into channel distortion and the potential causes for such impairment. Amplitude ripple generated by reflections, filter narrowing, and DGD due to multiple cascading of wavelength multiplexers are a few examples of the information obtained.

If one follows conventional practices for monitoring and managing coherent optical systems, one could rely on external devices such as optical spectrum analyzers and RF instrumentation. However, in the access environment where the number of coherent links could easily be two orders of magnitude to what is found in the backbone, a more scalable management strategy could be implemented. In coherent optical links, the baseband representation of the optical field (amplitude and phase) in the electrical domain, and its digitization, leads to effective post-detection processing techniques in digital domain as introduced in above section. These digital equalizer structures embedded in coherent transceivers can not only compensate for all deterministic linear channel impairments, but can also enable a comprehensive optical PNM. That matters because it provides information about the fiber linear parameters in a simple, cost- and power-effective way. Expensive external devices are not required to evaluate optical properties or to tap the optical signal, which eventually reduces the effective received optical power. In addition, DSP-based OPM techniques are adaptable to varying data rates and modulation formats, and are capable of realizing and jointly monitoring different parameters. Therefore, coherent systems provide a better way to support fault forecasting, detection, diagnosis, and localization. Additionally, they provide a resilience mechanism in addition to basic monitoring capabilities of optical signal power level and wavelength for both traditional analog and intensity modulated access networks.

4.3. Basic Operation Principle

There are a number of technical papers that demonstrated the CD, DGD, and OSNR monitoring techniques by analyzing a bank of finite-impulse response (FIR) filters arranged in a butterfly structure in the digital domain [5], which is shown in Figure 22 below.

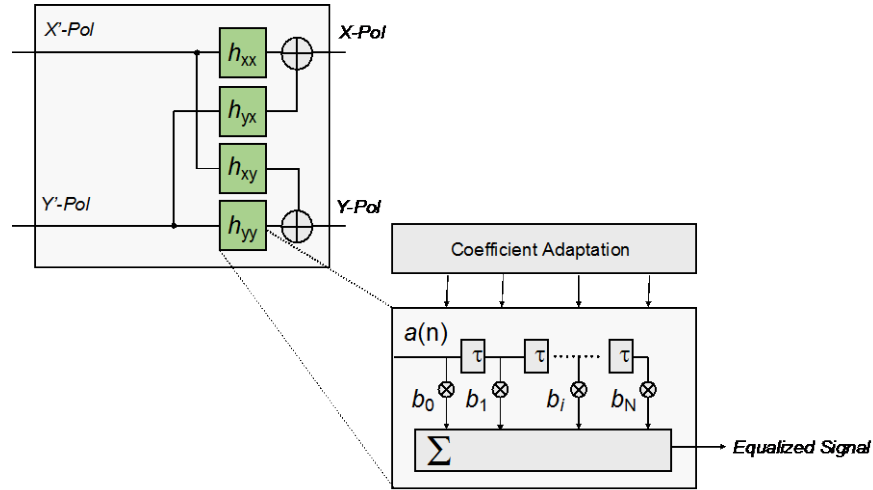


Figure 22 – Butterfly-structured Equalizer for Coherent Optical System

In such an equalizer, once the tap-setting algorithm for blind adaptation has converged, the filter's transfer function $H^{-1}(f)$ can be assumed as the inverse response of the fiber link $H(f)$. This equalizer consists of four complex-valued FIR filters arranged in this butterfly structure, which can be described with a single Jones matrix:

$$H^{-1}(f) = \begin{pmatrix} h_{xx}^{-1}(f) & h_{xy}^{-1}(f) \\ h_{yx}^{-1}(f) & h_{yy}^{-1}(f) \end{pmatrix}$$

On the other hand, the fiber channel is modeled by a concatenation of the following basic elements (the linear rotation of polarization is not included here):




Figure 23 – Simplified Fiber Model including Major Transmission Elements

Among them,

- CD: in frequency domain, CD can be described with $e^{j\psi f^2}$, where $\psi = 2 * \pi^2 \beta_2 z$ and β_2 is the propagation constant of fiber and z is the fiber length.
- DGD: only the first-order DGD is considered here, in frequency domain, $e^{+ \frac{j2\pi\tau f}{2}}$ and $e^{- \frac{j2\pi\tau f}{2}}$ are for two polarizations, where τ is the group delay.
- PDL: causes attenuation of $k(0 < k \leq 1)$ to the X-polarized optical field. The Y-polarized tributary remains unperturbed.

Now, the butterfly matrix can be expressed with the cascade of fiber channel elements:

$$H^{-1}(f) = D^{-1}(f) * E^{-1} * U^{-1}(f)$$

Where $D^{-1}(f)$ relates to CD, $U^{-1}(f)$ is the inverse DGD matrix, and E^{-1} is PDL vector. The phase and amplitude response of $H^{-1}(f)$ can be used to estimate the amount of DGD, CD, and PDL. For example, to estimate CD,

$$\begin{aligned} & \arg(h_{xx}^{-1}(f)h_{yy}^{-1}(f) - h_{xy}^{-1}(f)h_{yx}^{-1}(f)) \\ &= \arg(D(f)^2) \\ &= \arg((e^{-j\psi f^2})^2) \\ &= -2\psi f^2 \end{aligned}$$

We know that $\psi = 2 * \pi^2 \beta_2 z$, therefore the distance and chromatic dispersion value can be estimated from the tap value in this equalizer. An example of the given versus the estimated CD is given in Figure 24 [6]. As the estimation error is relatively small with respect to the absolute values, the deviation of the estimation is only visible with large magnification (see inset).




Figure 24 – Given versus Estimated CD Example

For PMD, the monitoring system needs to ensure it will not exceed values beyond which the adaptive filter can compensate. While the monitoring of SOP speed is of paramount importance in proactive failure detection, “slow” SOP monitoring would provide useful information on possible outages -- such as those that appear as a function of aging equipment. The idea is to monitor SOP fluctuations, over time, to evaluate the amplitude of this variation and assess the risk of outage, in the case of a fast SOP variation. Additionally, other impairments such as laser frequency offset and carrier phase are estimated and compensated in a digital coherent receiver. Other techniques such as artificial neural network (ANN) are also proposed for the monitoring of CD, PMD, and especially OSNR, using proper training sets.

4.4. Comparison Between Direct Detection and Coherent Detection

Unlike wireless networks, where all the necessary networking issues (such as link setup, optimization, and testing) are performed automatically, such tasks are currently handled manually in optical cable access networks -- a reality requiring substantial manual intervention. This is because the existing optical access networks are not yet capable of acquiring real-time information about the physical state of the network or the health of the signals propagating through the network. A number of OPM techniques have been proposed involving the time domain, frequency domain, or polarization domain for traditional direct detection systems, but the only digital technique is BER monitoring. When the migration from direct detection systems to digital coherent systems happens, many commonly used OPM techniques proposed for direct detection systems, such as interpolation-based out-of-band OSNR monitoring or polarization-nulling based in-band OSNR become practical.

Monitoring techniques are no longer suitable for coherent detection with very tight channel spacing and polarization multiplexing. Therefore, we have to take another look at the coherent system. Different from direction detection optical system, CD, PMD, PDL, and PSP are linear transmission effects that can be accurately estimated and fully compensated by linear digital filters at the optical coherent receiver. These fiber transmission parameters can be essentially monitored simply by reading the filter taps, as presented in the previous section and which come with almost no additional cost. In contrast to direct detection, the acquisition of channel parameters in general are inherent and integral in the coherent receiver. Table 4 compares the roles of various OPM functionalities in direct-detection to coherent systems.

Table 4: Comparison of OPM Functions in Non-coherent and Coherent Systems

	CD	PMD	PDL	PSP	Frequency Offset	OSNR	Power
Direct Detection	Additional Function Block Required						
	Analog performance detection in time domain (synchronous/asynchronous sampling); frequency domain (optical/RF domain); polarization domain (polarization nulling)						
Coherent Detection	Inherent in DSP process for estimation and compensation					Embedded DSP	

Through the inherent optical performance monitoring from coherent optical technology, operators can now understand exactly how much margin is currently present in the network, as well as the optimal capacity they can deploy. Combined with software-defined optics or networking analytics, applications such as predictive link failure now become possible, allowing operators to reconfigure their resource allocations.

4.5. Flexible resource allocation

In conventional direct detection-based optical access networks, channels, once initially provisioned, are seldom reconfigured until they are retired at the end of life cycle. The data rate, modulation format, capacity, and reach of a given provisioned channel is static and dependent on the specific transceiver interface being used, as well as the network environment. This static feature forces operators to maintain considerable safety margins, in order to provide a reasonable level of reliability. This results in an unintended waste of precious network resources. DOCSIS 3.1 specification-based analog optical channels introduced features to leverage the OFDM-based PHY layer, including variable bit loading and the option to define multiple modulation profiles on downstream and upstream channels. DOCSIS 3.1 OFDM (and OFDMA) profiles provide a wide range of modulation choices that can be used to fine-tune the CMTS's (and CM's) transmissions, to get the best performance from the current network conditions.

Now, the development of coherent optical technology will enable a similar operational style as DOCSIS technology, through the design of adaptive coherent transceivers, which are built to support a number of possible operational configurations, selectable by software. Such software-defined transceiver configurations can create specific modulation formats to support sets of data rates, corresponding tolerances to system impairments, and sets of electronic digital signal processing schemes chosen to function best in a given network environment. They can increase the network capacity and the spectral/energy efficiency, while providing a future-proof and flexible solution for an increasingly heterogeneous cable access network, from fixed to variable symbol rate and bit rate per channel. As an example, Figure 25 shows different modulation formats for different application scenarios, such as at 50 GHz or 100 GHz optical spacing, with or without optical amplification and DWDM Mux and DeMux. In other words, differing configurations can trade off the optical link margin with the data rate/capacity and power efficiency when coherent optics are introduced into cable access networks.




Figure 25 – Adaptive Coherent Transceiver to Support Different Scenarios

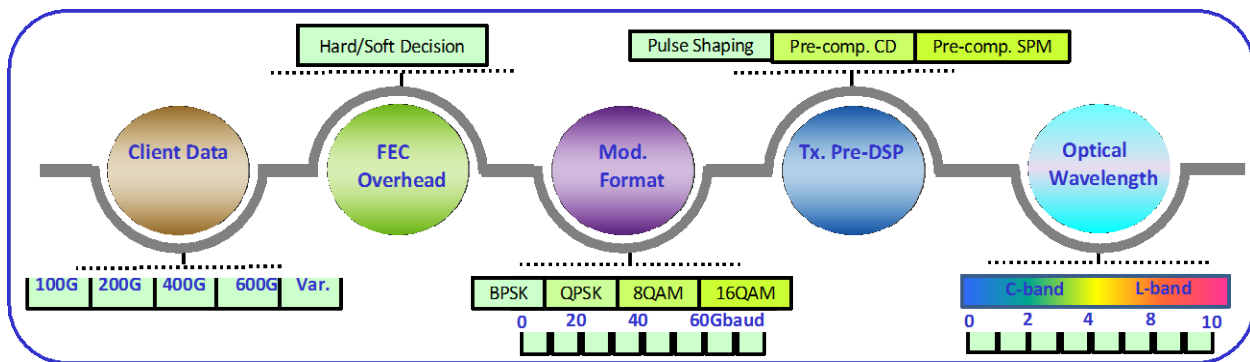


Figure 26 – Approaches for Flexible Data Rates and Software-defined Optics

Approaches exist for implementing flexible data rates with software-defined transceivers, which could operate at a constant symbol rate or could change between two or more symbol rates. Figure 26 illustrates the SDO concept. On the transmitter side, SDO may support variable client services. It can have different overhead coding schemas (hard decision or soft decision FEC) and configurable modulation formats. On the corresponding, transmitter-side DSP, they can support different numbers of optical wavelengths. At the receiver side, a universal DSP is needed for supporting different modulation formats and FEC coding schemas at different baud rates. In terms of ASIC design and implementation, optimization is required between performance and power consumption. The current implementations in the optical industry have demonstrated partial programmable capabilities in terms of modulation format (BPSK/QPSK/8QAM/16QAM), symbol rate, and FEC overhead adaptation.

5. Operational Strategy

An important aspect of PNM in coherent optical networks is that it can reduce the total cost of ownership (TCO) for operators. While it may be attractive to rely on existing PNM cost models, there are some differences that require a fresh look. One considerable difference in cost models is that with coherent optical networks, no benchmarks yet exist to baseline maintenance costs. This paper proposes that PNM capabilities will be available at the time of conception, so there will be no reduction in costs over time -- rather, it will be a matter of cost avoidance and operational efficiency. Similar to traditional coaxial PNM, the same problems exist that can make it difficult to demonstrate a favorable cost picture for PNM. This is because it involves making a claim about something that hasn't actually happened yet or possibly never will. Fortunately, having a history of PNM cost modeling in traditional HFC coaxial networks allows us to extrapolate certain cost avoidance. The model can be adjusted for construction, repair and maintenance of optical components, instead of active and passive coaxial components.

To begin understanding PNM cost avoidance, several stages of the network lifecycle will be considered. The lifecycle begins with finance, design and construction, which also includes the cost of materials. Next is activation and provisioning, which includes customer turn-up. Finally, and as important, is the ongoing support and maintenance of the network, including outage repair and customer disruption times.

Beginning with construction, there is essentially no additional expense to accommodate PNM features. Because PNM exists as an embedded capability, it is made up mostly of software components, which may require additional memory and processing. However, the size of these components is very small in the context of modern computing, so any additional cost to materials or construction can be considered negligible.

Conversely, the construction phase of the optical network can benefit from PNM in several valuable ways. The most significant factor in construction costs is labor, which can improve from automation of the post-construction quality control and certification checks. As we've learned in traditional coaxial PNM, many of the network faults began as small construction defects, which eventually deteriorated over time, such as corrosion and resultant micro-reflections. A similar paradigm exists in optical networks, in certain types of connectors, interfaces and fusion splice defects. While it's true that optical components do not corrode, connectors and optical interfaces can still get dirty over time. Likewise, a low-quality splice may deteriorate because of environmental influences, such as wind, water, stretching or enclosure contraction and expansion. These types of problems will be identified and located almost immediately, using the proposed PNM techniques. In addition to automated quality control checks, the amount of labor spent doing post-construction certification will be improved. This time savings could be realized as improved efficiency to allow for more splicing and reduced construction rework.

Most of the automation and quality control benefits will be realized during the network and customer activation stage. This is the point where multiple receivers will be activated, and will start providing valuable PNM information about the quality and condition of the optical links. Similar to contemporary DOCSIS operational support systems (OSS), the optical receivers will be providing remote telemetry data that is vital to assess the conditions of the optical signal.

Ongoing support and maintenance will benefit from the continuous reporting about network conditions at the receiver locations.

These models can be used to approximate the cost of common proactive repairs, which require rework of fusion splices and mechanical connector cleaning. Given the cost fusion splice repairs seen in Figure 27, the following model can be used to approximate the cost avoidance of proactively repairing splices.

Benchmarking Fusion Splice Time

The cross-tabs below, indicates the expected **Middle Position** of an achievable average.

Splice Protection Closures				
1 x fibre tech per joint	Cable size	Preparation	Splice and Coil	Total
	4-fibre	20-min	15-min	35-min
	8-fibre	20-min	25-min	45-min
	12-fibre	25-min	35-min	1-hr
	24-fibre	35-min	55-min	1-hr 30-min
	48-fibre	40-min	1-hr 30-min	2-hr 10-min
2 x fibre techs or a fibre tech and assistant per joint	Cable size	Preparation	Splice and Coil	Total
	72-fibre	1-hr 30-min	4-hr	5-hr 40-min
	96-fibre	2-hr 30-min	6-hr	8-hr 40-min
	144-fibre	4-hr	8-hr	12-hr

Unpopulated Patch Panels				
1 x fibre tech per panel	Cable size	Preparation	Splice and Coil	Total
	4-fibre	30-min	20-min	50-min
	8-fibre	35-min	30-min	1-hr 5-min
	12-fibre	40-min	40-min	1-hr 20-min
	24-fibre	45-min	60-min	1-hr 45-min
	48-fibre	50-min	2-hr 20-min	3-hr 15-min
2 x fibre techs or a fibre tech and assistant per panel	Cable size	Preparation	Splice and Coil	Total
	72-fibre	2-hr 30-min	6-hr	8-hr 30-min
	96-fibre	3-hr 30-min	7-hr	10-hr 30-min
	144-fibre	5-hr	9-hr	14-hr

Figure 27 – Average Time to Fusion Splice

Operational Support Systems (OSS)

Cable operators will always need systems to help them support their networks. This remains true in the case of coherent systems. A lack of remote diagnostics and reporting can be a costly mistake when deploying field-based network technology, resulting in significant costs and inefficiency associated with manual labor. One handy example of this condition is found in the case of traditional hybrid fiber-coaxial (HFC) optical nodes. For the vast majority of fiber nodes deployed in cable networks, a technician is usually required out in the field to take measurements and adjustments when needed. Of course, there are examples of fiber nodes which have been instrumented for OSS and remote management, but these are the exception and not the rule. Understandably, it was cost and power prohibitive to embed this type of remote monitoring several decades ago -- but that certainly is not the case today.

Engaging a standards body to construct management information bases (MIBs) is useful to achieve ubiquitous and consistent implementation. A good example of this can be found in CableLabs® and the PNM capabilities that are available in the DOCSIS standards, starting with version 2.0. Because of the well-understood benefits of proactivity, it's now standard practice to conceive these capabilities early in the design phase of the product. Furthermore, when the product eventually becomes deployed in the field, refinements to the specification or implementation may be needed. Considering the many facets of

unanticipated conditions out in the operating environment, a mechanism for refactoring requirements is almost always necessary.

Another important consideration for OSS is the timely delivery of information from the network sensors to the management systems. The most common network management protocols provide on-demand telemetry to satisfy immediate diagnostic and reporting needs. There is also usually an event-based reporting capability, to facilitate the time-sensitive, 24x7 monitoring of transient network events, like ingress, in traditional radio frequency (RF) networks. Many of these venerable protocols, such as Simple Network Management Protocol (SNMP), have been around for decades and tend to have their own scale and maintenance problems as a result. Fortunately, there are vastly improved protocols that work in a web-scale cadence, which address the shortcomings found in many Network Management Systems (NMS) and OSS systems. These protocols typically employ web sockets, streaming, and lightweight data models such as JavaScript Object Notation (JSON) that dramatically improve network agility, scalability, and maintainability.

Additionally, a mechanism is required to configure PNM functions such as event thresholds, northbound messaging URIs and OSS registration information.

The following Figure 28 illustrates the software stack, including PNM instrumentation, within the coherent optical receiver.




Figure 28 – Example of PNM Software Stack

Wavelength Detection and Inventory Management

Perhaps one of the most intriguing challenges and opportunities presented by coherent optical networking is the significantly increased density of available wavelengths. The intrinsic value of any network is its ability to transport payload on behalf of customers, measured in bits-per-Hertz. By increasing the bandwidth available to customers, the value of the network increases proportionately. However, the challenge is that contemporary optical networks lack the agility required to automatically detect and allocate wavelength spectrum utilization. This creates significant manual overhead and predictable failures to maintain consistent, timely records of the network.

By adding PNM enabled coherent optical receivers, operators will benefit from automated wavelength detection and mapping. Even when beginning with sparsely distributed receivers, operators will begin to populate their wavelength inventory over time due to the strategic location of these receivers. In fact, it's

conceivable that operators may discover lost or forgotten available spectrum due to the aforementioned failures created by manual process. This addresses the long-term cost of operation by dramatically reducing the cost and speed associated with its manual counterpart. It also provides added assurance to operators that they will be able to maximize the revenue potential of their costly investment.

Automated Service Activation with Software Defined Networking

As previously noted in Figure 28, the embedded PNM and Software Defined Networking (SDN) modules enable the remote monitoring and configuration required to automate the activation of wavelength services. This is especially attractive to operators interested in providing self-service provisioning directly to customers, thus avoiding additional time and expense associated to manual touch points within the system.

There are many popular open source SDN platforms that would be suitable to pair with coherent optical PNM, such as Open Network Operating System (ONOS) and OpenDaylight. This would enable wavelength programmability by the operator's NMS and drive further long-term cost reduction in the support and maintenance of the optical network. Likewise, it would increase the operator's ability to compete by removing days or weeks from the time it takes to activate new services that drive revenue growth. In many cases, new wavelengths could be allocated, activated and delivered to customers instantaneously, without a single human touch.

Access Network Resource Optimization

A fully instrumented monitoring and configuration mechanism also provides for highly elastic resource optimization -- such as for capacity management, redundant path routing and Quality of Service (QoS). A robust network optimizer will be capable of executing multiple optimizations for bandwidth, latency, cost and availability to satisfy the different use-cases previously mentioned.

This is particularly interesting when evaluating the relative value of these network resources to customers. With the new availability of additional capacity, traditional constraints may become a thing of the past. For example, when capacity becomes congested, adding additional fiber capacity may not be a practical option. However, additional wavelengths may now be allocated and provisioned in anticipation of peak congestion to deliver temporary relief of surge traffic. A good example of this might be to accommodate for a major sporting event or holiday. Then, when the anticipated surge traffic has abated, the NMS may tear down and release the resources back to the capacity pool. In this case, the surge traffic was accommodated in a seamless manner to the customers.

Similarly, operators may now offer new, high-value products to their customers that support improved performance and SLAs. As proposed in the elastic capacity example above, higher-order SLAs may now be offered due to the additional capacity and wavelength agility. By creating redundant access network paths, customers could enjoy higher, more reliable service performance by avoiding latency caused by adjacent, shared or impaired wavelength resources.

Co-Existence and Network Reliability

Last but not least, it is also important to consider the extended operational value of introducing coherent optical PNM to operators existing optical networks. Given that coherent signals co-exist well with legacy signals, the operator may quickly achieve improved network reliability by adding just a single coherent signal.

By sparsely distributing coherent receivers, this allows operators to realize the benefits of PNM across their pre-existing optical network. Wavelength detection, splice mapping, fault detection, increased performance, SLAs, QoS, capacity management and redundant access paths all become instantly available to the operator.

With this new visibility, old problems and new opportunities will become illuminated within the already sunk costs of the pre-existing optical networks.

Conclusion

This paper has reviewed the use of coherent optics links for the management, troubleshooting and assessment of health metrics. Health metrics not only of the coherent signal and the fiber paths it is transported in, but also indirectly of the health of all the other non-coherent optical signals, such as IM-DD and analog optical signals that share the same fiber segments with the coherent signal. Coherent links provide a rich set of metrics that provide insight on distortion, noise, link length, polarization state, loss and reflections.

Coherent links have higher sensitivity and higher robustness than IM-DD and analog optical link types. They can sense impairments as soon as the fiber path is affected and provide feedback when non-coherent links may have already ceased to operate. Coherent transceivers become effective health probes of fiber access networks.

Once these rich metrics are correlated with fiber topology, analytics can be leveraged to determine the type of impairment, severity and location. This allows the operator to assess impact and prioritize repairs.

Coherent transceiver capabilities enable the implementation of network embedded instrumentation. Embedded instrumentation will provide cable operators significant CAPEX reduction and ubiquitous coverage through already deployed probes. Traditional long haul and metro transceiver don't provide all the metrics that have been mentioned here because their environment does not require them. It is up to the optical transceiver manufacturers to meet the cable industry's need for certain parameters in the chip so that operators and management systems can analyze and extract the required functionality.

Flexibility and adaptability that can be achieved through SDO enables optimization of performance. The flexibility to remotely configure the optical network along with accurate record keeping helps manage resources better and avoid leaving them stranded.

As cable provides optical connectivity services, in order to maximize existing infrastructure, it has to migrate from fiber services to wavelength services. Cable service providers with few fiber strands available have to be mindful how to use this precious resource. Effective wavelength management rather than fiber management results in a much longer lifespan of cable infrastructure.

Abbreviations

ADC	analog to digital converter
ANN	artificial neural network
ASE	amplified spontaneous emission
ASIC	application-specific integrated circuit
BER	bit error rate
bps	bits per second
BPSK	binary phase shift keying
CAPEX	capital expense
CD	chromatic dispersion
CM	cable modem
CMTS	cable modem termination system
DAA	distributed access architecture
dB	decibels
dBm	decibels milliwatt
DEMUX	wavelength demultiplexer
DFB	distributed feedback laser
DGD	differential group delay
DOCSIS	data over cable system interface specification
DP-QPSK	dual polarization-quadrature phase shift keying
DP-QAM	dual polarization-quadrature amplitude modulation
DSP	digital signal processing
ECL	external cavity laser
EDF	erbium-doped fiber
EDFA	erbium-doped fiber amplifier
EPON	ethernet passive optical network
EVM	error vector magnitude
FEC	forward error correction
FP	Fabry-Perot laser
FWM	four wave mixing
GHz	giga-hertz
GPON	gigabit passive optical network
HFC	hybrid fiber-coax
IF	Intermediate frequency
IQ	in-phase and quadrature
JSON	javascript object notation
LO	local oscillator
MER	modulation error ratio
MIB	management information base
MUX	multiplexer
MZM	Mach-Zehnder modulator
NMS	network management system
ODC	optical distribution center

OMA	optical modulation analyzer
ONOS	open network operating system
OOK	on-off keying
OSA	optical spectrum analyzer
OSS	operations support systems
OTDR	optical time domain reflectometer
OVA	optical vector analyzer
N+0	node plus zero amplifiers
PBC	polarization beam combiner
PBS	polarization beam splitter
PDL	polarization dependent loss
PIN	p-type, intrinsic and n-type layer diode
PM	polarization multiplexing
PMD	polarization mode dispersion
PNM	proactive network maintenance
PON	passive optical networks
PSP	principal state of polarization
QAM	quadrature amplitude modulation
QoS	quality of service
RF	radio frequency
RFoG	radio frequency over glass
Rx	receiver
S	signal
SCM	sub-carrier multiplexing
SDN	software defined networking
SDO	software defined optics
SLA	service level agreement
SNMP	simple network management protocol
SOA	semiconductor optical amplifier
SOP	state of polarization
TCO	total cost of ownership
TIA	trans-impedance amplifier
TRx	transceiver
Tx	transmitter
XI	in-phase x-polarization component
X POL	x-polarization
XQ	quadrature x-polarization component
YI	in-phase y-polarization component
Y POL	y-polarization
YQ	quadrature y-polarization component
WDM	wavelength division multiplexing

Bibliography & References

- [1] Z. Jia, "Impact of Access Environment in Cable's Digital Coherent System – Coexistence and Full Duplex Coherent Optics," *SCTE Cable-Tec Expo*, 2018
- [2] Z. Jia, L. A. Campos, C. Stengrim, J. Wang, C. Knittle, "Digital Coherent Transmission for ext-Generation Cable Operators' Optical Access Networks," *SCTE Cable-Tec Expo*, 2017
- [3] L. A. Campos, Z. Jia, T. Liu, "Leveraging deployed fiber resources for the implementation of efficient scalable optical access networks," *Sept. SCTE/ISBE Cable-Tec Expo'16*, 2016
- [4] C. K. Chan, "Advanced Optical Performance Monitoring for Next Generation Access Networks," *OECC*, 2013
- [5] Z. Dong, F. N. Khan, Q. Sui, K. Zhong, C. Lu, A. P. T. Lau, "Optical Performance Monitoring: A Review of Current and Future Technologies," *J. Lightw. Technol.*, vol. 34, no. 2, pp. 525–543, Jan. 2016
- [6] F. N. Hauske, M. Kuschnerov, B. Spinnler, B. Lankl, "Optical Performance Monitoring in Digital Coherent Receivers," *J. Lightw. Technol.*, vol. 27, no. 16, pp. 3623–3631, August 2009
- [7] R. A. Soriano, F. N. Hauske, N. G. Gonzalez, Z. Zhang, Y. Ye, and I. T. Monroy, "Chromatic Dispersion Estimation in Digital Coherent Receivers," *J. Lightw. Technol.*, vol. 29, no. 11, pp. 1627–1637, June 2011

Quality-of-Experience Monitoring, Optimization and Management: A Unified End-to-End Solution

A Technical Paper prepared for SCTE•ISBE by

Zhou Wang, PhD, FIEEE
Chief Science Officer, SSIMWAVE Inc.
Professor, University of Waterloo
zhou.wang@uwaterloo.ca

Abdul Rehman, PhD
Chief Executive Officer
SSIMWAVE Inc.
abdul.rehman@ssimwave.com

Table of Contents

Title	Page Number
Table of Contents	2
Abstract	3
Introduction.....	3
Content.....	4
1. End-to-End Visual QoE Monitoring, Optimization and Management.....	4
2. Objective QoE Metric	5
3. QoE-Driven Optimization	8
Conclusion.....	10
Abbreviations	11
References.....	11

List of Figures

Title	Page Number
Figure 1. Unified end-to-end QoE monitoring, optimization and management framework in a video distribution system.....	4
Figure 2 - Applications of unified end-to-end QoE monitoring, optimization and management system.....	5
Figure 3 - Quality prediction accuracy performance evaluation of objective QoE metric	6
Figure 4 - Illustration of how bandwidth saving is achieved for given target quality (SSIMPLUS=90) by using a QoE metric that is able to provide consistent cross-content evaluation.....	9
Figure 5 - Illustration of how bandwidth saving is achieved for given target quality (SSIMPLUS=90) by using a QoE metric that is able to provide consistent cross-resolution evaluation.....	9
Figure 6 - Illustration of how bandwidth saving is achieved for given target quality (SSIMPLUS=90) by using a QoE metric that is able to provide consistent cross-device evaluation.	10

Abstract

Traditional quality assurance methods for large-scale video distribution networks operate independently at different points along the video delivery chain, reporting partial and incoherent measurements, leading to poor and fragmented understanding about how multiple stages of quality degradations affect the final quality-of-experience (QoE) of end users. We propose a framework that uses a unified end-to-end solution to produce consistent QoE scores at all points along the delivery chain under the same evaluation criterion. The novel solution produces a clear and complete picture instantaneously about how video QoE degrades over the network, allows immediate issue identification, localization and resolution, enables quality and resource optimization, and provides reliable predictive metrics for long-term strategic resource and infrastructure allocations. The main challenge in the implementation of the solution is to create a unified QoE metric that not only accurately predicts human perceptual QoE, but is also lightweight and versatile, readily plugged into multiple points in the video delivery chain. The QoE metric should produce real-time QoE scores across a wide range of bitrates, resolutions, frame rates and dynamic ranges, and combine presentation picture quality with the perceptual impact of video freezing and adaptive streaming events. We show that the SSIMPLUS metric offers the best promise to meet all the challenging demands.

Keywords

Quality-of-experience, video distribution system, video delivery chain, video quality assessment, video streaming, end-to-end quality assessment, video encoding, adaptive streaming

Introduction

There has been a remarkable growth of video distribution services in the past few years [1]. While common consumers are enjoying the video streams delivered to their TVs, smart phones and tablets, they often complain about the quality of the video they are experiencing [2]. Meanwhile, content producers are concerned about whether their creative intent is properly preserved during the video distribution process [3], [4]. Quality assurance (QA) is an essential component to warrant the service of video distribution systems. Traditionally, QA has been network-centric, focusing on the quality-of-service (QoS) [5] provided to the users, where the key metrics are determined by the network service level parameters such as bandwidth, package drop rate, and network delay. However, QoS metrics have fundamental problems in tracking what the users are actually experiencing. Recently, Quality-of-Experience (QoE) [6], which measures “the overall acceptability of an application or service as perceived subjectively by the end-user” [7], has been set to replace the role of QoS. In practice, the actual meaning of “QoE” measurement could vary significantly from one solution to another. For example, simple device playback behaviors such as statistics on the duration and frequency of video freezing events, may be employed to create a crude estimate of visual QoE. Such simple measures only provide a rough idea about how certain components of the video delivery system perform, but are distance away from what we really need in terms of accuracy, comprehensiveness and versatility. Moreover, the perceptual artifacts that affect picture quality are not properly measured, and the large perceptual differences due to viewing conditions are not properly taken into consideration. Consequently, they are at best “pseudo-QoE measures” or “QoS measures at the client”, and are difficult to be used to localize quality problems, to optimize system performance, and to manage the visual QoE of individual users.

We propose a unified end-to-end framework for QoE monitoring, optimization and management. The general philosophy is to align all measurements with the visual QoE of end users. Keeping this in mind, any design and resource allocation in the video distribution system, regardless of if it is for the whole system or for any individual component at the head-end, media data center, network, access server, or user device, should be evaluated, compared and optimized for one criterion, i.e., the impact on end users' QoE. To make such a system work properly, the most challenging task is to find a highly accurate, efficient and versatile QoE metric. Such a QoE metric, deployed throughout the video distribution system, establishes the basis for unified QoE monitoring, optimization and management.

Content

1. End-to-End Visual QoE Monitoring, Optimization and Management

Figure 1 illustrates a general framework of modern video distribution systems. When the source video content is received, it passes through a sophisticated video delivery chain consisting of many processing, encoding, transcoding, packaging, routing, streaming, decoding, and rendering stages before it is presented on the screen of individual users' viewing devices. To ensure the video is faithfully and smoothly delivered to the consumer device, the ideal quality assurance method would be to have human inspectors placed at all transition points along the chain, so that any quality issue can be identified instantaneously, and all measurements can be compared directly. In practice, however, this is infeasible because it requires thousands of source video streams and millions of derivative streams to be evaluated continuously by human inspectors, a non-scalable resource in the real-world. A viable solution is to replace humans with objective QoE monitoring probes, as illustrated in Fig. 1, which constantly predict human QoEs based on objective QoE metrics at the corresponding inspection spots.

There are two essential properties of such QoE monitoring probes. First, they should “see” and “behave” like human inspectors. More specifically, they should “perceive” all the actual pixels of all video frames like humans, and they should produce QoE scores just like what humans would say about the video quality when seeing the same video streams. Second, they should provide a “unified end-to-end” monitoring solution in the sense that the QoE evaluation methods at all transition points along the video delivery chain are designed under the same evaluation framework and compatible methodology to produce consistent quality scores that are directly comparable.

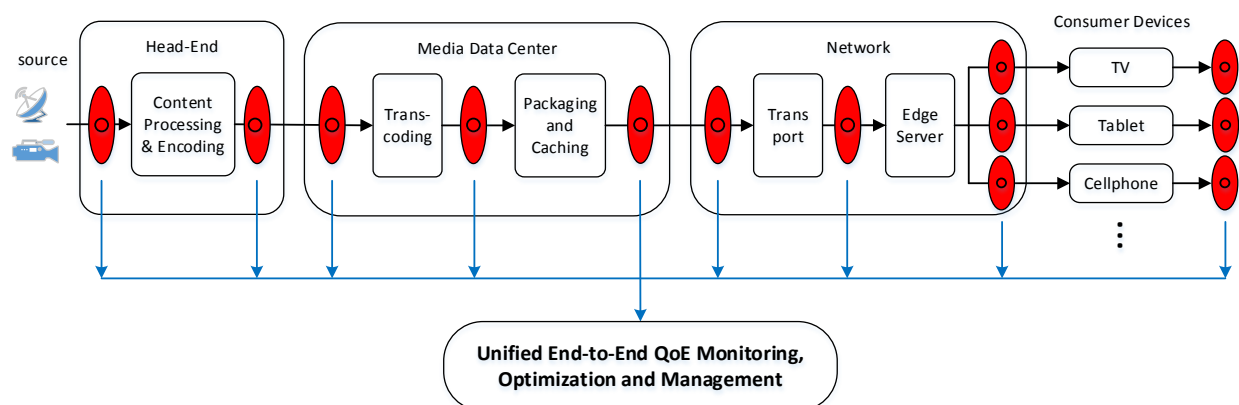


Figure 1. Unified end-to-end QoE monitoring, optimization and management framework in a video distribution system.

Once QoE monitoring probes are deployed throughout the video delivery chain, QoE data can be collected instantaneously and continuously. Subsequently, statistics can be computed at different time-scales (minutes, hours, days, weeks, months, years). These lead to many valuable benefits, as described in Fig. 2. More specifically,

- Operation engineers are able to gain immediate awareness about how video QoE degrades along the video delivery chain. As such, quality problems can be immediately identified, localized and resolved.
- Design engineers are able to closely observe the QoE variations between the input and output of individual components or the whole video delivery system as a whole. This helps them perform better design and optimization that target at improving and stabilizing the QoE of end users.

Managing executives are able to obtain a clear picture about how video quality evolves throughout the video distribution system and over long time scales. When long-time, large-scale data has been collected, big data analytics can be performed to help make intelligent strategic decisions on the operations of the system.




Figure 2 - Applications of unified end-to-end QoE monitoring, optimization and management system.

2. Objective QoE Metric

At the core of the end-to-end QoE monitoring framework is the QoE quality metric, which mimics human behaviors in evaluating video quality, and is the most challenging technical problem to solve. A good objective QoE metric combines deep understanding of the human visual system with advanced computational models and algorithms. It also requires smart design and efficient implementation of the algorithms and systems. Traditional approaches such as peak signal-to-noise-ratio (PSNR) have been shown to have poor correlations with perceptual video quality. More advanced perceptual video quality assessment (VQA) methods such as the structural similarity index (SSIM) [8], [9], multi-scale SSIM (MS-SSIM) [10], video quality model (VQM) [11] and video multi-method assessment fusion (VMAF) [12] improve upon PSNR but are still limited in prediction accuracy. More importantly, these traditional VQA approaches have fundamental limitations in their application scopes, functionalities and/or computational cost. These limitations largely impede them from being deployed broadly in real-world video distribution systems. When they are faced with the unified end-to-end QoE monitoring challenge we are targeting here, these disadvantages become even more pronounced.

To meet the challenge in a unified end-to-end QoE monitoring system, an objective QoE metric requires to have a number of must-have features. These include:

- *Accurate and light-weight.* The QoE metric must produce quality scores that accurately predict human visual QoE. The metric should be verified using independent, large-scale subject-rated video databases with diverse content and distortion types, and show high correlations with the opinions of an average human subject, as demonstrated by the scatter plot produced by the SSIMPLUS metric [13], [14] shown in Fig. 3. Meanwhile, the metric needs to be light-weight, allowing for real-time computations of high resolution videos (e.g., full high definition (HD), ultra-high definition (UHD) and 4K videos) with moderate hardware configurations. Such light-weight

and speed requirement is critical in large-scale video distribution systems to reduce the overall cost and to maximize the flexibilities in terms of deployment, integration, scaling, and customization.




Figure 3 - Quality prediction accuracy performance evaluation of objective QoE metric

Each point in the scatter plot represents a test video. The horizontal and vertical axes are the quality prediction from an objective quality metric (in this case the SSIMPLUS metric [13], [14]) and the mean opinion score (MOS) obtained from subjective test, respectively. A good quality metric should produce a narrow-band cluster extending from low to high quality ranges, regardless of the mixed video content, resolution and viewing devices, as exemplified by the SSIMPLUS metric [13], [14] shown in the figure. The spearman rank-order correlation coefficient (SRCC) between SSIMPLUS and MOS is 0.97.

- *Easy-to-understand and easy-to-use.* The QoE metric must be easy-to-understand, directly producing QoE scores that linearly scale with what an average consumer would say about the video quality. For example, if the quality score range of the metric is between 0 and 100, then the total scale range may be divided into five even segments corresponding to five perceptual QoE categories of bad (0-19), poor (20-39), fair (40-59), good (60-79), and excellent (80-100), respectively. The QoE metric must be deployed with an easy-to-use user interface (UI), where the presentation is simple and intuitive, focusing on the most important trending information. Such an easy-to-understand and easy-to-use QoE metric defines a common language, under which engineers can identify/fix quality problems and optimize system performance, and executives are able to make critical business decisions.
- *Applicable and consistent across resolutions, frame rates, dynamic ranges, user devices and contents.* In addition to accuracy and speed, another critical problem that hinders the wide usage of existing well-known video quality metrics (PSNR, SSIM, MS-SSIM, VQM, VMAF) is their limited applicability. In particular, when videos are of different spatial resolutions, frame rates, and dynamic ranges, these metrics are not applicable, because all of them require pixel-to-pixel correspondence. Moreover, when the same video stream is displayed on different viewing devices (e.g., TV vs. tablet vs. smartphone), the perceptual QoE could be significantly different. However, all traditional metrics fail to make meaningful device-dependent QoE predictions.

Furthermore, these quality metrics often produce inconsistent scores across different content types (e.g., sports vs. news vs. animations), strongly limiting the usefulness of such metrics in large-scale distribution systems that operate on thousands of video service channels to make resource allocation decisions across the whole systems. Therefore, to implement a unified end-to-end quality assurance framework for many real-world video distribution systems (e.g., for multi-screen and adaptive bit rate (ABR) video delivery networks), consistent and cross-resolution, cross-frame rate, cross-dynamic range, cross-viewing device, and cross-content QoE assessments are essential.

- *Versatile for usage in single-ended, double-ended and more sophisticated scenarios.* Single-ended and double-ended video quality assessments refer to the different application scenarios where a reference video may or may not be available when assessing the quality of a test video. Double-ended or full-reference (FR) quality measures assume the reference video is accessible and of perfect quality. They are essentially signal fidelity measures and PSNR, SSIM, MS-SSIM, VQM and VMAF all belong to this category. On the other hand, single-ended or no-reference (NR) measures do not assume access to the reference video. Double-ended quality measures typically have higher quality prediction accuracy than single-ended approaches, but are much more difficult to apply. Very often, the reference videos are completely inaccessible. Even when they are accessible, for example, at video transcoders, the reference videos are often not well aligned with the test videos both in space and time. Moreover, the source videos received from content providers are often distorted themselves, creating even more complex scenarios where the reference videos are already degraded. In order to provide consistent QoE assessment at all points along the video delivery chain, the QoE metric has to be extremely versatile. The QoE metric needs to be easily plugged into single-ended, double-ended and more sophisticated scenarios. It also needs to make the best use of all resources to produce the most accurate QoE predictions. For example, at the transcoder, the QoE metric needs to precisely align the source and test videos before applying double-ended fidelity assessment. It also needs to appropriately handle the case when the reference video quality is already degraded.

All of the above are critical features for a QoE metric to work effectively in a unified end-to-end quality monitoring framework. Conventional and well-known video quality metrics (PSNR, SSIM, MS-SSIM, VQM, VMAF), however, are distant away from meeting these requirements. In practice, their usage is often limited to laboratory-testing environment, restricted to small-scale, non-time-critical use cases, e.g., encoder comparison on videos of the same content, spatial resolution, frame rate, and dynamic range.

The large gap between the limited performance and functionality of the well-known video quality metrics and the essential requirements of large-scale unified end-to-end QoE monitoring systems has motivated the development of the SSIMPLUS video QoE metric, which has been set to meet all the requirements throughout its design and implementation phases [13]. A recent study using 10 independent publicly-available subject-rated video databases (created from a collection of hundreds of thousands subjective ratings) evaluates conventional and state-of-the-art video quality metrics (including PSNR, SSIM, MS-SSIM, VMAF, SSIMPLUS and several other metrics), by comparing the quality predictions of these metrics against subjective mean opinion scores (MOS) [14]. The results showed that SSIMPLUS achieves the highest QoE prediction performance in terms of its correlation coefficients against MOS. It appears to be the only QoE metric that achieves an average correlation coefficient higher than 0.9. The same study also found that the SSIMPLUS metric to be 16.4 times faster than the VMAF metric, allowing SSIMPLUS to be computed in real-time in real-world applications [14]. The SSIMPLUS metric is applicable and produces consistent scores across resolutions, frame rates, dynamic ranges and content types. For every single video stream, it generates multiple QoE scores corresponding to a wide spectrum of viewing devices, from small screens on cellphones to large-size TVs. When applied to ABR encoding, SSIMPLUS simultaneously computes single-ended QoE scores of the source video input, together with

double-ended scores for all the derivative video output produced by transcoders with different bitrates and resolutions. As well, it provides the absolute QoE scores of the derivative streams considering that the source input does not have perfect quality. At the client side, SSIMPLUS combines picture presentation quality with the impact of switching and stalling events to produce an overall QoE assessment for each individual user on a per-view basis [15], [16]. All of these computations are done at a speed faster than real-time. Due to these features, SSIMPLUS has been successfully deployed in large-scale operational environments, running 24/7 reliably and affecting the viewer experience of millions of users.

3. QoE-Driven Optimization

Many benefits come naturally once a unified end-to-end QoE monitoring solution is in place. The benefits are usually maximized through QoE-driven optimization. Here we use bandwidth optimization as an example. Bandwidth reductions without maintaining the right level of visual QoE makes little sense. Due to the lack of proper QoE assessment tools, currently most bandwidth optimization approaches in the industry result in inefficient and unstable results. The first step to success is to adopt a reliable QoE metric of superior accuracy and speed performance, and broad and powerful functionality. For example, it needs to perform meaningful and consistent video QoE assessment across resolutions, frame rates, dynamic ranges, viewing devices and video content.

Here we use SSIMPLUS as an example to illustrate how the cross-content, cross-resolution and cross-device features of a QoE metric may be employed to produce large bandwidth savings. Figure 4 plots the rate-quality curves of two video content (titles) at the same full-HD (1080p) resolution, assuming they are viewed on the same TV device. The rate-quality curve (or alternatively rate-distortion curve) is a widely used tool in the video coding technical community to evaluate and compare the performance of video encoders. Given a video title, together with its resolution and the quality evaluation criterion, each point on the rate-quality curve represents an operation point of the encoder in terms of a bitrate-quality combination. Thus, when we attempt to encode two titles, we end up with two rate-quality curves, as shown in Fig. 4. The gap between the two curves reveal the difference in encoding difficulty between the titles. To reach a target QoE quality level (e.g., SSIMPLUS = 90) using a fixed bandwidth (e.g., 4Mbps) to encode both videos would be a waste. Indeed, while 4Mbps is necessary for Title 1 to achieve the target quality level, only 3.1Mbps is necessary for Title 2 to achieve the same target, leading to a significant bandwidth saving.




Figure 4 - Illustration of how bandwidth saving is achieved for given target quality (SSIMPLUS=90) by using a QoE metric that is able to provide consistent cross-content evaluation




Figure 5 - Illustration of how bandwidth saving is achieved for given target quality (SSIMPLUS=90) by using a QoE metric that is able to provide consistent cross-resolution evaluation

For the same video content (title), when it is converted and then encoded to multiple resolutions, each resolution produces a different rate-quality curve, as exemplified in Fig. 5, where HD (720p) and Full HD (1080p) resolutions are used. It is commonly observed that the rate-quality curves for different resolutions cross at certain bitrate, as illustrated in Fig. 5. This is because when bitrate is high and compression artifacts are hardly visible, higher resolution video produces better sharpness and perceptual fidelity, but when bitrate gets lower, the quality of higher resolution video drops faster due to its high encoding difficulty. A good QoE metric that reflects such trend can help pick the most cost-effective resolution to

achieve the target quality while saving large bandwidth. For example, for the same target quality (SSIMPLUS=90), a bandwidth reduction from 3.1Mbps to 2.4Mbps is obtained by switching from 1080p to 720p resolutions, as shown in Fig. 5.

For the same video content (title) encoded at the same resolution, the perceptual QoE could still vary significantly when the video is presented on different viewing devices. This is demonstrated by the rate-quality curves for a TV and a cellphone shown in Fig. 6. When the user is known to use a cellphone rather than a TV to watch the video, a bandwidth of 0.8Mbps is sufficient to achieve the same target quality level (SSIMPLUS = 90), down from 2.4Mbps on a TV. With all the content, resolution and device factors are combined (from Fig. 4 to Fig. 6), a total of 80% bandwidth savings may be obtained.

The example given here is for demonstration purposes only. In practice users may be able to explore more or fewer than the three factors above for maximum cost-savings. Our study suggests that for most video content and most common usage profiles, an average cost saving of 20%-60% is typically achieved by properly adopting QoE metric-driven bandwidth optimization. Such bandwidth savings can be obtained in both live and file-based video distribution systems by smart operation of video encoders and transcoders at the server, and may also be incorporated into adaptive streaming frameworks to achieve similar goals in a dynamic way.




Figure 6 - Illustration of how bandwidth saving is achieved for given target quality (SSIMPLUS=90) by using a QoE metric that is able to provide consistent cross-device evaluation.

Conclusion

We propose a solution for unified end-to-end QoE monitoring, optimization and management in large-scale video distribution systems. The principle behind the solution is to start with end user's QoE in mind, such that all the QoE monitoring points should produce instantaneous scoring that reflects the end user's

QoE up to the monitoring point in the video delivery chain. The QoE scores need to be accurate, consistent and directly comparable, such that the monitoring solutions of the entire video distribution network speaks the same language. Such a unified end-to-end solution laid the groundwork for the subsequent operations for great benefits. Specifically, operation engineers will be able to immediately identify, localize and fix quality problem, design engineers will be able to perform effective and accurate optimizations on the video delivery chain and its individual components, and managing executives will have a clear picture about how video quality evolves throughout the distribution network and over long time scales, so as to make intelligent strategic decisions to manage the QoE of end users.

The most challenging task in implementing the proposed solution is to create an objective QoE metric that is not only accurate, fast, easy-to-understand and easy-to-use, but also applicable and consistent across resolutions, frame rates, dynamic ranges, viewer devices and contents. Moreover, it needs to be highly versatile for use in single-ended, double-ended and more sophisticated scenarios. Conventional and well-known video quality metrics such as PSNR, SSIM, MS-SSIM, VQM and VMAF fall short of meeting these requirements. As a result, their usage is limited to lab-testing environment or small-scale use cases. This has motivated the recent development of novel video QoE metrics such as SSIMPLUS [13], [14], which has been deployed in real-world large-scale QoE monitoring systems.

To further demonstrate the benefits of adopting the proposed framework and QoE metric, we use bandwidth optimization as an example, which demonstrates that large bandwidth savings can be obtained by adopting a QoE metric such as SSIMPLUS. With the wide deployment of the proposed solution and QoE metrics in large-scale video distribution networks. The QoE data collected in large and varying space and time-scales constitutes a valuable source for big data analytics and strategic intelligence, which is an interesting direction for future investigations.

Abbreviations

ABR	adaptive bit rate
FR	full-reference
HD	high definition
MOS	mean opinion score
MS-SSIM	multi-scale structural similarity
NR	no-reference
PSNR	peak signal-to-noise ratio
QA	quality assurance
QoE	quality of experience
QoS	quality of service
SRCC	Spearman rank-order correlation coefficient
SSIM	structural similarity
UHD	ultra-high definition
UI	user interface
VMAF	video multi-method assessment fusion
VQA	video quality assessment
VQM	video quality model

References

1. Cisco Inc., “Cisco Visual Networking Index: Forecast and Methodology 2015-2020”, 2016.

2. Limelight Networks, "The state of online video," <https://www.limelight.com/video/>, 2016.
3. C. Curtis, *et al.*, "American Society of Cinematographers Technology Committee Progress Report 2016," *SMPTE Motion Imaging Journal*, vol. 125, no. 7, pp. 43-58, Sept. 2016.
4. Z. Wang, "New quality-of-experience measurement technologies: streamlining how videos are delivered to consumers," *IEEE Signal Processing Society Blogs*, July 2017.
5. M. Seufert, S. Egger, M. Slanina, T. Zinner, T. Hobfeld, and P. Tran-Gia, "A survey on quality of experience of HTTP adaptive streaming," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, Sept. 2014.
6. O. Oyman, S. Singh, "Quality of experience for HTTP adaptive streaming services," *IEEE Communications Magazine*, vol. 50, Apr. 2012.
7. ITU QoE Recommendation ITU-T P.10/G.100, Amd.1, *New Appendix I Definition of Quality of Experience (QoE)*, 2007.
8. Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, Apr. 2004.
9. Z. Wang, L. Lu, and A. C. Bovik, "Video quality assessment based on structural distortion measurement," *Signal Processing: Image Communication*, vol. 19, Feb. 2004.
10. Z. Wang, E. P. Simoncelli and A. C. Bovik, "Multi-scale structural similarity for image quality assessment," *IEEE Asilomar Conference on Signals, Systems and Computers*, Nov. 2003.
11. M. H. Pinson, "A new standardized method for objectively measuring video quality," *IEEE Transactions on Broadcasting*, vol. 50, no. 3, pp. 312-322, Sept. 2004.
12. Z. Li, A. Aaron, I. Katsavounidis, A. Moorthy and M. Manohara, "Toward A Practical Perceptual Video Quality Metric," *Netflix TechBlog*, June, 2016
13. A. Rehman, K. Zeng and Z. Wang, "Display device-adapted video quality-of-experience assessment," *IS&T/SPIE Electronic Imaging: Human Vision & Electronic Imaging*, Feb. 2015.
14. SSIMPLUS: The most accurate video quality measure, <https://www.ssimwave.com/from-the-experts/ssimplus-the-most-accurate-video-quality-measure/>
15. Z. Duanmu, K. Zeng, K. Ma, A. Rehman, and Z. Wang "A quality-of-experience index for streaming video," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 1, pp. 154-166, Feb. 2017.
16. Z. Duanmu, K. Ma and Z. Wang, "Quality-of-experience of adaptive video streaming: exploring the space of adaptations," *ACM Multimedia*, Mountain View, CA, Oct. 2017.

Real-Time Analytics for IP Video Multicast

A Technical Paper prepared for SCTE•ISBE by

Dr. Claudio Righetti

Chief Scientist

Telecom Argentina

Agüero 2392, Buenos Aires, Argentina

Phone: +5411 5530 4468

crighetti@teco.com.ar

Emilia Gibellini

Data Scientist

Telecom Argentina

egibellini@teco.com.ar

Florencia De Arca

Data Scientist

Telecom Argentina

fdearca@teco.com.ar

Mariela Fiorenzo

Data Scientist

Telecom Argentina

mafiorenzo@teco.com.ar

Gabriel Carro

Senior VP R&D

Telecom Argentina

gcarro@teco.com.ar

Table of Contents

Title	Page Number
Abstract	4
Contents	4
1. Introduction.....	4
2. Motivation and Backgrounds.....	5
2.1. Definitions.....	5
2.2. Background	6
3. Systems Overview and Data Description.....	6
3.1. Data Description.....	7
4. TV User Behavior Analysis	8
4.1. Comparison on Live TV and VoD User Behavior.....	8
4.2. Regular Weekday.....	10
4.3. Major Events	12
4.4. Variation of Rankings in Time	14
5. Multicast Gain.....	15
5.1. Analysis at CDN Level	16
5.2. Analysis at Service Group Level	18
6. Real-Time Analytics	23
6.1. K-means Clustering.....	23
6.2. K-means Clustering applied to the selection of multicast channels.....	24
Conclusion.....	26
Abbreviations	27
Bibliography & References.....	28

List of Figures

Title	Page Number
Figure 1 - [a] Concurrence of OTT devices from Flow, colored by type of request. [b] Proportion of Live TV and VoD tunings.	9
Figure 2 - [a] Distribution of requests from STB of the Legacy system. [b] Proportion of Live TV and VoD tunings on March 14, 2018 from 8 p.m. to EOD.	9
Figure 3 - [a] Concurrence of Chromecast devices. [b] Concurrence of other OTT devices, on Flow, every 10 minutes. May 24, 2018.	10
Figure 4 - [a] Hourly access frequency to Live TV in Legacy system, on May 24, 2018. [b] Legacy STB playing Live TV channels simultaneously, on May 24, 2018 from 8 p.m. to EOD.	11
Figure 5 - [a] Concurrent tunings from Legacy STB. [b] Concurrent tunings from Flow STB. May 24, 2018 between 10 p.m. and 11 p.m. (busy hour).	11
Figure 6 - Comparison of observed distribution to Zipf-Mandelbrot.....	12
Figure 7 - [a] Concurrence of Chromecast devices. [b] Concurrence of other OTT devices of Flow, every 10 minutes. March 4, 2018.....	13
Figure 8 - [a] Hourly access frequency to Live TV in Legacy system, on March 4, 2018. [b] Legacy STB playing Live TV channels simultaneously, on March 4, 2018 from 8 p.m. to EOD (end of the day).....	13
Figure 9 - Concurrent tunings from Legacy STB on March 14, 2018 between 9 p.m. and 10 p.m. (match hour).	14

Figure 10 - [a] Correlation between the top 10 channels on July 1, 2017 and the top 10 on the 180 following days. [b] Correlations between the top 10, top 20 and top 30 channels, on July 1, 2017 versus the same rankings on the following 180 days.	15
Figure 11 - Multicast gain, as a percentage of the capacity needed with 100% unicast scheme.	17
Figure 12 - Multicast gain versus number of channels that are set to multicast. Based on data from May 20 to May 28, 2018 gain calculated for hour slots from 8 p.m. to midnight.....	17
Figure 13 - [a] Distribution of service group's size (HHP) by region. [b] Maximum multicast gain versus service group size.	18
Figure 14 - [a] Maximum multicast gain at service group level, colored by service group size. [b] Mean multicast gain at service group level.	19
Figure 15 - Popularity in Buenos Aires region versus other regions, on May 24, 2018. [a] Buenos Aires versus Córdoba. [b] Buenos Aires versus La Plata.	20
Figure 16 - Average multicast gain at service group level for different scenarios.	20
Figure 17 - [a] Capacity needed at service group level versus multicast channels count, by SG size. [b] Multicast gain distribution by region and scenario on May 23, 2018 from 9 p.m. to 10 p.m.....	20
Figure 18 - [a] Size of the cluster that groups the high access frequency channels -multicast cluster- by date, colored by type of event. [b] Access frequency versus date, channels colored by cluster. Data from July 1, 2017 to December 31, 2017.	25
Figure 19 – [a] K-means clustering applied to the views per channel by hour for OTT devices. [b] K-means clustering applied to the access frequency per channel by day for the Legacy system. Algorithm used to classify the signals between multicast and unicast. Blue dots represent multicast channels and red dots unicast.	25

List of Tables

Title	Page Number
Table 1 – DTV (Legacy) log sample	7
Table 2 - Flow log sample	7
Table 3 - Fixed parameter estimation for the mixed-effects model.....	21
Table 4 – Estimation of the capacity (Mbps) for a 500 HHP service group, by multicast channel count and region.	22
Table 5 - Estimation of the capacity (Mbps) for a 128 HHP service group, by multicast channel count and region.	22
Table 6 - Estimation of the capacity (Mbps) for a 64 HHP service group, by multicast channel count and region.	22

Abstract

In order to understand the impact of multicast implementation, it is necessary to collect data on key indicators such as the number of concurrent streams, the average bitrate, and the average bandwidth, among others. We use these indicators to estimate the gain, in terms of bandwidth, at a service group level. The aim of this paper is to analyze the way in which the gain varies according to the service group size and its location, and to obtain –through the usage of statistical modeling– a model that describes and quantifies this relationship. In addition, the gain is estimated under a wide variety of scenarios, to know how many channels should be set to multicast, and if there is any gain in having a real-time analytics system that updates what channels should be delivered using Multicast.

Contents

1. Introduction

It has been more than thirty years since the IP (Internet protocol) Multicast standardization work started [RFC] [1]. Much research has been conducted into the benefits of IP Multicast versus Unicast for Live video in access networks with xDSL (Digital Subscriber Line), FTTH (Fiber To The Home), DOCSIS (Data Over Cable Service Interface Specification) and Wireless technology. In particular, cable operators have been using technology for years to distribute digital video over IP backbone networks to multiple head-ends and hubs to feed broadcast QAMs (Quadrature Amplitude Modulation).

CableLabs-IP Multicast Working Group- published a document (“IP Multicast Adaptive Bit Rate Architecture Technical Report” [2]) describing how to put together two network concepts: Multicast and Adaptive Bitrate delivery, in what is called M-ABR (Multicast Adaptive Bitrate).

This approach enables IP video subscribers in the same node to consume a common linear video stream over the access network, thus reducing access network bandwidth requirements over Unicast delivery (where a separate stream is delivered to each subscriber). The adaptive video streaming is a type of technology responsible for delivering video through the Internet in an efficient way. This is done by selecting the image quality according to the resources of each user. Adaptive Bitrate streaming technologies are almost exclusively based on HTTP (Hypertext Transfer Protocol).

However, in the world of cable operators there are still some questions with regard to the benefit of implementing M-ABR in their networks. How convenient is that Multicast migrate to IP Video Service? If service areas tend to be reduced, does that situation justify the implementation of this technology? What policy is used to define what channels are Multicast and what are Unicast? Should it be reached with a static policy or a dynamic policy in real time? If this assignment is adaptive, must the analysis of the demand be done in real time? Must we apply machine-learning technologies? Do client behaviors change significantly from one service area to another? Through an updated analysis of the behavior of video subscribers and the incorporation of machine-learning (ML) technologies, our work is aimed at finding the answers to the above-mentioned questions.

There are several works related to the video subscribers' behavior in HFC (Hybrid Fiber Coaxial) networks. In most cases they have been made by the vendors with samples of some operators – Cable Labs in 2009 and 2012 as well [3].– Our analysis includes the behavior of Legacy STB (Set Top Box), Hybrid STB (Video QAM and Control IP) and the behavior of our OTT (Over The Top) subscribers –the latter are part of our service called “FLOW”.– In this paper, we also include the behavior in major events, such as the 2014 and 2018 FIFA World Cup.

2. Motivation and Backgrounds

Telecom Argentina (former Cablevisión Argentina) has already moved from legacy Digital TV (DTV) to Hybrid (DTV+IP) and OTT system and now, we are finally starting to deploy Full IP Video delivery. Our biggest challenge in migration to full IP video is to deliver fully managed linear TV services to any device. The primary motivation for this migration to be based on IP Multicast is the expected improvement in efficiency over Unicast.

IP Multicast WG defines *Best Practices as the techniques that the working group has identified as generally being the preferred design approach in a specific area*. In this work, we seek to see how we can apply these best practices in light of the analysis of our clients behavior.

2.1. Definitions

The IP Multicast CableLabs Working Group suggests multicast live linear TV as the best practice and identified three main approaches to determine what content should be delivered using Multicast:

- *Viewership Driven Multicast*: any stream with more than one consumer will be multicast regardless of bit rate.
- *Policy Driven Multicast*: n configured channels are available for request via multicast (typically, these are the n most popular channels for a given time period and location)
- *Hybrids*: There are hybrids between the two previous models, the two possible ones that the working group would like to highlight are:
 - *Viewership Driven with Maximum Number of Multicast Channels*: the set of multicast channels at any given time is driven by *real-time requests* for content. However, like Policy-Driven multicast, there is a maximum number of channels allowed to be multicast.
 - *Viewership Driven with Limited Bit Rates*: This hybrid model adds to the pure Viewership Driven model a policy component that limits the number of bit rates which are available for multicast. Typically, in this model, bit rates are limited to HD-only or HD- and SD-only.

2.2. Background

Maximizing efficiency was the motivation for the development of IP Multicast. This efficiency is directly related to our video subscribers' behavior. This means that we must determine what the most popular channels are, and those will be the ones delivered using Multicast. The Pareto principle –or the 80-20 rule– is often referred to when describing video popularity and the concentration of user interest towards a few popular programs [2] [3].

Many authors have adjusted this popularity following a Zipf distribution [4], and based on that, they have determined the gain of using Multicast in the most popular channels. The distribution is as follows:

$$P_i = \frac{1}{\sum_{i=0}^N i^{(1-\alpha)}}$$

Where α is the skew factor and i is the rank.

Multicast gain is a measure of the efficiency of multicast delivery compared to unicast. The multicast gain achieved depends on a variety of factors, especially, the number of viewers per service group and the popularity of the programming.

With $\alpha = -1$, there are just a few very popular channels at a particular time and the potential for high Multicast gain $\gg 8$ [5]. Multicast gain of 8 indicates that the Unicast approach requires 8X the numbers of streams.

If $\alpha = 0.5$, we will have more popular channels at a particular time and potential for low Multicast gain $\gg 3$. For example, in [6] it was reported a gain of 5 under certain SG size conditions, popularity, etc.

Through this example, we want to illustrate in a simple way how the skew factor influences the gain; having a *long tail* and a *tall head* in the distribution. The tall head during prime time – observed in [6], for instance–corresponds to 60% of viewers watching the top 10 channels.

Zipf-Mandelbrot is the most appropriate model to replicate video popularity distributions –as presented in [7] and subsequent work [8].–

3. Systems Overview and Data Description

Telecom Argentina S.A. provides Live TV (or linear TV) and VoD (Video On Demand) services over two systems: Flow and DTV. There are about 500 Live channels and over 50,000 videos available. The users of both systems pay a monthly subscription fee to use Live TV and VoD services and they have to pay extra fees for some VoD contents. There are many differences between the systems; by way of example, Flow has functions as Catch up TV, Restart TV and NDVR (Network Digital Video Recorder), while Legacy platform has a TV guide where users can choose a Live channel or search for a specific VoD content by browsing into a couple of folders.

3.1. Data Description

In order to analyze the TV user behavior, we collected a large amount of logs from the two platforms from July 2017 to July 2018 and then selected particular days and weeks to conduct our study. There are about 3 million subscribers, taking into account STB –Legacy and Hybrid–and OTT devices, and the average number of daily records is about 55 million, so the sample that has been chosen is representative of general TV system users.

The logs contain many fields and those differ according to the type of system. Table 1 shows the format of Legacy system logs and Table 2 shows the format of Flow logs.

Table 1-DTV (Legacy) log sample

Fields	Examples
Date	06/26/2018
Hour	00:00.0
IP Address	10.132.34.53
Flag	w
Set Top ID	0004c96740
Service ID	788
Channel Number	4612
Time	61
Idle	61
Data	<i>Telediario 10 minutos</i>
Region	SANTA_FE_8

Table 2-Flow log sample

Fields	Examples
Account ID	3101671
Customer ID	788840
Device ID	3632563
Type	PHONE
OS Type	ANDROID
OS Version	7
Brand	SAMSUNG
Model	SM-G610M
Firmware	1.10.1-173531
Channel ID	277
Channel Name	DISNEY XD
Program ID	MV00000000153771
Program Title	<i>Un gran dinosaurio</i>
Quality	SD

Tunein	25/05/2018 09:28
Tuneout	25/05/2018 09:30
Duration	17

4. TV User Behavior Analysis

In this section, we explain some of the analysis we carried out related to TV user behavior from our Flow and DTV systems. As we are planning the migration to a Full IP Video platform, we focus our attention on Multicast gain at CDN (content delivery network) and SG (service group) levels. In order to estimate the impact on the CDN and SG sizing, we studied some parameters, described as follows:

- *Concurrence*. Number or percentage of STB or OTT devices using a service at the same time (day, hour, minute, etc.).
- *Access frequency*. Number of tunings of each channel or videos during a certain time window.
- *Type of requests*. It refers to Live TV or VoD.
- *Bitrate*. Streaming bitrate of Live channels or VoD videos (in bps). It depends on the quality of the contents, the quality of the channel and the type of device that reproduces the content.
- *Popularity*. Probability of tuning a certain channel or video. It is calculated as the number of tunings to this channel or video divided by the total tunings.
- *Busy hour*. Hour slot with the greatest concurrence in all day, which generally happens from 9 p.m. to 10 p.m. or from 10 p.m. to 11 p.m. It differs from the US's prime time, because in Argentina people tend to have dinner after 8 p.m.

For Legacy STB, we analyzed access frequency and concurrence by hour, and for Flow STB, only concurrence by 10 minutes. The reason why we studied different indicators for each system is the structure and complexity of each log. Calculating the concurrence for Flow is quite simple through an elaborated algorithm but it is not possible to reproduce the same algorithm for Legacy system. Therefore, for Legacy STB we calculated access frequency that is the most similar indicator to the concurrence in the lowest time possible that is an hour.

4.1. Comparison on Live TV and VoD User Behavior

To understand user behavior and traffic of both systems, we show in Figure 1 the concurrence of OTT devices on Flow ([a]) and the proportion of them in Live TV and VoD ([b]). In Figure 2, we represent the Live TV and VoD tunings, expressed as the percentage of total STB on the Legacy system([a]) and the proportion of STB that were streaming Live TV and VoD from 8 p.m. to 11 p.m. ([b]).

We conducted a weekly analysis to study concurrence in both cases and we found clear differences between the systems, mainly due to the granularity of data and the types of devices analyzed in each case. For the week represented in Figures 1 and 2, we observed that Legacy STB follows the

typical pattern of access frequency all the week while OTT devices present an irregular pattern, but they all have in common a peak in the middle of the week, which appears very pronounced in the Flow system, due to a soccer match.

Then, we carried out a daily analysis to study the proportion of Live TV and VoD tunings –we show the soccer match day and the previous day. – In both cases, the proportion of Live TV tuning is higher than VoD. For Flow, VoD tunings are about 20% and in the Legacy system they represent less than 1%. This result is in line with the configuration of the platform, which is able to support up to 12K simultaneous VoD tunings, which is 1% of all active Legacy STB.




Figure 1-[a] Concurrence of OTT devices from Flow, colored by type of request. [b] Proportion of Live TV and VoD tunings.




Figure 2-[a] Distribution of requests from STB of the Legacy system. [b] Proportion of Live TV and VoD tunings on March 14, 2018 from 8 p.m. to EOD.

One of the questions we have to answer is if it is necessary a multicast configuration for some channels and if so, how many channels are needed to be configured as multicast. We founded through the analysis that live contents are what users tend to view the most on the STB. Figures 1 and 2 show that while VoD represents around 20% of the views in OTT, in STB the same

percentage is around 1%. In [9] it was observed that more than 80% of viewers were found to be watching live TV between 7 p.m. and 9 p.m.

As Multicast is the best practice for Live TV, we conclude that it is necessary to implement it.

In the next sections, we make a deeper analysis of Live TV users behavior from both systems. We focus our attention on a day with a major event and on a regular day. The results of this analysis answer most of the above-mentioned questions.

4.2. Regular Weekday

In order to understand if the users behavior changes depending on the day of the week or when some particular event happens, we performed a study during regular weekdays –‘regular’ means a day with no soccer match or any other major event. We selected May 24, 2018 as an example.




Figure 3-[a] Concurrence of Chromecast devices. [b] Concurrence of other OTT devices, on Flow, every 10 minutes. May 24, 2018.

We can see in Figure 3 that the subscribers who accessed the live contents via Chromecast tended to choose sports channels –such as *Fox Sports*– or general interest channels –such as *Telefe*, *America TV*, *Arter*, among others. On the other hand, when we look at the rest of the OTT devices, we observe a substantial difference in users behavior, as the views are distributed among many channels.




Figure 4 - [a] Hourly access frequency to Live TV in Legacy system, on May 24, 2018. [b] Legacy STB playing Live TV channels simultaneously, on May24, 2018 from 8 p.m. to EOD.

The conclusions we get from the observation of tunings on the Legacy platform are similar to the ones driven from the Chromecast case. It is clear from Figure 4 that the majority of the requests go towards the sports channel, *Fox Sports*, and the general interest and news channels *America 2*, *Telefe*, *El Trece* and *TN*. It is important to mention that on that night, *Fox Sports* was transmitting a Spanish soccer match from 8 p.m. to 11 p.m., which is not a major event in Argentina, but still gets many viewers.




Figure 5-[a] Concurrent tunings from Legacy STB. [b] Concurrent tunings from Flow STB. May 24, 2018 between 10 p.m. and 11 p.m. (busy hour).

Figure 5 shows that there is not much difference among the percentage of views that the top channels get on a regular day. One channel concentrates 10% of the views on Legacy and 12% on Flow. It is followed by a set of four channels that get between 9% and 6% of the views. On the

legacy system, we observed that 10 channels get 50% of all simultaneous views, while on Flow this happens with eight channels.

We compared the distribution of the simultaneous views to several Zipf and Zipf-Mandelbrot theoretical distributions. The Zipf-Mandelbrot is:

$$p_i = \frac{C}{(i + b)^a}$$

Where p_i is the probability that a certain STB would tune in the i -th most popular channel, $a > 1$ and $b \geq 0$. The parameter C is a normalizing constant that depends on a and b :

$$C = b^{1-\frac{1}{a}} \cdot (a - 1)^{\frac{1}{a}}$$

We found that the one that approximates the most to the data observed is a Mandelbrot-Zipf with parameters $a=1.11$ and $b=0.56$, which is shown in Figure 6.




Figure 6 - Comparison of observed distribution to Zipf-Mandelbrot.

This agrees with the results obtained in related works, and shows that under normal circumstances a few channels –10 or less– concentrate a high percentage of the total views. We proceed to show how this distribution is affected when a major event occurs.

4.3. Major Events

In Argentina, soccer matches really drive TV usage and can introduce several variations to the channel ranking. In order to investigate the impact of these sports events, we analyzed two world cups and other important sports events. We picked two dates: March 4, 2018 (a typical soccer Sunday) and March 14, 2018 (Argentinian Super Cup). This event faces the winning teams from

previous tournaments and, in the latest edition, it faced *Boca Juniors* and *River Plate*, the two soccer teams with the most fans.




Figure 7 - [a] Concurrence of Chromecast devices. [b] Concurrence of other OTT devices of Flow, every 10 minutes. March 4, 2018.

Figure 7 shows how Flow Live TV consumption occurred via Chromecast and other OTT devices (phones, tablets and computers) on a typical Sunday with matches. The most visited channels were *TNT Sports* and *Fox Sports Premium*, around the prime time. The other two channels that stand out are *ESPN* and *TNT Series*. Therefore, three out of the four signals that accumulate the most views are sports channels.




Figure 8 - [a] Hourly access frequency to Live TV in Legacy system, on March 4, 2018. [b] Legacy STB playing Live TV channels simultaneously, on March 4, 2018 from 8 p.m. to EOD (end of the day).

With regard to the Legacy system, Figure 8 leads us to reach similar conclusions: the channels that were tuned the most around the prime time are *TNT Sports* and *Fox Sports Premium*, and during the afternoon, *ESPN*. It is easy to see in the second chart that during the night there are some other series (*TNT*), news (*TN*) and general interest (*America 2* and *El Trece*) channels that concentrate views.




Figure 9 - Concurrent tunings from Legacy STB on March 14, 2018 between 9 p.m. and 10 p.m. (match hour).

On March 14, the match started at 9.10 p.m., so we analyzed the concurrence around that time. According to the Legacy system data, showed in Figure 9, the most visited channel during that time slot is *Fox Sports Premium*. It consolidates 125K simultaneous views, which represents around 15% of the STB. Meanwhile, the second most viewed channel is *El Trece* –general interest, – with an 8% concurrence.

We would like to highlight the difference with respect to the regular day ranking. In this case, the first channel doubles the second channel’s concurrence. Besides, seven channels concentrate 50% of the views. It should be taken into account that on a regular day, around 10 channels get 50%.

To sum up, we conclude that major events not only introduce a variation in the channels that appear on the top of the ranking, but also modify the distribution of the tunings during a certain time interval.

4.4. Variation of Rankings in Time

To explore the variations of the rankings through time, we calculated each channel’s popularity in one-hour intervals, and established a daily ranking based on the maximum popularity achieved by each channel in one hour. Then, to compare the rankings from different dates, we used the Spearman’s rank correlation coefficient [10]. Figure 10 [a] shows the correlations obtained after comparing the top 10 channels on July 1, 2017 versus the top 10 calculated for the following 180 days. This data corresponds to the Legacy system.

It is easy to notice that during the first 40 days the correlation is high, indicating that the list of the top 10 generally consists of the same channels. After that, a series of soccer matches and TV series appear on screen. Correlation is lower when we compare July 1, 2017 to soccer days or when a popular series is first aired.




Figure 10 - [a] Correlation between the top 10 channels on July 1, 2017 and the top 10 on the 180 following days.[b] Correlations between the top 10, top 20 and top 30 channels, on July 1, 2017 versus the same rankings on the following 180 days.

In Figure 10 [b], we compared the correlation series when we use the top 10, top 20 and top 30 channels to calculate it. The top 10 series is the one with the highest variation, meaning it is more sensible to changes in the ranking. Nevertheless, the daily top 10 is similar throughout the days, even when an important series is transmitted. What introduces more changes is the transmission of a major sports event.

From the user behavior analysis, we conclude that it would be enough to set by multicast a limited set of signals, consisting of the most popular general interest, news, movies and sports channels. According to the Legacy data, there are 11 general interest and news, 6 movies and 8 sports channels that regularly appear among the top 10.

After that, we continued to analyze the gain, in terms of capacity, that multicast implementation would bring. This analysis aims at looking for an optimum number of channels that should be delivered using multicast.

5. Multicast Gain

We evaluated the multicast gain at CDN and at service group levels as a percentage of the capacity needed under a 100% Unicast scheme, which we define as follows:

$$Capacity\ 100\% \ Unicast = \sum_{\substack{All \\ channels}} Concurrence \cdot Avg\ bitrate$$

When working with data from the Legacy system, we counted on the access frequency to approximate the concurrence, and we assumed that the average (Avg.) bitrate is 4 Mbps.

We analyzed three scenarios: in the first one, called *Top 10*, the 10 most popular channels are delivered to Multicast, and the rest remain Unicast. For the second and third, named *Top 20* and *Top 30*, we set the 20 and 30 most popular channels to Multicast and the rest of the grid in Unicast.

The capacity that we would need at CDN level is calculated as follows:

$$Capacity\ Top\ "X"\ Scenario = \sum_{\substack{Top\ "X" \\ channels}} Avg\ bitrate + \sum_{\substack{Other \\ channels}} Concurrence \cdot Avg\ bitrate$$

Finally, we defined the multicast gain as:

$$Multicast\ gain\ for\ Top\ "X"\ Scenario = \frac{Capacity\ 100\% \ Unicast - Capacity\ Top\ "X"\ Scenario}{Capacity\ 100\% \ Unicast}$$

To calculate the gain at CDN level, we used the total concurrence –or total access frequency– for each channel. On the other hand, to calculate the gain at service group level we used the concurrence observed within the service group.

In addition, to estimate the maximum possible multicast gain, we proposed a theoretical scenario in which all the channels are transmitted via multicast. This is useful to determine whether the gain in other scenarios is close to the maximum or not.

$$Capacity\ 100\% \ Multicast = \sum_{\substack{All \\ channels}} Avg\ bitrate$$

So, the maximum gain is estimated as:

$$Maximum\ Multicast\ gain = \frac{Capacity\ 100\% \ Unicast - Capacity\ 100\% \ Multicast}{Capacity\ 100\% \ Unicast}$$

We would like to highlight that all the results in this section are based on the Legacy data. As seen in the previous section, the VoD views represent less than 1% of all views, so the results may differ in case the distribution of views between VoD and Live is other than 1%-99%.

5.1. Analysis at CDN Level

In order to address the question of how many signals should be delivered using Multicast, we analyze how multicast gain varies according to the scenarios and time in this section. Figure 10 shows the variation of the gain at CDN level under the three scenarios, during one week, from May 20 to May 28, 2018. We observed, as expected, that gain is greater at peak times, for all scenarios.




Figure 11 - Multicast gain, as a percentage of the capacity needed with 100% unicast scheme.

In Figure 11, it should be noted that the Top 30 line overlaps with the maximum. It seems that there is not much difference when there are 20, 30 or even if all channels are delivered using multicast. The difference between the Top 10 scenario and the maximum is, on average, 6% with a standard deviation of 2.8%.

It is intuitive that when a few channels are delivered using multicast, there is gain, but then if we add more channels, after a certain point the gain does not suffer a drastic increment. It is our task to find out where that cutoff is. In order to do that, we plotted the gain by the number of channels sent via multicast, as if each hour slot was a new sample. It is clear from Figure 12 that such cutoff should be between 10 and 30.




Figure 12 - Multicast gain versus number of channels that are set to multicast. Based on data from May 20 to May 28, 2018 gain calculated for hour slots from 8 p.m. to midnight.

If 25 channels were delivered using multicast, as we proposed at the end of section 4, the multicast gain at CDN level would be near its maximum.

5.2. Analysis at Service Group Level

We know that at service group level the gain is subject to the service group size. On the other hand, service areas tend to be smaller in time. Therefore, it is a frequent question whether there is multicast gain at this level. In this section, we tried to find the answer.

Figure 13 shows the distribution of Telecom Argentina S.A. service groups' size, in terms of households passed (HHP), by region ([a]), and the relationship between the SG size and the maximum possible gain under a 100% multicast scheme ([b]). The STB count per service group is higher in Buenos Aires than in other regions, because this is a highly populated area. It is clear from the scatter plot in Figure 13 [b] that the relationship between the gain and the service group size has a logarithmic shape.




Figure 13 - [a] Distribution of service group's size (HHP) by region.[b] Maximum multicast gain versus service group size.

In Figure 13, we plotted the maximum multicast gain at service group level, and in [a] we colored each line according to the service group size. It should be notice that plot [a] is not an area plot, it just contains so many series that it looks like one. The warmest colors, that are the biggest service groups' series, indicate higher maximum possible gain. Chart [b] summarizes chart [a] by showing the minimum, maximum and average gain at service group level.

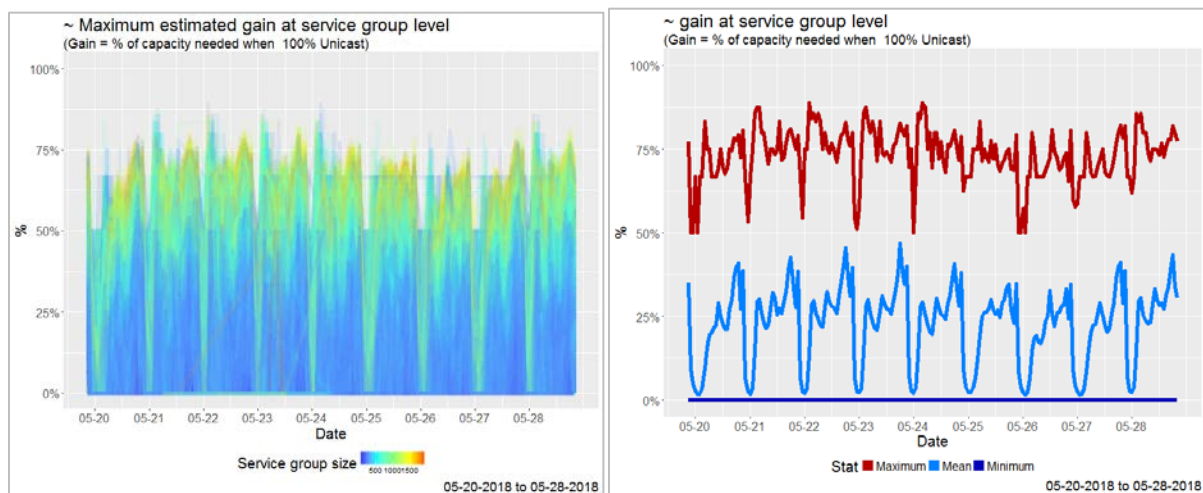


Figure 14 - [a] Maximum multicast gain at service group level, colored by service group size. [b] Mean multicast gain at service group level.

The gain may also be determined by the region where the service group is located. We know that channels' popularity tends to vary according to the region. By way of example, Figure 14 compares the popularity in Buenos Aires versus Córdoba ([a]) and the popularity in Buenos Aires versus La Plata ([b]). When a point is near the diagonal line, this means that the channel is as popular in the other region as it is in Buenos Aires. Otherwise, when a point gets far from the diagonal and near the edges, this means that the channel is very popular in one place but very unpopular in the other.

Local versions of the news and general interest channels tend to be popular within the regions where they are from. In the case of May 24, 2018 data, there are some news channels that are popular in Córdoba and not so much in Buenos Aires, and the sports channels are popular at both locations. Then, in La Plata there are more coincidences –most of the higher popularity points are near the diagonal.

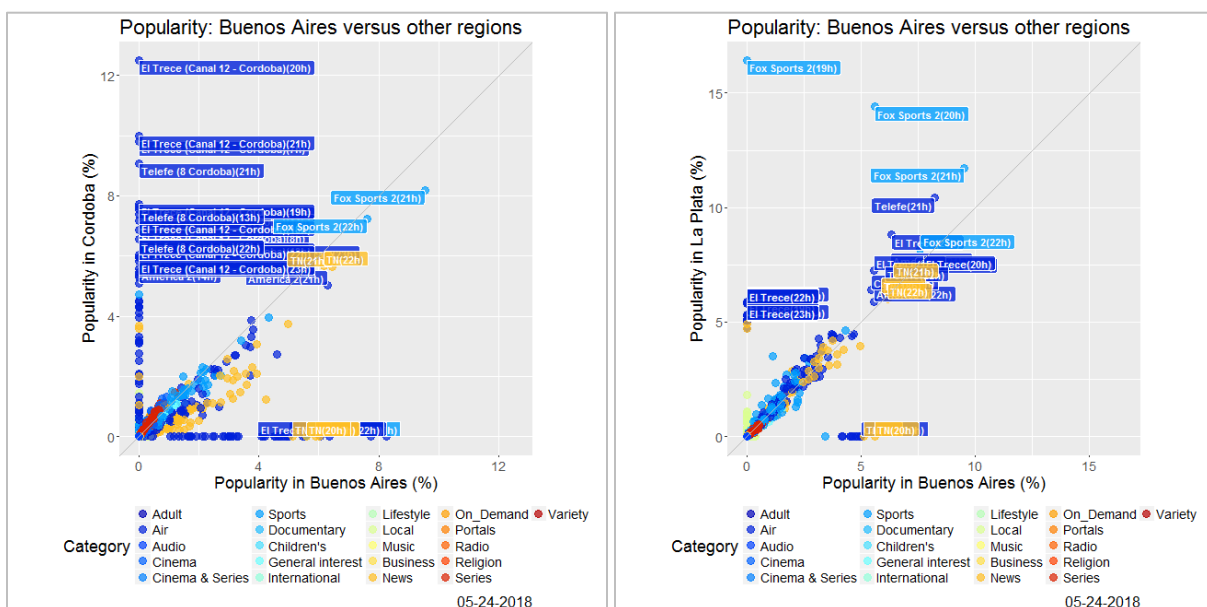


Figure 15 - Popularity in Buenos Aires region versus other regions, on May 24, 2018.[a] Buenos Aires versus Córdoba. [b] Buenos Aires versus La Plata.

To get an idea of multicast gain under the different scenarios, Figure 15 shows the average during the week from May 20 to May 28, 2018. The *Top 30* series overlaps with the *Maximum - 100% multicast* series, indicating that when there are 30 channels delivered using Multicast it may not make a difference to continue to add more signals.




Figure 16 - Average multicast gain at service group level for different scenarios.

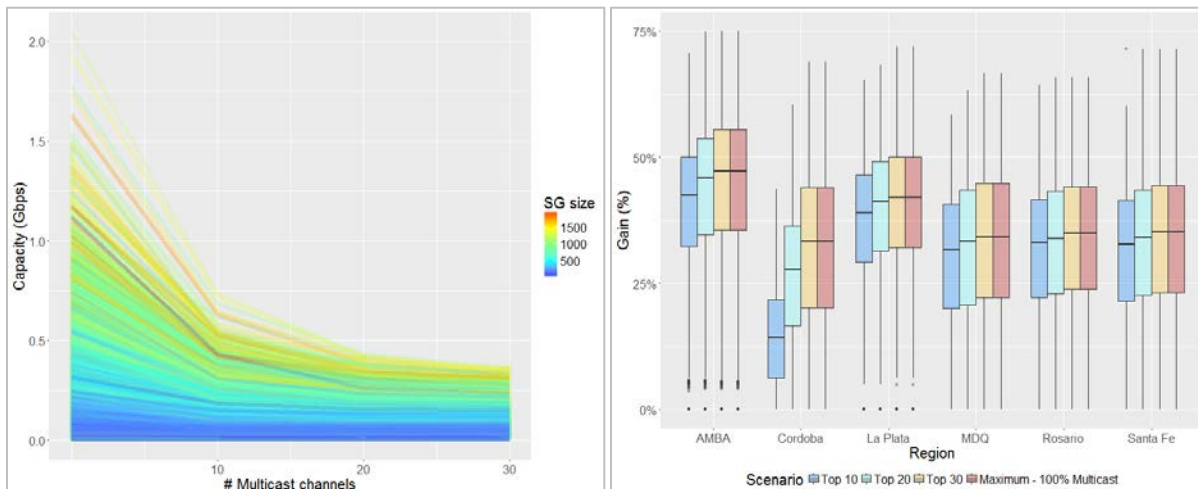


Figure 17-[a] Capacity needed at service group level versus multicast channels count, by SG size. [b] Multicast gain distribution by region and scenario on May 23, 2018 from 9 p.m. to 10 p.m.

From Figure 17 we conclude that when there are more HHP in the SG, there is a need for more capacity to support their activity, but also the decrease after setting channels by multicast is drastic.

In addition, the region variable seems to have a greater impact on gain when the region is Córdoba and there are fewer channels delivered using multicast.

To summarize the relationship between gain, region and area size, we estimated a statistical model [11] that has the general form:

$$Capacity = \exp(\alpha + SG_i + \beta \cdot X_i + \gamma \cdot \log HHP_i + \delta \cdot X_i \cdot \log HHP_i + \varepsilon_i)$$

Where HHP_i represents the HHP count in the i -th service group, it should be noted that we used the logarithm to denote the relationship described in Figure 12 [b]. The variable $T10_i$ should be replaced by 1 for the Top 10 scenario –and 0 in other cases–, $T20_i$ should be replaced by 1 for the Top 20 scenario and $T30_i$ should be replaced by 1 for the Top 30 case. The term SG_i is a random effect, which varies from one service group to the other, and the random error ε_i . This linear mixed effects model is subject to the assumptions of Gaussian distribution of the error and the SG_i random effect, in other words:

$$\varepsilon_i \sim N(0, \Sigma)$$

$$SG_i \sim N(0, \sigma_{SG}^2)$$

What is actually fruitful to get from the model is the parameter estimation, which is shown in Table 3. According to this information, for a SG of size 500 located in Buenos Aires, under 100% unicast scenario, the capacity needed would be of 210 Mbps. For the Top 10 case, we would need 164 Mbps (46 Mbps less) to sustain all of the STB's activity. With the Top 20 scheme, the saving is on average 83 Mbps, which means that if we add ten more channels, we can save an extra 37 Mbps. Following the same logic, when we add ten more (totaling 30 channels) we reduce the need, on average, by an extra 28 Mbps. It should be noted how the marginal difference between one scheme and other declines when the total count of channels delivered using Multicast increases.

Table 3- Fixed parameter estimation for the mixed-effects model

	Buenos Aires	Córdoba	Other
α	1.66	0.50	1.96
β	0.02	0.02	0.02
γ	0.59	0.74	0.50
δ	-0.01	-0.01	-0.01

For the case of a SG with the same size (500 HHP) but located in Córdoba, when passing from a 100% unicast scheme to the Top 10, the capacity that we would need would be reduced by 29 Mbps (approximately a 60% of what we observed for the Buenos Aires area). The Top 20 plan would sum another 24 Mbps to the saving –less than what we estimated for Buenos Aires. The Top 30 scenario would add another 20 Mbps.

In the Córdoba region, the percentage of gain when adding more multicast channels does not decrease as fast as in Buenos Aires. It is a subtle difference, but the explanation for it resides in the fact that Córdoba is a big region yet a very different one in terms of habits. The most viewed

channels may differ to the ones that are popular elsewhere but may still appear in the top 10, 20 or 30.

For the other regions, the interpretation is quite similar to Buenos Aires, except that on average, the impact of multicast implementation would be slightly reduced. Going from all-unicast to the *Top 10* scenario, would reduce the demand on average about 31 Mbps. Adding ten more channels, with the *Top 20* scenario, would add another 25 Mbps to the saving. Then, with the *Top 30* case, this would increase by another 20 Mbps. More details about this example can be found in Table 4.

Table 4–Estimation of the capacity (Mbps) for a 500 HHP service group, by multicast channel count and region.

Multicast channels\Location	Buenos Aires	Córdoba	Other
0	210	163	159
10	164(-22%)	134(-18%)	128(-19%)
20	127(-17%)	110(-15%)	103(-16%)
30	99(-13%)	90(-12%)	83(-13%)

Table 5 and Table 6 show capacity estimations for service groups of size 128 and 64, respectively. Throughout the examples, it is clear that when the SG is smaller, the gain –as a percentage– decreases. In addition, Córdoba is increasingly different from the rest of the regions. It should be noted that the examples in Table 4, Table 5 and Table 6 correspond to the blue lines in Figure 17 [a].

Table 5 - Estimation of the capacity (Mbps) for a 128 HHP service group, by multicast channel count and region.

Multicast channels\Location	Buenos Aires	Córdoba	Other
0	94	60	80
10	81 (-14%)	54 (-10%)	70 (-13%)
20	70 (-12%)	49 (-9%)	61 (-11%)
30	60 (-10%)	44 (-8%)	54 (-10%)

Table 6 - Estimation of the capacity (Mbps) for a 64 HHP service group, by multicast channel count and region.

Multicast channels\Location	Buenos Aires	Córdoba	Other
0	62	36	57
10	56(-9%)	34(-5%)	52(-9%)
20	51(-8%)	32(-5%)	47(-8%)
30	46(-8%)	30(-5%)	43(-7%)

The percentage of gain in the 128 HHP service group in Córdoba is practically the same as in the 64 HHP in Buenos Aires. Therefore, we conclude that not only the service group size but also its location determines the multicast gain. In order to boost the gain to its maximum, the regional channels –especially the ones that are popular in Córdoba– should be considered.

6. Real-Time Analytics

According to Gartner’s definition: “**Real-time analytics** is the discipline that applies logic and mathematics to data to provide insights for making better decisions quickly. For some use cases, real time simply means the analytics is completed within a few seconds or minutes after the arrival of new data. **On-demand real-time analytics** waits for users or systems to request a query and then delivers the analytic results. **Continuous real-time analytics** is more proactive and alerts users or triggers responses as events happen”.

We observed certain consistency in the rankings through time, we identified the sports channels and we know they should be delivered using multicast as the best practice. In addition, we believe it is convenient to consider the channels that are popular in other regions –especially Córdoba– and that may not appear on top of the general ranking. Therefore, it seems to be a better strategy to set a policy driven multicast approach and use the real time analytics to monitor its functioning.

In this section, we propose the idea of continuous real-time analytics in order to create alerts related to linear TV channels that are not delivered using multicast but in some cases –due to a major event– they would need to be. In order to do so, we looked for an unsupervised machine-learning algorithm to detect changes in user behavior, so that it allows us to make a decision about how to adjust the multicast plan.

6.1. K-means Clustering

Cluster analysis is a concept that encompasses a variety of machine learning techniques that aim to group a set of units or objects so that the ones within the same cluster have similar characteristics and the clusters are as different as possible from one another. This technique is unsupervised, which means that data is not previously labeled; it is used to find hidden underlying structure in the data.

One of the most well-known clustering algorithm is *k-means*. It is very useful to execute exploratory analysis on a large number (millions) of cases, when we want to classify them but the classes are unknown a priori. It has been applied in the past on pattern recognition, image analysis, data compression, among others.

A good cluster analysis has two main characteristics:

- *Efficient*: uses as few clusters as possible.
- *Effective*: captures all statistically and commercially important clusters.

The k-means method seeks to minimize the distances between the observations in the same cluster, and maximize the distances to observations in other clusters [12][13].

The algorithm consists on the following steps:

1. Place K points into the space determined by the variables measured on the units that we want to cluster. These points represent initial group centroids.
2. Assign each unit to the group that has the closest centroid.
3. After assigning all units, recalculate the positions of the K centroids.
4. Repeat Steps 2 and 3 until the centroids no longer move.

There is a variety of definitions of the distance that can be used to execute a k-means analysis. In this paper, we use the Euclidean distance.

6.2. K-means Clustering applied to the selection of multicast channels

In this application case, we will look for two groups: one that contains the channels with higher access frequency –that should be delivered using multicast, – and the other with the rest of the channels –that should remain delivered using unicast.– This means that the parameter has to be $K=2$.

If we wanted to split the complete channel list into unicast, multicast and variable multicast, we would set $K=3$. Provided that we found that the top channels are most of the time the same, we do not want to have a variable section, and we only want to monitor that the top channels are delivered using multicast in all cases, so the results showed in this section were obtained for $K=2$.




Figure 18- [a] Size of the cluster that groups the high access frequency channels - multicast cluster- by date, colored by type of event. [b] Access frequency versus date, channels colored by cluster. Data from July 1, 2017 to December 31, 2017.

Figure 18 shows part of the results obtained after running k-means on six months of data. On the left, figure [a] shows that the sample size of the multicast cluster varies mainly between 4 and 9 channels. Major events do not influence the number of channels that should be set to multicast; this situation is expected since the algorithm only takes into account the access frequency. On the right, we show a daily detail. As it was previously stated, a few channels capture most of the views. To understand what channels they are, we expose a sample with Flow data and another with Legacy data, in Figure 18.




Figure 19-[a] K-means clustering applied to the views per channel by hour for OTT devices. [b] K-means clustering applied to the access frequency per channel by day for the Legacy system. Algorithm used to classify the signals between multicast and unicast. Blue dots represent multicast channels and red dots unicast.

Note in Figure 19 that the names of the channels grouped as multicast tend to repeat execution after execution. In [a], common labels are *Telefe*, *Canal 13*, *America 2*, *TN*, *TNT Sports* and *Fox Sports*. In [b], which is from a different date and the k-means was executed based on the daily ranking, we see some repeated channels: *Telefe*, *El Trece* (which is the same as *Canal 13*), *America 2*, *TN*, *Canal 9* and *Fox Sports*. This leads us to think that channels that concentrate the most views have little variations in time, and a continuous monitoring system would be good enough to keep the system on track.

After taking into consideration these observations, we suggest the following algorithm:

```
Every 10 minutes repeat:

    Calculate {Cluster_Unicast; Cluster_Multicast}

    If Cluster_Multicast not in Multicast_fixed then:

        PrintCluster_Multicast
```

The clusters should be based on the concurrence variable, calculated using the k-means method and $K=2$. We would like to clarify that this is a theoretical design and the implementation of the alerts represents a new challenge.

Conclusion

Through the analysis of user behavior, we have found that the proportion of Live TV versus VoD tunings, as well as the most viewed channels vary according to time slot, region and the device used by the subscribers. The percentage of VoD tunings is around 20% on the Flow platform, and below 1% on the Legacy system.

The distribution of concurrent views on regular days follows a Zipf-Mandelbrot distribution, as observed in related works. Major events modify the ranking, and tend to alter this distribution. When one of these events takes place, it increases the difference between the most tuned channels and the rest.

In a more general approach, we have used the Spearman's coefficient to study the relationship of rankings from different dates, for a six-month period. We have found that the correlation is, most of the time, high (above 50%), except when a sports event or a series is aired.

The following conclusions about multicast gain are based on the subscribers' behavior while using the STB. The variation of the device may introduce alterations.

After estimating the multicast gain at CDN level, we observed that when there are around 10 channels delivered using multicast, the gain is around 50% during the busy hour. If more channels are delivered using multicast, the marginal gain tends to decrease. We found that with 25 channels delivered using multicast, the gain approximately reaches its maximum.

We studied the relationship between multicast gain at SG level and its size. We found that, given the variation of channels' popularity among regions, the gain is not only conditional to the STB count but it also depends on its location. For a SG located in Córdoba, the marginal gain is lower than a similar-sized SG in Buenos Aires. Meanwhile, in other regions, the tendency is the same as in Buenos Aires but on a lower scale, the absolute capacity and gain are in all cases smaller due to smaller SG. We concluded that regional channels should be taken into consideration to boost the gain.

In order to explore the channel count that would be delivered using multicast if it was an automatic and unsupervised process, we executed the k-means algorithm on ranking data. After analyzing six months of data, we found that this technique grouped between 4 and 9 channels as the most popular.

Given the high correlation between rankings from different days, the sports channels being already identified and considered for multicast as the best practice, and the fact that it would be necessary a complex process –hence non-scalable in real time– to capture the regional specificities, it seems inconvenient to apply real time analytics for a viewership driven multicast approach.

Nevertheless, real time analytics provide an efficient alternative for monitoring a policy driven multicast approach, since there could be a special event not considered so far (not a sporting event, or a popular TV series) which could drastically shift the ranking for a few hours and then return to usual.

We proposed a continuous process, which consists of a k-means clustering algorithm to be executed every 10 minutes. The program looks for the channels that get the most views and checks whether they are included in the multicast channel list. In case there is one or there are more channels that are being accessed aggressively, and do not belong to the multicast list, it sends an alert and reports the list of the most popular channels.

Abbreviations

ABR	adaptive bitrate
avg	average
bps	bits per second
CDN	content delivery network
DOCSIS	data over cable service interface specification
DSL	digital subscriber line
FTTH	fiber to the home
HD	high definition
HFC	hybrid fiber coaxial
HHP	household passed
HTTP	hypertext transfer protocol
IP	Internet protocol
IPTV	Internet protocol television
ML	machine learning
NDVR	network digital video recorder
OTT	over the top
QAM	quadrature amplitude modulation
SD	standard definition
SG	service group
STB	set top box
VOD	video on demand

Bibliography & References

- [1] S. Deering, "Host Extensions for IP Multicasting," Network Working Groupb-RFC 1112 , Stanford University, August 1989.
- [2] Cable Television Laboratories, "IP Multicast Adaptive Bit Rate Architecture Technical Report," OC-TR-IP-MULTI-ARCH-C01-161026, October 26, 2016.
- [3] Ron Reuss, "IP Unicast v. Multicast Modeling Overview," CableLabs, Liousville, Colorado, September 2012.
- [4] Kunwadee Sripanidkulchai, Bruce Maggs, and Hui Zhang, "An analysis of live streaming workloads on the Internet," in *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement (IMC '04)*, ACM, New York, NY, USA, 2004.
- [5] Guillermo Wilkinson, "IP Video Topics," in *Seminario Internacional de Telecomunicaciones - SIT 2016*, Buenos Aires, March 2016.
- [6] J. Horrobin and G. Shah, "Pioneering IPTV in Cable Networks," in *SCTE Cable-Tec Expo*, October 2013.
- [7] Ulm and P. Maurer, "IP Video Guide - Avoiding Pot Holes on the Cable IPTV Highway," in *SCTE Cable-Tec Exp*, October 2009.
- [8] Weisenborn, Hildebrand J, "Video Popularity Metrics and Bubble Cache Eviction Algorithm Analysis," PhD thesis, University of Essex., <http://repository.essex.ac.uk/22350/>, 2018.
- [9] Amit Eshet, John Ulm, Uzi Cohen, Carol Ansley, "Multicast As A Mandatory Stepping Stone For An IP Video Service To The Big Screen," in *NCTA/SCTE Technical Sessions*, spring 2014.
- [10] S. Sprent and N. C. Smeeton, *Applied Nonparametric Statistical Methods*.--3rd edition., Chapman & Hall/CRC, 2001.
- [11] G. M. Fitzmaurice, N. M. Laird and J. H. Ware, *Applied Longitudinal Analysis*, Hoboken, Hudson, U.S.: J. Wiley & Sons, 2004.
- [12] J. A. Hartigan, *Clustering Algorithms*, [New Haven, Estados Unidos]: John Wiley & Sons, 1975.
- [13] D. Peña, *Análisis de datos multivariantes*, España: McGraw-Hill Interamericana de España S.L., 2002.

- [14] H. Yu, D. Zheng, B. Zhao, and W. Zheng, "Understanding User Behavior in Large-Scale Video-on-Demand Systems," in *In Proceedings of EuroSys2006*, Leuven, Belgium, 2006.
- [15] M. Cha, H. Kwak, P. Rodriguez, Y. Ahn, and S. Moon, "Analyzing the Video Popularity Characteristics of Large-Scale User Generated Content Systems," *IEEE Transactions on Networking*, vol. 17, p. 1357–1370, October 2009.

Running a Multi-Tenant Hybrid Cloud for Large Scale Cable Applications

A Technical Paper prepared for SCTE•ISBE by

Neill A. Kipp

Distinguished Engineer and Cloud Software Architect
Comcast

1401 Wynkoop Street, Suite 300

Denver, CO 80202

m: 720.530.6917

Neill_Kipp@cable.comcast.com

Table of Contents

Title	Page Number
Table of Contents	2
Abstract	3
1. Dream of the Hybrid Cloud.....	3
2. Data Center Snowflakes	4
3. Use Case: Single Tenant Private Cloud	6
4. Lure of the Public Cloud.....	7
5. Hybrid Cloud Management Platform	9
6. Beyond Virtualization: Containers, Orchestration, Serverless	11
7. Use Cases of the Hybrid Cloud.....	12
8. Hybrid Cloud Community	13
9. Summarizing the Hybrid Cloud	14
Abbreviations	15
Bibliography & References.....	16

List of Figures

Title	Page Number
Figure 1 - The idealized “hybrid cloud” provides a runtime fabric of computing infrastructure, networking and storage.	4
Figure 2 - Private clouds come with a diversity of management issues.	6
Figure 3 - The IP linear application stack deploys thousands of components into the private cloud.	7
Figure 4 - Public clouds normalize runtime infrastructure.....	9
Figure 5 - The hybrid cloud portal lets users browse relevant documentation and manage tenancy, showback, and permissions.	10
Figure 6 - Containers virtualize the operating system with increased compute density.	11
Figure 7 - Kubernetes attaches to a cluster for management, but does not broker communication from application clients.	12
Figure 8 - During high usage, applications could auto-scale into the public cloud.	13

Abstract

The cloud—with its automation and virtualization of compute, network, and storage—has fundamentally changed the way software engineers design, develop, and deploy applications. Instead of by-hand configuration, software programs deploy virtual routers, firewalls, databases, and application servers. Software causes application servers to respond to load changes, change network topology, and scale deployments accordingly. As a result, applications can be effectively tested while they are running in production!

Before public cloud, cable companies invested in data centers and managed their own compute, storage, networks, and applications as their own private cloud. The public cloud alternative offers a compelling agility but often comes at an increased price. For the best of both worlds, cable companies can implement a hybrid cloud solution that leverages the existing capital investment in their private cloud infrastructure alongside the increased agility of the public cloud.

Running a hybrid cloud is challenging. The hybrid cloud must manage multiple tenants, each represented by multiple users. Users must be able to request cloud resources for their tenants in any private cloud region and from multiple public cloud vendors. All cloud deployments must secure video media, customer data, and application services.

The Comcast cloud team has built a hybrid cloud for large-scale cable applications. We publish hybrid cloud architectures and tools to help product owners upgrade their applications to be virtualized, containerized, and orchestrated. We provide user permissions, a network security framework, and automated controls that provide guardrails to streamline software development and minimize risks of security breaches. Our hybrid private cloud is running in eight regions and uses three public cloud providers. It hosts hundreds of tenants, thousands of users, and the software it hosts serves tens of millions of customers.


1. Dream of the Hybrid Cloud

Years ago, when server closets overflowed with gear, cable companies expanded into data centers to deploy their computing resources for television delivery and software applications. While a “data center” typically provides racks of servers, connectivity, power, cooling, and technical staff, a “cloud” extends these services to provide virtual servers, networks, and storage. It also includes software that lets application teams configure their own deployments.

Demand for computing has increased such that these on-premises data centers with virtualization (“private clouds”) are usually operating at maximum capacity. As an alternative to improving compute facilities in the private cloud with additional capital expenditures, companies are increasingly deploying their expanded operations into seemingly unbounded “public clouds” including those offered by Amazon Web Services, Microsoft Azure, and Google Cloud.

Private clouds retain certain perceived advantages, including physical security, operational control, capitalized infrastructure, and simple familiarity. As such, full migration to the public cloud may never be warranted. As a result, the natural evolution from private closet to public cloud has left the largest companies with an amalgamation of software deployments spread across networks, firewalls, hand-racked servers, and virtualized instances. Private infrastructure runs in well-known locations, is operated by esteemed colleagues, and the story of each data center contributes to a collective corporate history. Public infrastructure runs in undisclosed, regional locations (with nonspecific names like “US East”).

For the foreseeable future, it makes sense that companies will operate “hybrid cloud” environments—transient combinations of private and public virtualized services. Widely distributed, heterogeneous infrastructure will support a broad range of custom interconnected applications. Toward that mission, cloud operators are seeking to normalize the cloud technology substrate and realize their vision for the cloud: a homogeneous, distributed, secure, and fault-tolerant computing infrastructure, networking, and storage “fabric” for all application deployment needs (Figure 1).



Homogeneous Fault-tolerant Distributed Runtime Fabric

Figure 1 - The idealized “hybrid cloud” provides a runtime fabric of computing infrastructure, networking and storage.

2. Data Center Snowflakes

The private cloud, a key component of the hybrid cloud, comes with its own challenges. Data centers are not homogeneous; in practice, they are as unique as snowflakes. Application teams must manage their own resiliency and high availability. Quarters are usually cramped.

Comcast has deployed servers, network, and storage to more than 50 data centers. Historically, deployment teams reserve space in a specific data center, then purchase and ship servers, storage, and networking equipment to be racked there. Teams took latitude in their purchasing choices, and negotiated with individual operators in each data center to resolve specific conflicts with the local autonomy. As a result, each data center grew into its own operating “silo” of unique heterogeneous sprawl. This diversity has become challenging to operate, as the corporation must hire and retain operational expertise for each vendor and model of networking equipment, storage unit, and compute node.

A typical hardware deployment in a single facility requires months of lead time and an extensive approval process. The team orders hardware to be shipped and works with the operators to get the hardware racked, powered, and networked. The operators install the requested operating systems and provide specific network addresses per server. Before automation such as PXE boot and Foreman, servers were hand-configured by typing in values from spreadsheets. When the data center operations team completed its phase of work, application teams were left to verify and debug each hand-typed deployment. Inflexibly

numbered servers, by convention, became the “pets” of application teams, each server having a personality of its own.

Growth in these data centers sometimes required servers to be re-racked in rows close together so that network configuration could be normalized. Otherwise, routing became more complex and therefore subject to accidental misconfiguration. For example, an application deployment in the first year might require only one rack, but due to increased usage, may need additional racks in year two to satisfy demand. The operations team would choose between re-planning the footprint or reconfiguring its network routers. Sprawling router configuration is its own problem, where manual mistakes lead to long down times.

While some data centers are owned (and therefore operated) by the corporation, others are made available by vendors, and provide “remote hands” to do physical installations, maintenance, and upgrades. Service-level agreements vary by facility, and therefore increase the “snowflake” problem.

Software deployment automation similarly evolved. In the earliest days, an operator had to insert an installation disk into each node. On-site manual operations were quickly replaced by remote deployments such as secure shell (SSH) and secure copy (SCP), then automated by deployment tools such as Puppet or Chef. Even with automation, these agent-based deployments too often get “stuffed up” and fail to reach the desired state. Thus, the versions of deployed software in the tier of application servers could diverge and, as a result, the application service would be degraded.

Enter virtualization. With virtualization, each physical host server is configured with “hypervisor” hosting technology that allows multiple “virtual machine” guests to be deployed atop. A centralized console lets teams manage their virtual machines, networking, and storage attachments for each data center. Virtualization provides a clean division between hardware operations and software applications.

Ideally, virtualization consoles would span data center installations, and application teams could rely on a single console, command line, and programming interface to manage their national deployments. Virtualization consoles have only recently provided the feature to coalesce and manage multiple deployments as a homogeneous private cloud.

Even the virtualization technology substrate has diverged. Today, the Comcast private cloud supports virtualization technologies that include VMWare and OpenStack. Application teams must therefore pick a substrate technology or learn the console and programming interface for each virtualization system on which they wish to deploy.

One of the goals of application deployment is high availability. Single points of contact (SPOC) directly imply single points of failure (SPOF). Applications, servers, networking, power, air conditioning, must each be redundant. Multiple fiber optic connections must be provided for each rack, row, and data center—when a backhoe cuts one fiber connection, the redundant fiber must suffice.

For reliability, storage must also be redundant. Initially, redundant arrays of independent disks (RAID) technology provided this redundancy. Unfortunately, a rack, row, or whole data center might be partitioned from the rest of the network, rendering the entirety of compute and storage there unavailable. In today’s private cloud, individual application engineers must have the additional expertise to design for availability and reliability. Cloud storage is increasingly realized by “just a bunch of disks” (JBOD) and employ software that implements robust storage.

Ideally, cloud infrastructure configuration and subsequent software application deployments are developed “as code.” Infrastructure configuration tools such as Terraform can be launched using files that

are stored in familiar source code control systems such as Git. Applications can then be pushed to servers using tools such as Ansible. “Configuration as code” significantly reduces the complexity faced by application teams when deploying software, and as such has become a software engineering “best practice.”

Private clouds also suffer from being visibly finite and ultimately cramped.

If an application team was prescient enough to provision for growth, a majority of its physical resources would remain underutilized until demand met supply. If not, growth is crippled by delays in purchase, delivery, rack, and configuration.

With virtualization, teams are selfishly motivated to secure quota above current usage. Unmanaged “land grabs” constrains the ability for new teams to obtain resources. Worse yet, to accommodate requests, private cloud operators can allow over-commitment of virtualized resources (as much as 8:1 virtual CPU to CPU!) in the hopes that actual utilization will not simultaneously peak on each hypervisor. Unfortunately for cable companies, actual loads peak at the same times every day, every week, and during specific events like championship sports. When a virtual machine instance hogs the compute or network resources for the whole hypervisor, other instances must suffer the “noisy neighbor” problem. On the other hand, strict quota enforcement leads to stranded compute and underutilized bandwidth (Figure 2).




Figure 2 - Private clouds come with a diversity of management issues.

3. Use Case: Single Tenant Private Cloud

Some applications at Comcast are central to the business and continually serve nearly every subscriber, including IP Video on Demand (VOD), IP Linear Channel delivery, and Content Delivery Network (CDN). VOD, linear, and CDN each represent a single “tenant” in the private cloud, and because of their demand volume, each runs on dedicated infrastructure.

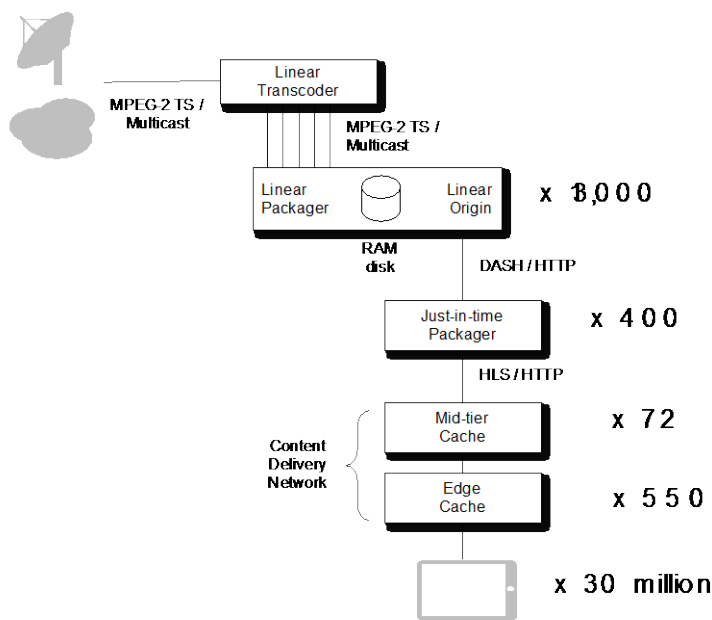
The IP VOD system runs on dedicated servers hosted in national data centers. End to end, the VOD library holds seventeen years of video content, with titles rotating continually. Network-attached video

storage is in the petabyte range. VOD server software for packaging and encryption run on virtual instances and are automatically deployed with Puppet.

Similarly, 13,000 linear channels are served from a combination of dedicated servers in national data centers and regional installations. The content delivery network protects these servers, and thus allows the software that packages and originates linear channels to run at a steady resource state.

Servers that implement our two tiers of content delivery network (CDN) run on hardware configured specifically for network and storage performance. For compute efficiency, CDN caching runs without a layer of virtualization. Comcast has deployed the second largest CDN in the country, delivering terabytes daily. We support CDN clusters in more than 50 data centers nationwide.

These full-stack deployments require dedicated operations teams to provision and maintain these mission-critical applications. The hybrid cloud must continue to support large, single-tenant deployments such as IP Linear Video both for Comcast and syndicated to other providers in the US and Canada (Figure 3).



IP Linear Single Tenant Use Case

Figure 3 - The IP linear application stack deploys thousands of components into the private cloud.

4. Lure of the Public Cloud

Often cable applications teams can code and deploy applications in less time than it takes to purchase and rack the servers these applications would be hosted on. Procurement need not risk private cloud augmentation for applications that might not be embraced by the market. Developers can approach a homogeneous infrastructure that closely resembles the cloud vision of a generic compute fabric. These statements are true, when using a public cloud.

In the public cloud, all resource reservation and allocation is done using consoles and automation interfaces. Each public cloud is a homogeneous entity spread across multiple regions. Within each region is a set of “availability zones” that provide a singular failure domain. Within each availability zone,

application teams are unable to learn exactly where their virtual machines are deployed, and if restarted, if they have been migrated. Even virtual server internal network addresses should be expected to change.

In the public cloud, everything is “software defined” and all configuration is stored “as code.” Thus, configuration can be automated by Terraform or similar proprietary public cloud software. Deployments that took months now can take hours and upgrades can happen in minutes.

Software defined compute means that configuration files ensure that enough virtual machines are running at all times; should one fail, it will be recycled and another will be online in minutes. Advanced deployments can combine load monitors with compute configurations—the public cloud can add virtual machines when loads spike upward and remove them as load reduces.

Software defined networking (SDN) means that routing is accomplished as follows: Software configures the domain name system (DNS), global server load balancing (GSLB), high availability load balancing (LB), virtual private clouds (VPCs), and peering between application stacks. And once a hardware router is installed between public and private clouds, software configures routes between public and private networks. Software configures intrusion detection systems and distributed denial of service mitigation systems.

Teams from many large companies, including cable companies, have begun to consider the public cloud, overcoming this often-repeated sentiment: “Our proprietary data should be stored within our premises.” With that consideration, significant security challenges must be overcome, as follows:

“Our proprietary
data should be
stored within our
premises.”

Networks must be secured. Similar to the risks with storage, access to servers within the public cloud must be secured and controlled. Inadvertently making firewall changes such that protected servers become public can lead to disastrous consequences.

Data must be secured. Storage “buckets” in the public cloud can be converted from protected to public with one click. While malicious intent is possible, much more likely is that someone on the application team inadvertently assigned incorrect configurations. Furthermore, data stored in buckets should be encrypted at rest, and the keys maintained separately.

Permissions must be managed. As a key part of access control, companies should apply a blanket DENY permission to any resource or permissions changes (especially global ones!) and carefully add ALLOW permissions for specific team members or services assuming specific roles to make specific changes to specific resources. This implements the “principle of least privilege,” where access is provided only to users (or pseudo-users) at the level of that entity’s legitimate purpose in the larger system. A dangerous shortcut in the public cloud skirts researching and applying specific permissions for specific resource types, instead giving blanket authority to “do anything, just get it done.” Furthermore, permissions should be managed “as code” as well, to ensure each account complies with human and automated audits.

Access must be controlled. Unless further restricted, permission to modify any resource within a cloud account is by username and password. Because teams share the account, it is far too easy to share the passwords to it. Advanced access protection includes the following:

- Requiring multi-factor access (MFA) on each login mitigates password sharing.
- Binding cloud account access to a corporate single sign-on (SSO) expires users when they change teams or leave the corporation.
- Monitoring and alerting on access, including times and locations of login, helps ensure a reasonable workplace context.
- Monitoring and alerting on permissions changes helps prevent breaches of data and network.
- Restricting root access to a select few trusted operators, ensuring passwords are extremely long and complex, rotating these passwords often, and storing these passwords in a “vault” that requires MFA and SSO to access significantly reduces the risk of accounts and data being accessed inappropriately.
- Implementing VPCs (micro-segmentation) and peering in the private cloud reduces the impact of any single breach and protects resources in the private cloud.

Costs must be kept within budget. With the public cloud, everything on the menu is available for a team to very quickly buy and deploy. Inefficient architectures or inadvertent deployments can lead to huge cost overruns. Public cloud costs should be continually monitored and tracked, and alerts sent before spending goes too far awry. For example, many subversive crypto-currency mining operations have been discovered and dismantled in the public cloud. In data ingress and egress can be free or very expensive. Extreme care must be taken to ensure the cost of video egress to subscribers is well managed (Figure 4).




Figure 4 - Public clouds normalize runtime infrastructure.

5. Hybrid Cloud Management Platform

Every cloud needs a console. The Comcast cloud engineering team is developing a “OneCloud Portal” with the following specific features for hybrid cloud users:

6. Beyond Virtualization: Containers, Orchestration, Serverless

Cable companies can quickly take advantage of recent advances in compute resource virtualization including containerization, orchestration, and serverless computing, discussed below.

Containerization. A “virtual machine” is an emulation of a server running with the assistance of a hypervisor application. Hypervisors can provide a compute density of approximately 10 virtual machines per host. A “container” such as those provided by Docker or Rkt technologies is a single process or small group of processes running directly on the underlying operating system, without the overhead of a hypervisor. A single host can run as many as 100 containers. Thus, the efficiency gained by containerization can increase a server’s compute density by an order of magnitude. Note that containerization does not improve network resource efficiency, and compute density might be constrained by the physical network capability of the underlying host (Figure 6).




Figure 6 - Containers virtualize the operating system with increased compute density.

Orchestration. Like an operating system schedules processes, an orchestration system such as Kubernetes or Apache Mesos schedules containers across multiple servers. In this way, developers can create robust, long-lived “microservices” that scale independently alongside short-lived “jobs” useful for big data analytics. Note that orchestration systems themselves require installation and maintenance. Orchestration systems manage multiple tenants and container versioning, and as a result, can be used to deploy, install, test, and roll back updates *in production environments* prior to full application launch (Figure 7).




Figure 7 - Kubernetes attaches to a cluster for management, but does not broker communication from application clients.

Serverless. “Functions as a service” occurs when the runtime infrastructure has grown in sophistication such that the unit of application deployment is not a virtual machine with installed software, nor a container, but the application function that would run each time the service inside the container was called. Functions as a service are not ideal for every workload. Such functions typically have a wide deviation for startup latency, must operate within strict timeout periods, must work within random access memory (RAM) constraints, and must be written in specific languages such as Python or Node.js.

Applications must be reworked significantly to enable the move from virtual machine deployment to containers. Server applications that were once deployed on monolithic server such as those running Enterprise Java on a load-balanced Tomcat cluster can now be rewritten into containers as Go microservices and deployed on a Kubernetes cluster. Launch times are faster. Testing can be done directly in the production environment, particularly using strategies such as “blue/green” deployments and “canary” testing. Well-designed containerized applications can be updated seamlessly in production and do not require maintenance windows.

Even more so than containers, applications must be reworked to use functions as a service. However, by writing functions that integrate sophisticated cloud “software as a service” including document databases, messaging systems, and machine learning libraries, application time to market can be dramatically reduced. Today functions as a service are provided by the major public cloud providers, yet once written are not directly reusable between cloud providers.

7. Use Cases of the Hybrid Cloud

The hybrid cloud provides a substrate for innovative benefits to application developers and users. Those benefits are outlined here.

Fault tolerance. Private cloud resources can be lost during single points of failure. Adding redundant storage, network, and servers in the public cloud can make application services more robust.

Spot markets. Application teams with less time-sensitive workloads could delay processing until a public cloud provides variable pricing windows or auctions services outright. With the right tooling, application teams could take advantage of spot market pricing for resources.

Cloud bursting. For cable applications, compute and network load attributable to subscriber activity comes predictably every evening such as for prime time, every week such as for weekend football, and for special events such as the Olympics. Ostensibly, the bounded private cloud resources could handle the baseline application load and effectively unbounded public cloud resources could be used exclusively for additional, transient load. Bursting load onto public cloud requires additional engineering, but promises to be an effective way to manage variability in resource demand.

For these hybrid cloud use cases to become a reality, however, each development team must specifically engineer resource configurations, create unique scale triggers, extend application deployment scripts to run on hybrid infrastructure, while monitoring deployments in multiple clouds. Internally developed and third-party tools will lead the way to supporting these and other hybrid cloud use cases (Figure 8).

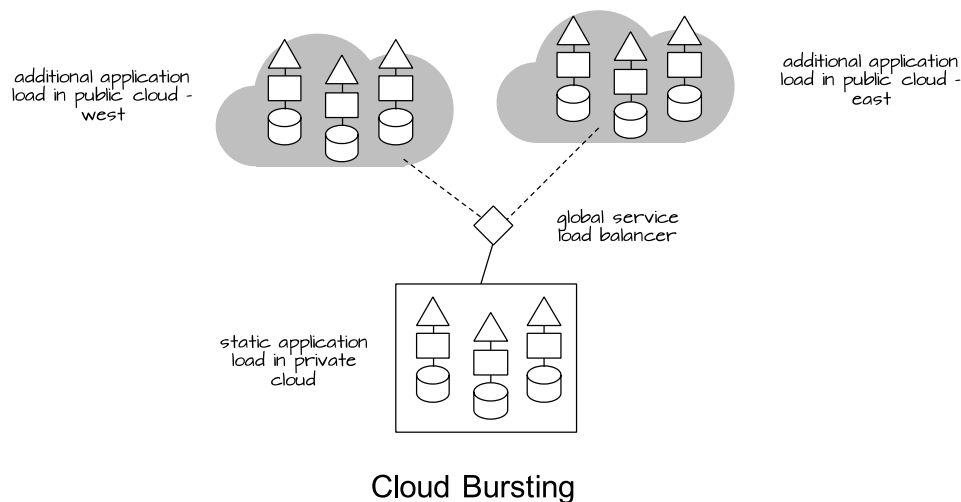


Figure 8 - During high usage, applications could auto-scale into the public cloud.

8. Hybrid Cloud Community

As with all popular technologies, conversations ensue and communities form. Here are some communication channels that we see building the communities that surround and support the hybrid cloud.

Documentation portal. Starting with the “home page” for the hybrid cloud, a documentation portal includes new feature announcements, governance deadlines, adoption instructions, case studies, and architectural design patterns. In the documentation portal, the entire community can read and contribute to the knowledge base for the hybrid cloud.

Email blasts. The cloud operations team regularly sends announcements to distribution lists concerning security governance, new feature announcements, and upcoming events.

Scheduled training. Not everyone is an expert in all aspects of the private cloud. Instructor-led and online training keeps engineers apprised of the newest and best things happening in the hybrid cloud. Training courses also inspire discussions and build bridges between teams.

Cloud Center of Excellence. Trained cloud architects, especially those with public cloud certifications, are the primary instructive resource for application teams throughout the enterprise.

Collaboration channels. Team-centric collaborative chat platforms such as Slack or Skype keep everyone up-to-date in the hybrid cloud discussion.

Cloud summit. Similar in form to an industry conference, a regularly scheduled corporate cloud summit provides a podium for leadership and technical content that coalesces discussion and provides a venue for far-flung, face-down teams to come together and share their experiences with the hybrid cloud. Catered lunches for attendees add a personal touch.

9. Summarizing the Hybrid Cloud

The hybrid cloud is a combination of private cloud and public cloud. A cable operator's private cloud has necessarily evolved and may have diverged between regions. Meanwhile the public cloud is purchased and used as a product that is largely consistent across regions. A cable-specific application stack that delivers IP linear video naturally deploys in the private cloud, provided enough physical infrastructure can be made available. For a cable company to embrace the public cloud, networks must be secured, permissions managed, access controlled, and costs must be kept within budget. Furthermore, companies will need a hybrid cloud software platform to manage tenants, accounts, showback, permissions, roles, and documentation.

Virtualization does not end with virtual machines. Teams are embracing new technologies for containerization, orchestration, and functions as a service, even though applications must be redesigned and rewritten for the cloud. Hybrid clouds provide innovative benefits to improve service or save money, especially using strategies including fault tolerance, spot markets, and cloud bursting. Hybrid clouds are fertile fields for communications and conversations, including documentation, news, training, and collaboration.

The Comcast hybrid cloud uses multiple private and public cloud technologies including OpenStack, VMWare, Amazon AWS, Microsoft Azure, and Google Cloud Platform. We run the multi-tenant portion of the private cloud in eight regions, and have single tenant deployments in more than 50 data centers. We have hundreds of tenants, thousands of users, and serve tens of millions of subscribers.

Abbreviations

CDN	Content delivery network
CPU	Central processing unit
DNS	Domain name system
GSLB	Global service load balancing
IP	Internet protocol
ISBE	International Society of Broadband Experts
JBOD	Just a bunch of disks
LB	Load balancer
MFA	Multi-factor authentication
RAID	Redundant array of independent disks
RAM	Random access memory
RBAC	Role based access control
SCP	Secure copy
SCTE	Society of Cable Telecommunications Engineers
SDN	Software defined networking
SSH	Secure shell
SSO	Single sign-on
SPOC	Single point of contact
SPOF	Single point of failure
vCPU	Virtual CPU
VOD	Video on demand
VPC	Virtual private cloud

Bibliography & References

Amazon Web Services, <https://aws.amazon.com/>

Ansible, <https://www.ansible.com/>

Apache Mesos, <http://mesos.apache.org/>

Docker, <https://www.docker.com/>

Git, <https://git-scm.com/>

Google Cloud, <https://cloud.google.com/>

Kipp, Neill, “A Highly Scalable Cloud Architecture for Delivering Linear IP Video,” SCTE Expo, Philadelphia, September 2016.

Kubernetes, <https://kubernetes.io/>

Microsoft Azure, <https://azure.microsoft.com/en-us/>

Node.js, <https://nodejs.org/en/>

OpenStack, <https://www.openstack.org/>

Puppet, <https://puppet.com/>

Python, <https://www.python.org/>

Rkt, <https://coreos.com/rkt/>

Terraform, <https://www.terraform.io/>

VMWare, <https://www.vmware.com/>

Securing a Hyper-Connected Society

A Technical Paper prepared for SCTE•ISBE by

Tom Conklin

Vice President of Business Development
Ericsson
6300 Legacy Blvd. Plano, Texas USA
+1 (703) 789-4574
Tom.conklin@ericsson.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
A Hyper-Connected Society.....	3
1. How did we get here? Agriculture as a case study	4
2. Changes in the way software and hardware products are designed	5
3. Intelligent Transport Systems.....	6
4. E-Health	7
5. Smart Grid	7
6. Manufacturing and Processing.....	8
Conclusion.....	8
Bibliography & References.....	9

List of Figures

Title	Page Number
Figure 1 - IoT Use Case Diversity	4
Figure 2 - Simple Farming Beginnings.....	4
Figure 3 - Global Dependencies	5
Figure 4 - Design Concept: Rich Execution Environment vs. Trusted Execution Environment	6

Introduction

Powerful and robust communication networks are a foundation of the global economy, and they are already sparking dramatic transformations in industry and society by enabling new ways of innovating, collaborating, socializing and communicating. While this shift to a hyper-connected, open society brings about many opportunities, it will also introduce many new threats, risks and obstacles. As greater value is extracted from networks and new business structures, the threats are also adapting, becoming more frequent, more sophisticated and more impactful.

A secure communications infrastructure which integrates multiple ecosystems, such as the Internet of Things (IoT) is the foundation for the hyper-connected society. Services for society and business will share similar infrastructure but with different security requirements. We expect next generation networks to enable greater reliability, faster throughput and lower latency as user, device & application demands continue to increase. These requirements call for a new generation of services that ensure end-to-end security across diverse/innovative architectural models. Data integrity and protection is one of the top concerns for operators, enterprises, governments and regulators. As data flows across organizational boundaries and nations, it must be protected at all stages - as it is generated, stored, transmitted, and used over both trusted and untrusted ecosystems.

Future networks will be designed to serve a variety of applications and solutions for people as well as business and connected industries such as manufacturing and processing, intelligent transport, smart grids and e-health. This will result in more complex management of security, privacy and trust across the “things” that make up an IoT ecosystem. And it will also result in far reaching dependencies on the data that is created in one IoT ecosystem and consumers of that data. This paper will examine steps and recommendations for ways to bring a practical approach to securing IoT devices, platforms and ecosystems.

A Hyper-Connected Society

Ericsson is currently predicting 29 billion connected devices by 2022 including about 18 billion IoT devices. Just this year it is expected that the number of IoT devices will surpass the number of mobile phones. We add new ones to our lives now without thinking. Fitness trackers, smart scales, baby and pet monitors, refrigerators, connected cars and other products differentiate from competition through software. And they connect to applications and databases, web pages and other devices with our explicit permission with little regard to security, privacy and trust. When everything takes a username and password as a basic method of authentication providing authorization to produce and consume information that may be very private, it is very tempting to use the same credentials for everything. Product vendors effectively distribute the blame for predictable security breaches to end users by adopting more complicated authentication methods.

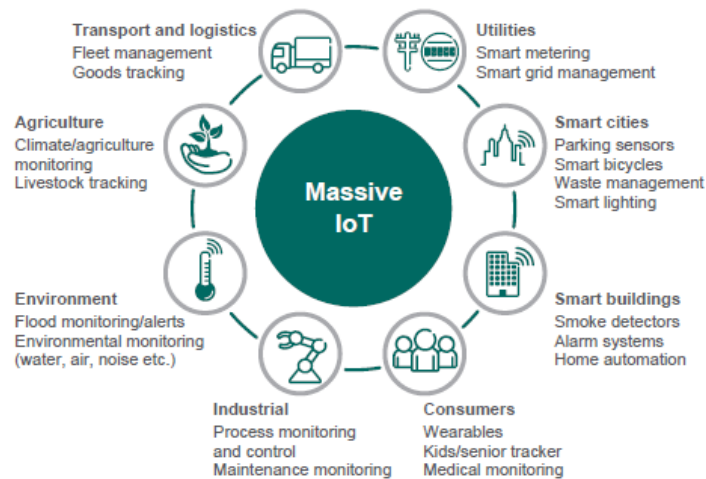


Figure 1 - IoT Use Case Diversity

1. How did we get here? Agriculture as a case study



Figure 2 - Simple Farming Beginnings

It is easy to see that things are changing, but sometimes the benefit of new devices, services, and methods obscures the complexity of their implementation and the risk that accumulates with added convenience and efficiency. Take for example, Agriculture. Originally, man was dependent on rainfall for both their survival and their crops'. Dry or wet years were life or death situations. But there was no interdependence between farmers. They adapted and invented irrigation, which was more predictable, but required an above ground water source like a river or stream. At this stage This limited their capacity to produce and therefore proliferate. Wells, wheels and wind allowed irrigation to be derived from distant or underground sources. There was now some dependence on others, for instance, to not install a dam upstream, or consume too much from the source. Much later diesel pumps and powered irrigation systems allowed better control of both the source and application of water. If it rained too much, the irrigation was not turned on. Dependence on the cost of fuel, availability of parts for the machinery was added to the equation, but more land could be used, and more crops were grown. The systems were improved with rain sensors and automation to reduce the labor requirements and chances for error. A Farmer's success was less more predictable, but at the cost of increased dependence on other factors outside of their control.

On the modern farm today, there are not just many dependencies, there are many interdependencies. Now, farms can be equipped with a matrix of optical and electrical sensors that are GPS located and actively test soil nutrients, moisture, airflow and other conditions. Farmers don't just control irrigation, fertilization, and harvesting dynamically for all parts of the farm based on data from these sensors. They also compare this data with information from multiple weather services to predictively plan for these functions. And they use the data collected to predict crop yield. Data from multiple farms is analyzed by industry and government agencies to predict crop yields and therefore prices, that the farmers can use to plan season to season. Packers, smart transport systems, manufacturers, wholesalers and retailers scale

their businesses based on models of supply and demand, changing costs of labor and fuel, and global consumption. Smart transportation systems, manufacturing and distribution robots and smart grids are tuned and configured based on what has to be trusted information. Governments negotiate treaties and levy taxes based on this data. Monetary transactions from purchasing and selling, to investing, borrowing, and lending happen based on the information. All these players in the economy fed by Agriculture provide and consume data that is critical to all the others.



Figure 3 - Global Dependencies

So while the first farmers would live or die by the rain and their ability to fend off rabbits and deer eating their crops, today's farmers are part of an enormous cascading set of dependencies on their secure connections to places, people, and things. Security of data to ensure its integrity is required to make sure all players in the economy are acting on the right information. Security of access to the network and systems using, transmitting and storing the data is complicated by the sheer quantity and difference of all the devices. Cyber physical-systems like the sensor networks in the Agricultural industry, the smart transport systems, and industrial robots can cause real damage if their security is breached through defect or hack. It is necessary to implement new methods to protect these systems as we move into this next phase. One break in the chain could have global implications if we do not. So, for example, it can't be possible for the farmer to open the door to this whole ecosystem to people with ill intent through deployment of defective, poorly designed or misconfigured sensors. The integrity of the data is too important. New methods must be embraced as our dependencies on the interworking of formerly disparate systems grow.

2. Changes in the way software and hardware products are designed

Agile has largely replaced waterfall software development methodology in many sectors. Small, iterative releases of products with few new features, bug fixes and security patches is how things work. This is

recognizable when apps are updated on mobile phones every few weeks. But with IoT, often the software on the devices cannot be upgraded after the product is deployed. Many of these products operate on 2AA batteries that are required to last for five or ten years. The device, say a moisture sensor, turns on rarely, connects to a wireless network, and transmits a small amount of data before going back to sleep. So a mistake in design could have critical impact on securing the systems the device is connected to. And the mistake of trusting the end user to configure secure authentication of the device on the network has been proven.

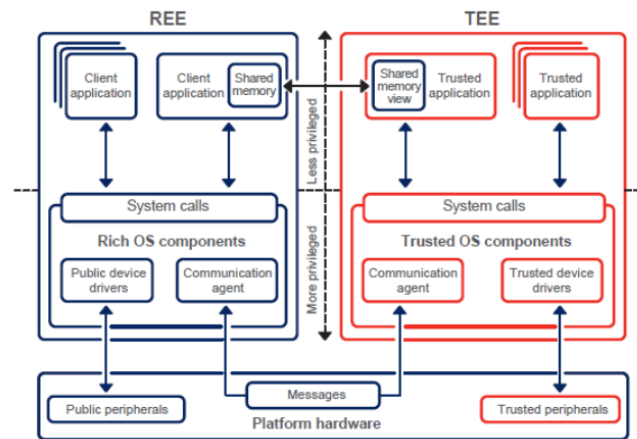


Figure 4 - Design Concept: Rich Execution Environment vs. Trusted Execution Environment

A better approach is to design products with the security built in. While cellular networks have built in authentication and device management, most IoT devices are connected to the network through other interfaces like WiFi, Bluetooth, Zigbee, and NFC from a hub that connects to the Internet through the cellular network. Platforms on the cellular network can facilitate secure bootstrapping of devices and remote application and management of authentication credentials the first time the devices connect. This eliminates user error. Also, the cellular network is a good place for protection from distributed denial of service attacks and rogue devices degrading the trustworthiness of the data provided by a device class. The cellular network is a good platform to apply proven and standards driven security protection to a vast array of IoT devices.

And the proper application of encryption at the device can eliminate the possibility of tampering with the data provided. Devices should be able to verify software updates and boot code cryptographically so that they can't be replaced with malware. A root of trust can be created through separation of these functions logically or even physically through purpose-built ASICs from applications run on the devices in the most sensitive applications.

3. Intelligent Transport Systems

While self-driving cars get much of the attention when it comes to Intelligent Transport Systems, we are still not at a point where systems enabling this technology can be practically hosted in the network. There is just too much at risk. This is an example of a security problem with available solutions that still are not trusted. Guidance and collision avoidance systems for a moving driverless automobile are still largely hosted in the automobile and not just for latency and network reliability concerns. But software is updated, configurations and routes are sent, and traffic and road data is downloaded wirelessly. So, security and trust structures shared between the vehicle manufacturer, the map and navigation source, and

the operator of the vehicle. As with the other use cases, authentication credentials cannot be left up to the end user and authorization profiles need to be shared in real time between public, commercial and private entities. When the technology truly takes off, there will also be privacy concerns about driver behavior and destinations. So, a car's key needs to authenticate the passenger biometrically to ensure privacy and security of the vehicle. Software updates must be cryptographically trusted with keys and block chain before being applied. Rollback to a prior build must be allowed for all systems. And mapping and navigation information must be tracked with block chain before used for any chosen route.

This category also extends to smart cities and control of variable speed limits, traffic lights, tolls, parking spots, and public transport. At the scale of a municipality, there is significant attraction for hackers to disrupt these systems. This is another case where cellular technology can significantly improve security. Standard, hardware-based security for authentication of users for charging for tolls, parking spaces, and express lanes comes with 5G. Also, distributed PCI compliant support for low latency microtransactions is a feature of modern cellular network billing platforms. 5G itself provides very low and predictable latency, high bandwidth, support for low power, low cost devices enabling an economical deployment of hundreds of thousands of sensors, switches, and signals. Automation comes with a risk of its own, though, in that anything that can be automated to fix or optimize itself can also be automated to continuously break itself. Trust structures for system configurations, algorithms, and policies must be in place and audits must be continually run to ensure public safety.

4. E-Health

Because of the high stakes for innovation in healthcare, this use case has always been top of mind for IoT. Saving lives and saving money are easy concepts to sell. At the same time, healthcare is a highly regulated industry with significant interest in privacy, security and trust of data consumed and created. The category extends far beyond IoT to expert systems, doctor / patient portals, and telemedicine. Privacy and trust of data are primary concerns. But the risks get greater when categories like health monitors, pharmaceutical dispensers, surgical and other hospital robots, and elderly patient tracking are considered. It is possible for the "things" in the internet of things to cause real, direct physical harm to people and property.

Care must be taken to design the devices used in E-Health with security, safety and privacy in mind as well as compliance to evolving laws and standards. The price point for many of these devices allows for implementation of hardware encryption, fail safes, and communication and data integrity checks. Because this is so important, government and professional organizations must and will expand regulatory oversight and require more advanced certification of devices against their standards.

5. Smart Grid

Smart grids provide efficiencies through coordination of production, sale, transmission, distribution and consumption of power using information and communications technology. As with the use case, these efficiencies come at the cost of dependencies. All components of this model are potential entry points for bad actors looking to create widespread chaos, demand ransom payments through threat of said chaos, or invade the privacy of consumers. Protection of the smart grid is a critical component of any country's defense strategy as an attack could be debilitating and difficult to reverse.

Many of the vulnerabilities of the smart grid come from the smart part ironically. Unlike the long-lived power equipment and infrastructure, IT equipment typically has a three to five year lifecycle. The mismatch of the lifecycles means that it is very likely that compliance to security standards for the IT equipment will drift as it ages. Automated audit and updating of the equipment and software to maintain

compliance is a good way to avoid this and modern automation and orchestration systems are a good way to counteract this. Besides this, best practices for securing the IT infrastructure, LAN, and WAN are critical as the stakes are high.

The addition of smart thermostats and home automation systems as a popular product category represent a privacy risk that many consumers do not recognize. As these systems adjust temperatures and activate lights and other devices based on the user's habits and presence, access to this information can allow criminals to understand when the house is or is not occupied. This can also be done from a hacked smart meter. Manufacturers of these devices, like others mentioned, cannot place the entire responsibility for security on the end user's ability to maintain and secure a user account.

6. Manufacturing and Processing

Autonomous systems for manufacturing and distribution of goods are becoming more justifiable than ever. In the past, robots replaced humans for jobs that required precision, speed, significant repetition, or hard to find valuable skill sets. As the cost of these robots and their control systems declined, they became useful for replacing humans in tasks humans find unpleasant. And in areas with unreliable labor forces, it is cheaper to automate a job than to train many different people for the same job over the course of a year as they leave too often. Unlike industrial automation systems of the past that required extensive retrofit of the facility, on site IT and automation platform staff, and significant investment, it is possible to buy one robot for a specific purpose. Robots are managed from central platforms on cloud infrastructure to eliminate the need for extra on-site staff. And automation systems can be coordinated between multiple facilities that could be thousands of miles apart. The coordination of order and inventory systems, real time location, collision avoidance and even intelligent transport systems can improve the payback on investment significantly while improving quality and customer satisfaction. But considerations must be made in network design and security to make this possible. There are significant safety concerns for on site personnel that must be mitigated by trusted distribution of system software updates, configurations and orders. Because these systems may be enhanced by adoption of edge networking technologies to reduce latency, distribution of platform logic provides an opportunity for hackers and considerable liability. Application of blockchain to ensure all players in the ecosystem are acting from trusted data is a good solution. And active assurance of compliance to security, privacy, industry standard, and customer service level agreements in a fungible cloudified network is vital. Orchestration and automation of this active oversight can make this efficient. Machine learning and Artificial Intelligence applied to the vast amounts of data created by these systems can aid in the detection and correction of anomalous behavior in real time.

Conclusion

The Internet of Things grows as it makes innovation, efficiency and convenience easy. But it comes at the price of dependencies between previously unconnected systems that significantly escalates the severity of any security breach. While the use cases described have their own set of dependencies, they also collectively have a common set that could be exploited. The GPS network, or public cloud platforms are examples. It is critical that security strategies be employed by the vendors of IoT products, providers of platforms, and deployers of services that consider all dependencies. As the IoT grows, it will be more common for hackers to exploit unexpected dependencies than products or services, as it will be difficult to track them. A start, though, is to employ design, management, and maintenance techniques that are developed with careful thought of these dependencies and ensure security, privacy, and trust of all players in the IoT ecosystem. And devices, networks and services should be audited and tested by independent labs and standards bodies to ensure they will be good citizens in the universe of dependencies.

Bibliography & References

Ericsson Mobility Report, Ericsson November 2016

Ericsson Mobility Report, Ericsson,

5G Security, Ericsson June 2017

Ensuring Critical Communication with a Secure National Symbiotic Network, Ericsson May 2018

Cellular Networks for Massive IoT, Ericsson January 2016

CTIA Cybersecurity Certification Test Plan for IoT, CTIA May 2018

Self-Service Dimensional Data Analytics

Scalable Patterns for Data-Driven Enterprises

A Technical Paper prepared for SCTE•ISBE by

Francesco Dorigo

Senior Manager

Comcast

1400 Wewatta Street, Denver, CO 80202

+1 720 512 3674

francesco_dorigo@comcast.com

Bao Nguyen

Principal Engineer

Comcast

1400 Wewatta Street, Denver, CO 80202

+1 720 512 3687

bao_nguyen@comcast.com

Daniel Howell

Senior Engineer

Comcast

1400 Wewatta Street, Denver, CO 80202

+1 720 512 3693

daniel_howell2@comcast.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Collection Tier: Headwaters	4
1. Collection Tier Components	5
2. HTTP-Collector: The First Layer of Data Acquisition	6
2.1. From semi-structured to structured data	6
2.2. Multi-tenancy	6
3. Schema Registry and Schema Evolution.....	6
4. Event-preprocessor: Data Validation and Enrichment.....	7
Aggregation Tier: Vortex	7
5. System Overview	7
6. Components Diagram	8
7. Vortex Aggregator	8
7.1. Ingress SDP/Avro Event Data.....	9
7.2. Vortex Processing Layer (EMR/YARN)	9
7.2.1. Enable Generic Multi-dimensions	9
7.2.2. Windowing and Watermarking	10
7.2.3. Publish Aggregate Results.....	10
7.3. Persistence/ETL Layer.....	10
7.3.1. Medium Term Persistence	10
7.3.2. Enrichment Sources.....	10
7.3.3. Lambda ETL.....	10
7.3.4. Cache	11
7.3.5. Lambda Protocol Specific	11
7.4. Service Layer	11
7.5. External System Requests.....	11
8. Vortex Manager.....	12
8.1. View saved or running configurations	12
8.2. Schema retriever & parser	12
8.3. Query builder & validator.....	12
8.4. Job execution status	13
8.5. Start new Vortex aggregation job	13
9. Vortex Analyzer	14
9.1. Sample event data	14
9.2. Analyze & Validate cardinalities	15
9.3. Cardinality threshold violation notification	15
10. Example of Vortex Application	15
Conclusion.....	16
Abbreviations	17

List of Figures

Title	Page Number
Figure 1 – Data Pipeline.....	4
Figure 2 – Collection Tier System Diagram	5
Figure 3 – Vortex Components Diagram	8

Figure 4 – Vortex Aggregator System.....	9
Figure 5 – Vortex Manager Components.....	12
Figure 6 – Vortex Analyzer Diagram.....	14
Figure 7 - Anomaly detection in example.....	16

List of Tables

Title	Page Number
Table 1 – Collection Tier System Description	5
Table 2 – Vortex Components	8

Introduction

The IP video analytics platform design described herein streamlines the conversion of event-based analytics telemetry into time series visualization. With this objective in mind, Comcast developed a platform with the flexibility to support a wide variety of data producers and data consumers, and with the scalability to provide an enterprise solution for real-time analytics.

Parametrized and configurable execution libraries automate repetitive data engineering tasks. In addition, our analytics system provides abstraction layers both for data ingress and data egress, which enables a seamless evolution of the ETL (Extract, Transform, Load) pipeline. This has two main advantages: First to simplify the efforts related to upgrading the pipeline by letting producers and consumers evolve at their own pace, and second, the underlying technologies can seamlessly evolve with zero impact for ingestion and aggregation layers (producers and consumers).

This document describes the major components of our IP video analytics data pipeline, with a specific focus on custom components. This design reduces the time between data ingress and insight.

Our custom components are shown in Figure 1 as the collection tier (internally called “Headwaters”) and the aggregation tier (internally called “Vortex”).

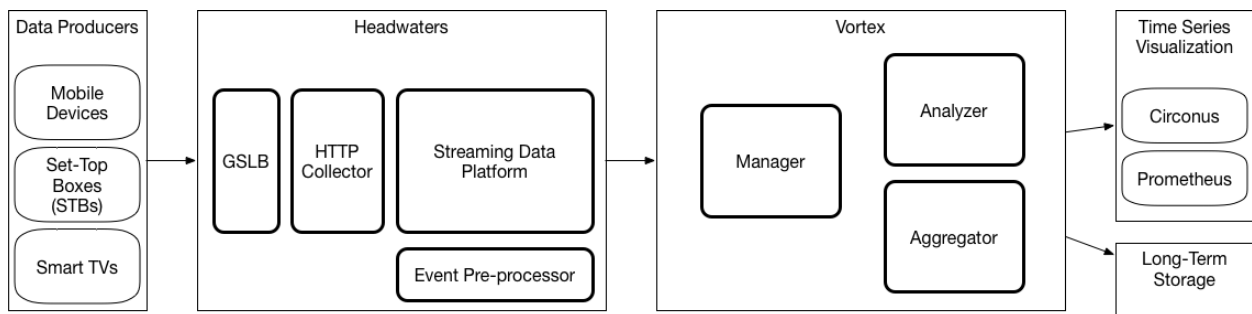


Figure 1 – Data Pipeline

Collection Tier: Headwaters

Our data collection tier was originally designed to ingest analytic events generated while streaming video content from an IP video player device. The expected growth and evolution of Xfinity TV applications call for a fast and reliable collection tier that can guarantee an actionable level of operational monitoring for system health and customers’ experience. These aspects require the collection tier to be highly configurable, reliable and scalable when adding new data sources.

The IP video analytics pipeline supports high volume, high velocity, semi structured data acquisition through a collection tier based on the HTTP protocol. The collection tier provides a REST-compliant HTTP endpoint, ensuring data extensibility, which allows most systems to natively use the REST (Representational State Transfer) architecture. For systems unable to provide data via the REST endpoint, it is possible to deploy a thin, client-side forwarder to bridge the gap.

Under the hood, Headwaters is a streaming data platform (SDP) comprised of a REST API to receive data, and is clustered using Apache Kafka, which handles the queueing for incoming data streams. Additionally, Headwaters enforces serialization and schema governance, using Apache Avro, and provides a Confluent schema registry for event serialization/deserialization.

1. Collection Tier Components

Table 1 – Collection Tier System Description

HTTP-Collector	Exposes a REST API to clients not supporting direct Kafka APIs for publisher connections with the Kafka cluster. Performs data validation and routes data to the appropriate Kafka topics.
Streaming Data Platform	Kafka-based cluster with regional and national clusters. Allows for replication and mirroring across geographic regions/zones.
Schema Registry	Avro-based Confluent schema registry used to regulate published data content to the data exchange cluster.
EventEvent Pre-processor	Spark streaming application converting ingested data into meaningful structured data events facilitating downstream consumption.

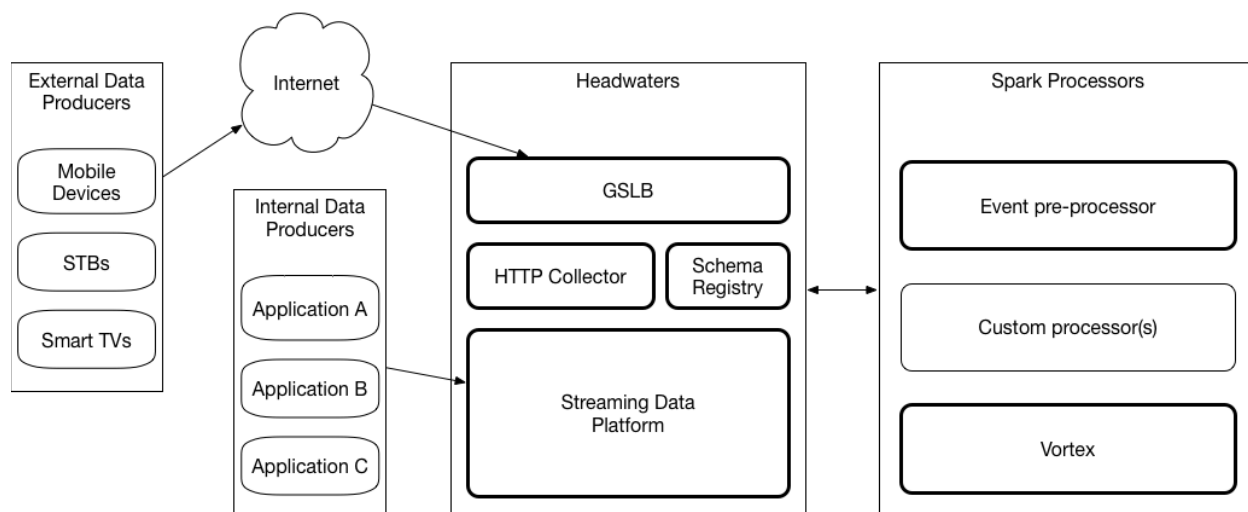


Figure 2 – Collection Tier System Diagram

At any given time, Xfinity TV applications are running on millions of devices and a wide variety of platforms (STB, iOS, Android, Roku, in-browser applications, etc.). Each of these devices are sending massive amounts of event data to the HTTP-Collector. These events represent a blend of the current internal status of the video player and the quality of the user experience as measured through predefined parameters such as startup time, bit-rate, and video format (SD/HD).

For instance, video player applications are required to send synchronous events (referred to as *heartbeats*) to periodically report current player activity and internal state. Heartbeat events are used for statistical analysis and trend forecasting and don't require immediate action from downstream consumers. On the other hand, asynchronous events are used to communicate unexpected status changes, such as warning and error conditions or user input channel changes, video playback pausing and fast forwarding, etc. The error events are dealt with immediately.

Beside the set of pre-defined events a client can send, Headwaters provides the flexibility to add new event types or include custom fields within an existing event. These customizations are performed without code changes to either the collection tier or the data consumer tier. For this reason, the architecture presented in this paper is utilized to process events generated by systems other than IP video player devices. Adding a new data source will not require any changes to the collection tier, which significantly lowers the barrier to entry for new data producers.

2. HTTP-Collector: The First Layer of Data Acquisition

The HTTP-Collector is the first layer in the data acquisition process and is bundled as a light-weight webserver. This service acts as an interface for ingesting player events into Headwaters/Kafka topics where raw data is collected and transformed. The interface enables multi-tenant REST interfaces with configurable endpoints, which allows configured clients to send event data into Comcast's analytics pipeline via HTTP for processing. The HTTP-Collector is built for horizontal scale, meaning that extremely high-volume data ingestion is supported at real-time latency. Received clients' requests are structurally validated and routed into Comcast's Kafka streaming data platform.

2.1. From semi-structured to structured data

The HTTP-Collector supports the variety of player clients by allowing both generic JSON-based payloads and specific Avro-encoded payloads. In the first case, the JSON data is wrapped in a defined Avro record before being published to Kafka.

A specific consumer, the event pre-processor (detailed in section 4), is used to perform deeper validation and transforms semi-structured events into structured Avro records. In other words, the pre-processor transforms the data received as JSON into fully qualified Avro records, according to the corresponding schema, which are routed back to the appropriate Kafka topics for downstream consumption.

2.2. Multi-tenancy

Multi-tenancy is achieved by exposing different endpoints corresponding to each data producing system. In other words, when logical separation between data sets is desired, separate URLs are used to support proper data routing for data posts. For example, Comcast's syndication partners have dedicated endpoints to post event data. Each of the endpoints is backed by dedicated topics in the streaming data platform (SDP). Other systems could take advantage of a similar approach and rely on post processing systems to correlate their events with other systems' events.

3. Schema Registry and Schema Evolution

Comcast's streaming data platform (SDP) handles continuous data flow from multiple systems; each system utilizes a dedicated topic for its data stream. The Headwaters data collection tier requires an Avro schema, per topic, to enforce governance and to ensure that the content of each topic is simpler to share (content discovery). We leverage Apache Avro to serialize data and manage schemas using a Confluent schema registry to store and govern schema evolution. Avro provides rich data structures that offer a fast and compact binary data format, which allows each datum to be written with minimal overhead. The result is a more efficient data encoding and faster data processing.

Schema evolution and governance are crucial in all IP video analytics pipelines and provides the automatic transformation of an Avro schema. This transformation is only applied during deserialization. If the reader's schema is different from the writer's schema, the value is automatically modified during deserialization to conform to the reader schema using default values. Different teams and organizations within Comcast manage their own schemas for data they produce. These schemas are added, modified, or removed frequently to meet the teams' requirements, which are reviewed by a governance body before being merged. The Confluent schema registry forces schemas to be registered and associated with the appropriate topic before data can be published into Kafka.

The Confluent schema registry contains Avro schemas which are associated on a per-topic level. Each schema is used by a consumer application when de-serializing Avro event topics, because the Avro schema itself is not supplied on the Kafka event data record, only the schema ID. Our SDP Kafka topics contain Avro-serialized data only, which was previously transformed using an associated SDP Avro schema. The schema registry is responsible for serving up the associated Avro schema to provide the ability to properly de-serialize each Avro record.

4. Event-preprocessor: Data Validation and Enrichment

The purpose of the event-preprocessor is to provide correct and consistent data to downstream processing systems. The event-preprocessor is a Kafka consumer application, ensuring that the data received from the HTTP-Collector conforms to the fields defined in the Avro Schema. The event-preprocessor serializes each of the JSON events into an Avro event and decorates each with derived or sourced data from external data repositories. If the data is not conforming to the expected required fields in the Avro schema, the resulting record is tagged with warnings or errors so that downstream consumers can independently decide how they should be used.

Aggregation Tier: Vortex

5. System Overview

Data collected through the Headwaters data exchange platform is made available for any client able to consume data from a Kafka topic. Vortex is a collection of consumers, which aims to simplify the task of creating data aggregations across team domains. As a result, it saves the teams' time and shortens feature delivery, while also reducing team/system overhead. These aspects allow teams to focus on what matters most and obviates the need to create *boilerplate* data pipelines.

Data ingested by the Vortex Aggregator is targeted for structured event data, which could be specified as Parquet, JSON or Avro in this case. The data at this layer has already been cleansed by the event preprocessor (described in section 4), which is responsible for cleaning and transforming JSON data into the appropriate Avro using the proper Avro schema from the enterprise schema registry.

6. Components Diagram

Table 2 – Vortex Components

Aggregator	Responsible for processing a stream of data into multi-dimensional aggregates based on a generic configuration, the core of which hinges on an Apache Spark SQL statement. The Vortex Aggregator provides the data to the persistence layers, where the data can be served up through a REST interface for external system requests.
UI Manager	Allows for system end-users to check which jobs are running for their respective group and the associated configuration for each job. When a user wants to create a new job based on a new generic configuration, the user can use the UI and build an aggregation query, because the UI has fetched the schema registry for element selection. Configuration error validation is also performed at this layer. Lastly, the Vortex Manager enables cardinality validation. Cardinality validation ensures the level of cardinalities being asked to execute for the Vortex Aggregator are within pre-defined tolerances, which are performed by requesting cardinality counts from the Vortex Analyzer.
Analyzer	Responsible for sampling various streaming data topics to determine qualifying dimensional cardinalities. This validation acts as an execution rule for the Vortex Aggregator, since a cardinality set too large isn't considered useful for the end user. Additionally, an outcome of this cardinality processing allows the system to obtain the top-level cardinalities. For example, after Vortex Analyzer has run for a short time interval, the system understands the top X (i.e. top 3 or top 10) dimensional cardinalities, which can then be used directly by the Vortex Aggregator for processing via the generic configuration.

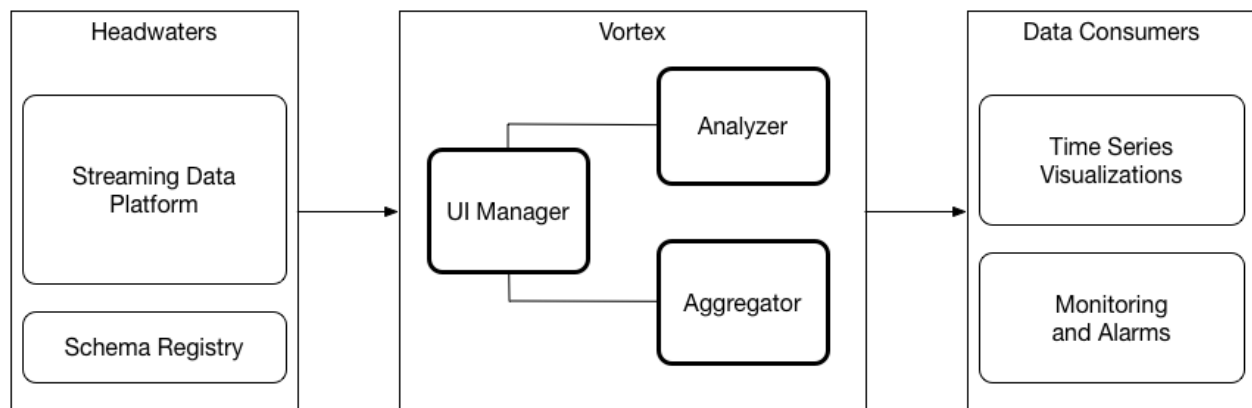


Figure 3 – Vortex Components Diagram

7. Vortex Aggregator

Vortex Aggregator is a collection of Spark applications consuming topic events from Headwaters to produce aggregations targeted for time series databases (TSDB). The Spark executables are highly parametrized so that the users can supply their specific business rules via configuration settings, rather than changing the code itself. An additional component, the UI manager (discussed later), further simplifies this task by providing a push button UI to build a generic query for a desired aggregation.

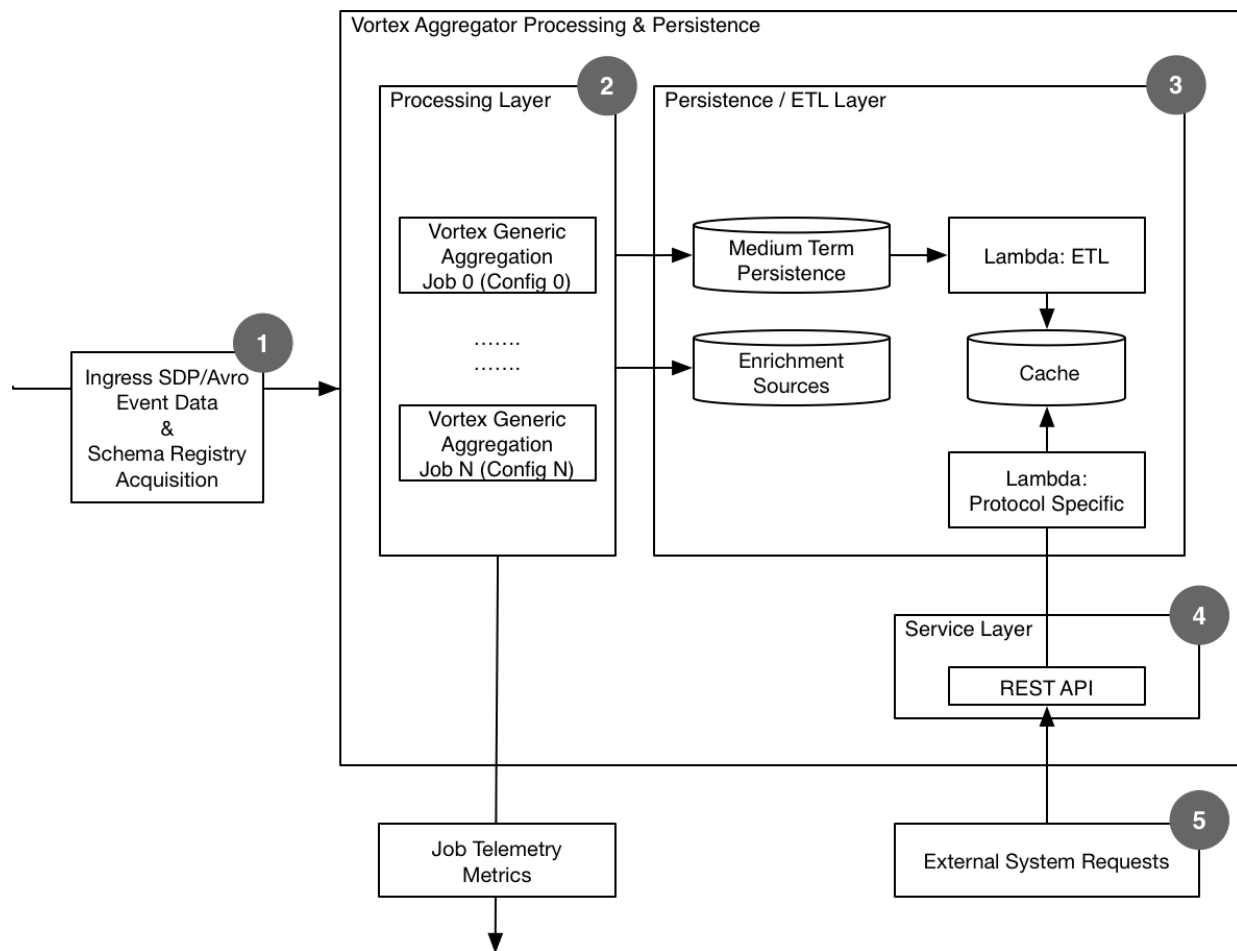


Figure 4 – Vortex Aggregator System

7.1. Ingress SDP/Avro Event Data

As described in the previous sections, the Avro-based events are streamed from any of Comcast’s SDP topics, where the events are then used to produce aggregates results. Avro events are de-serialized using the schema obtained from the Schema Registry.

7.2. Vortex Processing Layer (EMR/YARN)

The processing layer is where various Vortex Aggregation jobs execute. Vortex Aggregator is an Apache Spark application and uses a generic configuration for execution. The generic configuration relies on Spark SQL syntax to provide aggregations across multiple dimensions and cardinalities.

7.2.1. Enable Generic Multi-dimensions

Using the Vortex Manager query builder, a new generic configuration can be supplied to the Vortex Aggregator. The Vortex Manager triggers the creation of a new executable for the Vortex Aggregator. Aggregates are published within several minutes of deploying a new configuration file. This generic approach reduces the engineering effort necessary to go from “requirements” to “insights.”

7.2.2. Windowing and Watermarking

Windowing and watermarking apply directly to the Apache Spark aggregation terminology. In the Vortex Aggregator, a “window” is defined as the duration of time for an aggregation using the event creation timestamp (excluding time/server adjustments). Using these timestamp values allows the proper events to be included into the appropriate window duration. Windowing is also customizable in the system.

“Watermarking” refers to the ability to handle late arriving data. The late arriving data is customized based on a user customized value to determine how late/old the data should be aggregated until those events no longer apply to the corresponding time window. This feature is useful if an application falls behind (in terms of processing), or when events are received out of order or late, to ensure aggregations are calculated properly.

7.2.3. Publish Aggregate Results

An obvious point, but worth noting, is that the goal of the system is to publish well-formed multi-dimensional aggregates into a medium-term persistence. Associating metadata with the aggregates also provides the ability to debug/trace the data from a specific configuration down to the end system, which is often a Time Series Database (TSDB). At each data touch point, timings are recorded, which provides internal latency metrics for throughput and processing speed.

7.3. Persistence/ETL Layer

The persistence layer allows for storage, transformations and data availability. These features are described below in more detail:

7.3.1. Medium Term Persistence

As a consideration of the design, aggregations were deemed valuable to publish into a medium-term persistence. These aggregate objects are kept to ensure system-to-system and configuration tracking for debugging, quality control/assurance and performance analysis. All the data at this level is a combination of the aggregation and the metadata used to generate the aggregate, including the job’s YARN application ID, configuration ID and specific topic-based data. All the metadata enables backward tracing from any point in the pipeline, including from the external systems to the Vortex Manager/user generic configuration request.

7.3.2. Enrichment Sources

An enrichment source can be considered as data which provides some degree of value for the aggregation from an additional dataset. The enrichment data at this level is considered small and can be applied using a join to help avoid verbose data shuffling across executors/partitions. For example, the joined dataset could relate to Geo Location information, which is joined on IPv4/6 addresses contained within each event. The joined data does not identify customer details but allows for decoration of market level information in real-time.

7.3.3. Lambda ETL

When a new object is created in the medium-term persistence, a trigger is fired to invoke a serverless Lambda function, which transforms the aggregate by stripping off most of the metadata. The aggregate transformation (counter, gauge & histogram data) is then written into a short-term persistence/cache. At

this point, the aggregate data is ready for consumption by external systems, such as a time series database (TSDB).

7.3.4. *Cache*

The cache is used to hold the results for each Vortex aggregate application. When a new aggregate is created, only the most recent aggregate for the configuration ID is made available to downstream systems. This ensures that only the most recent version of the record is kept within the cache and made available to downstream systems.

7.3.5. *Lambda Protocol Specific*

When an external system requests aggregates from the service layer, a specific Lambda function will trigger. Each of the Lambdas are mapped directly to one endpoint (there are multiples) allowing for the same data in the cache to be transformed into a specific protocol. This flexibility allows for the same cached data to be served when called by each of the varying TSDB's where a specific format is required. An example of this conversion is when a system needs metrics in JSON vs. text-based formats supporting OpenTSDB standards.

7.4. Service Layer

The service layer provides a REST API which makes the distinct set of multi-dimensional aggregates available across one-to-many systems via HTTP GET request. This layer ensures aggregate data is extensible by design and provides the ability to transform the request into the proper response application/content-type.

7.5. External System Requests

External system requests made through the REST interface and enables:

- Time series database (TSDB), such as Circonus or Prometheus, where temporal aggregate data visualizations can be quickly created and augmented – supplying new actionable insights for a teams' workflow (alert, support, help to identify root cause, etc.)
- Any other system with the capability to make HTTP GET requests. These types of systems could be developer controlled or Quality Assurance systems for validating engineering-based changes and maintaining simulated canned test scenario performance.

8. Vortex Manager

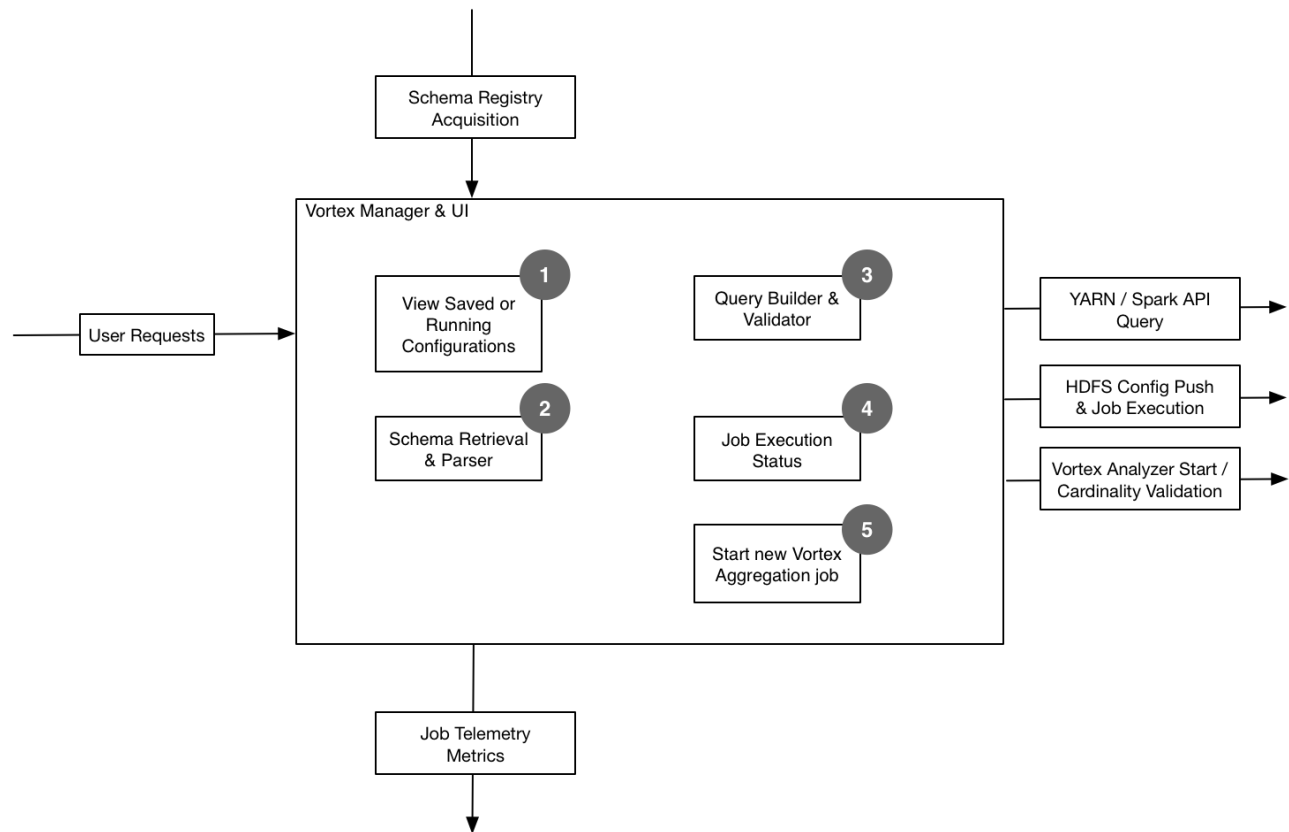


Figure 5 – Vortex Manager Components

8.1. View saved or running configurations

A user can view saved or running generic configurations. This helps to ensure that duplicate Vortex Aggregator applications/jobs are not executing. This is a combined view of the job status with the configurations each Vortex Aggregator application being processed.

8.2. Schema retriever & parser

Behind the scenes, the manager may already have the most recent version of the schema for a topic; however, when a user needs to browse a new topic, the manager requests the schema and parses the schema into a traversable “tree”. This “tree” is then made available to the end user to point/click on one-to-many elements, which saves the user from having to know the full schema object names (which can be quite complex due to nesting.)

8.3. Query builder & validator

The query builder and validator systems allow the user to build the desired aggregation using the point/click system (additionally, other fields can be set here, such as the window and watermark durations). Users can also apply various Spark SQL syntax expressions, via filters, conditionals & SQL-based functions. Once a user’s selection has been completed, the Spark SQL statement is presented for

review prior to submission. During this step, validation occurs and provides areas where the generic configuration may be invalid, requiring user correction.

Once there are no further errors in the generic configuration, the user submits the request, which probes the Vortex Analyzer for information on the validity of the desired dimension and/or cardinalities. The Vortex Analyzer determines if the cardinality for the requested dimensions meet the pre-determined thresholds. This check is in place to ensure that the configuration will be useful to the end customer and ensures the Vortex Aggregator can successfully process the desired request.

8.4. Job execution status

An additional feature built into the Vortex Manager is to obtain information about the Vortex Spark applications using the YARN and Spark APIs. This information can be provided to the user as an ad-hoc request or triggered as part of the workflow to determine if an identical job is still running, in addition to identifying the job running state after submission. Additionally, the generic configurations are coupled to each current/past job.

8.5. Start new Vortex aggregation job

When the above steps meet the criterion, the configuration is stored locally for the Vortex Manager and becomes securely copied into HDFS. Once the generic configuration is uploaded, the spark-submit script (with arguments) is provided for Vortex Aggregator execution. From here, YARN/Spark manages itself by acquiring the appropriate resources and copying the files/code to each of the executors for execution.

In the meantime, the Vortex Manager supplies the user with information about the status of the job using the typical YARN states. Once in “RUNNING” state, various information is recorded about the job and associated back to the user’s request.

9. Vortex Analyzer

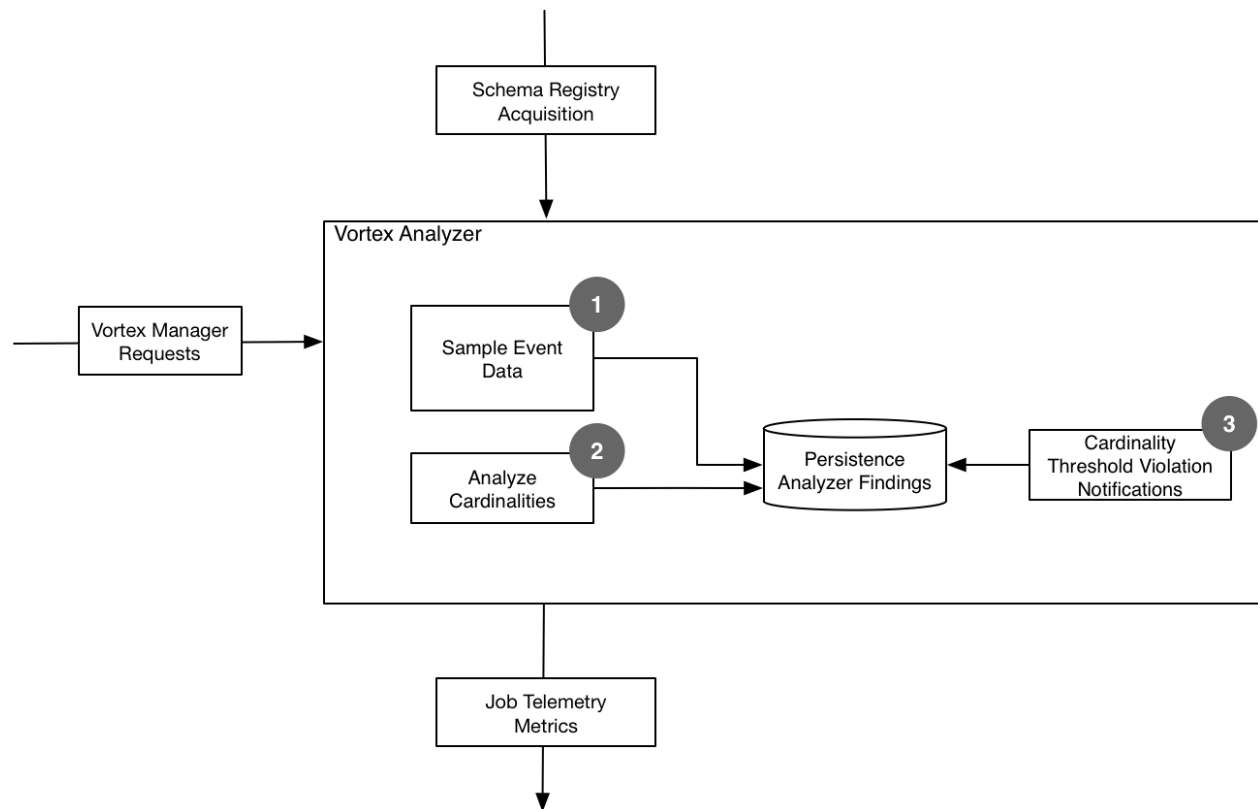


Figure 6 – Vortex Analyzer Diagram

9.1. Sample event data

Vortex Analyzer samples data for various topics to record the dimension and cardinality state (number of occurrences, i.e., depth of the dimensions). When the Vortex Analyzer runs in “normal” mode, it executes in the terms of hours per day, which is enough time to record cardinality state. “Quick” mode offers a fast inspection of the event topic data to determine the same dimension and cardinality state, which is a tradeoff between time to execution vs. confidence levels.

When an unsupervised Vortex Analyzer topic is requested, a user can apply two execution states – “quick” or “normal” -- both resulting in a new Vortex Analyzer job but with differing degrees of results. In quick mode, the Vortex Analyzer will run for just a matter of minutes, and, based on those findings, a validation decision will be made to determine if the Vortex Aggregator can execute for the requested dimensions and cardinalities. Quick mode allows the user to get aggregates flowing as quickly as possible for a new “Headwaters” topic. In normal mode, the same analysis is performed, but provides a much higher level of confidence for the witnessed data. Once the execution time has passed (usually measured in hours), the validation results will be returned to the Vortex Manager with a callback and will allow the user to then submit the job. The execution of the job can be performed autonomously, so the user doesn’t need to wait for the execution of validation results.

When either mode is selected by a user, the option exists to enable the Vortex Analyzer to run continuously, because other validations may be requested in the future, which makes the topic supervised by the Vortex Analyzer. All quick mode jobs will be scheduled into normal mode jobs when a user requests for continuous validation on the specific topic. This is a feature of the scheduler built into the Vortex Analyzer.

9.2. Analyze & Validate cardinalities

Dimensional cardinalities are examined to ensure the breadth of the data fits within the pre-defined tolerances and doesn't violate the expansion rules. The rules in place provide the ability to limit jobs from running thousands of cardinalities, which wouldn't be useful from a TSDB visualization perspective. However, across large data sets, there may be a need to examine only a handful of dimensional cardinalities. The Vortex Analyzer provides Vortex Manager with the top "X" dimensional cardinalities. This allows a filter to be applied for the Vortex Aggregator and enables a level of control that wouldn't have been possible by using a simple click from the Vortex Manager.

For a user to enable this feature, a selection in the Vortex Manager is applied for the top "X" data points available, where "X" is scalable up to a bounded limit. This way, the user can understand the top values for large datasets and encourage exploratory analysis using other tools within the data analysis ecosystem.

9.3. Cardinality threshold violation notification

One additional feature of the Vortex Analyzer is to analyze cardinality threshold violations for executing Vortex Aggregator jobs. This background processing handles post analysis so as to deeply analyze acquired data, which identifies once "thought to be good - valid" configurations to "bad - invalid." Such an occurrence could be caused by an upstream application release.

To account for this possibility, the topic data is continuously supervised/analyzed. When a threshold violation is triggered, a notification is sent to the manager (and to the internal telemetry metric system, which is distinct from external TSDB's). The Vortex Manager will indicate the problem with a level of confidence for the possibility of failure/impact. Additionally, the telemetry metric system will deliver an alert as part of SRE/OPS support. If the job fails, the application service manager for YARN/Spark (a standalone background component) will not attempt to restart the job, because it now violates the pre-defined rules of an acceptable configuration and requires a user to re-issue and validate the requested configuration. Typically, when this edge case occurs, a user can submit a new generic configuration using the maximum top X, which typically removes the violation. This feature provides a technique to protect the Vortex Aggregator and the user's end system from large degrees of data drift and cardinality explosions.

10. Example of Vortex Application

As data consumers add more and more dashboards to monitor their systems, they discover the inherent value of visualizing the data being collected. Correlations between measurements that belong to the same system enable deeper exploratory analysis. Real-time execution of aggregations shortens the mean time to detect (MTTD) and react to operational issues.

In the example below, several platforms were experiencing an unusually high spike in error rate per video playback start. The operation and engineering team was alerted immediately, and focused on the platform with the highest error rate, diagnosed the problem, and resolved it in a timely fashion. Figure 7 shows the session starts for iOS, desktop, and Android platforms in purple. The daily trend exhibits the expected

peak near the prime-time hours of the day. The error rate overlaid on the graphs below in green shows a sudden and sustained spike for all platforms.




Figure 7 - Anomaly detection in example

Conclusion

We invest heavily in growing and expanding the analytics capabilities of all components of the IP video delivery platform. This document summarizes the current architecture for a state-of-the-art and end-to-end data pipeline that can scale horizontally to adapt to the growing needs of the enterprise. Moreover, automating and simplifying the data engineering tasks required to aggregate and visualize event-based telemetry provides a low risk migration option for systems that have outgrown their own ad-hoc solutions.

A universally available REST API for data ingestion, a scalable data stream platform, a push button aggregation system, and powerful time series visualization tools are key elements for the successful evolution of a data exchange platform.

As an enterprise data exchange solution, Headwaters and Vortex provide unprecedented data sharing opportunities for all organizations within Comcast. The systems and practices described in this document reduce the effort necessary to collect, prepare, and share the data between internal groups. Having a common solution for most applications allows us to focus all data engineering resources on improving the performance and feature offerings of the data exchange, rather than providing a dedicated ad-hoc solution for each system.

Abbreviations

ETL	Extract, transform, load
SDP	Streaming data platform
STB	Set-top box
TSDB	Time series database

Shifting Left

Harnessing AI to Deliver a Consistent, Engaging Customer Experience

prepared for SCTE•ISBE by

Bradley May

Managing Director – Artificial Intelligence
Communications, Media & Technology

Accenture

1-678-657-4349

bradley.s.may@accenture.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. Three Areas of AI-Assisted Care	3
1.1. Reactive Care.....	4
1.2. Proactive Care.....	5
1.3. Predictive Care.....	6
Making the Shift Left	7
Abbreviations	7
Bibliography & References.....	7

List of Figures

Title	Page Number
Figure 1 - Shifting Left in Customer Care	3
Figure 2 - Components of Reactive AI-Powered Care	4
Figure 3 - Components of Proactive AI-Powered Care	5
Figure 4 - Components of AI-Powered Preventive Care	6

Introduction

Artificial intelligence and automation are coming to the foreground. As machine learning advances, artificial intelligence (AI) will continue to build an influence over all parts of the business and the customer experience. By 2021, nearly one in six customer service interactions globally will be handled by artificial intelligence. And, improved AI technologies will automate parts or all of up to 40 percent of customer service needs by 2019.¹ Indeed AI is a powerful tool at communications services providers (CSPs) disposal that can be deployed to differentiate service delivery.

Content

AI can have dramatic impact on engaging customers and delivering the expeditious and personalized experience they desire. Chatbots, for example, are becoming a common channel for customer interaction. However, the ability for AI to improve the customer experience goes far beyond the applications we typically imagine occurring at the point of customer contact. When a customer initiates an interaction, this “reactive” care is an important element of the customer journey that needs to run efficiently and with positive customer outcomes. It is a natural area for CSPs to apply AI, analytics and other digital technologies. But CSPs also need to be more proactive and predictive in managing the customer experience and herein lies the hidden AI opportunity.

Truly transforming the customer experience requires CSPs to “shift left” in their use of AI: from reactive to proactive to predictive customer care (Figure 1). Shifting left enables the CSP to move from one-to-one customer care to one-to-many customer care. This reduces the cost of care and, by pre-empting problems rather than reacting to them, CSPs can deliver a consistent and reliable experience that results in customer satisfaction.

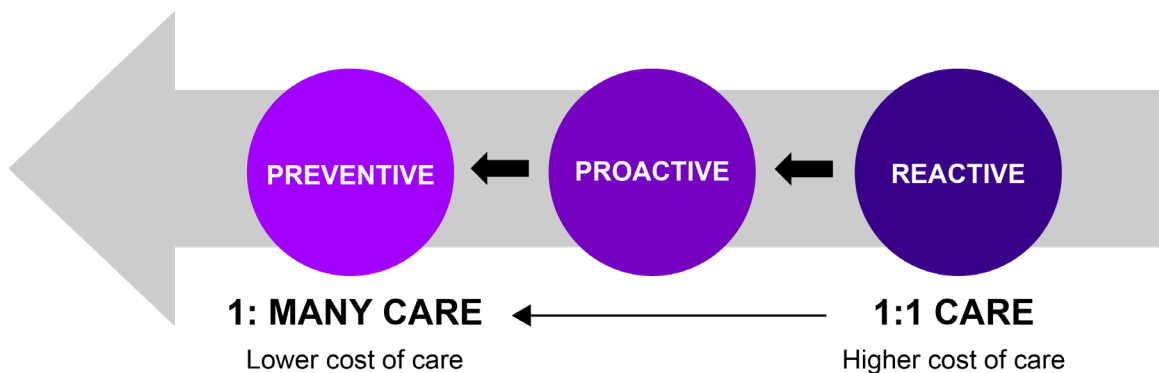


Figure 1 - Shifting Left in Customer Care

1. Three Areas of AI-Assisted Care

Three distinct areas are ripe for using AI in customer care. Each can be tackled independently, but all should be on the customer experience roadmap.

- Reactive care – Applying AI to predict why a customer is calling and get the caller to the right channel and method of interaction.

- Proactive care – Deploying AI to identify potential or likely issues with a customer’s service and taking proactive actions to resolve or advise the customer.
- Predictive care – Integrating AI to identify and resolve potential problems within the network before they materialize or, if an issue does occur, accelerate the resolution to minimize impact.

1.1. Reactive Care

By embedding Artificial Intelligence into reactive customer care, companies can anticipate, engage, and satisfy customers on a one-to-one basis and cost-effectively deliver best-of-breed customer engagement. Analytics and digital conversion capabilities enable CSPs to predict why the caller is calling, move them to the right channel and method of interaction, (preferably digital) and, where possible, automate the interaction. Virtual agents automate chat experiences – both transactional and informational. Human agents are empowered with intelligent automation tools to drive higher concurrency. And, contextual information is available across all channels to create a unique omni-channel experience.

Multiple human and technology components are needed to provide AI-powered reactive care (Figure 2).

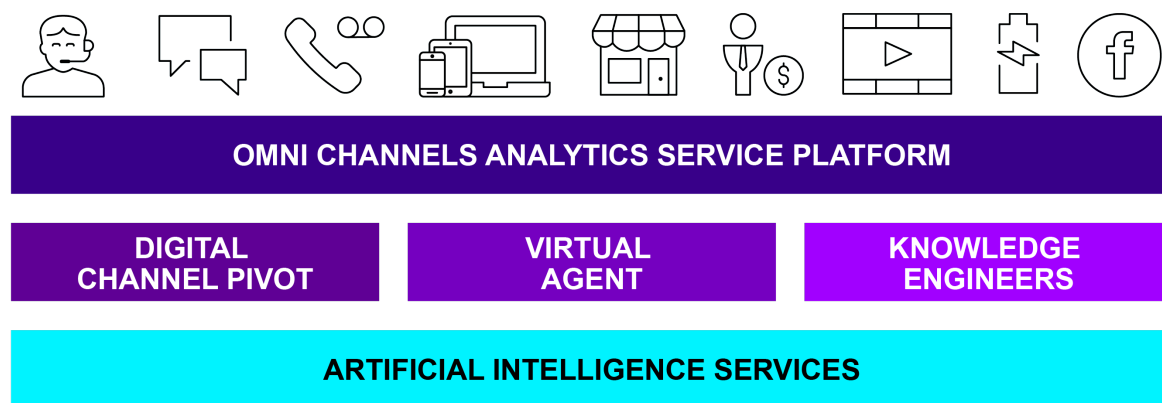


Figure 2 - Components of Reactive AI-Powered Care

- An omni-channel analytics services platform is used to analyze the effectiveness of the digital channel pivot and to inform the cognitive services with historic interaction details
- A digital channel pivot confirms the customer’s intent, determines the best channel for handling/resolving the customer’s intent and pivots the customer’s interaction accordingly
- A digital assistant (or Virtual Agent) leverages the AI services to execute back and forth conversations in natural language
- Knowledge engineers are highly specialized chat workers who serve three functions: AI trainer, AI enhancement designer and chat agent
- Artificial intelligence services are cognitive services that support the Digital Assistant in orchestrating a human-like and context-aware dialogue with the customer.

The use of AI in these ways can create many benefits, the first of which is a differentiated customer experience. CSPs can intelligently drive customers to digital experiences; provide conversational interactions through digital assistants, increase digital adoption and containment and ultimately eliminate calls to centers. AI agent-assist improves agent performance and customer outcomes. All of these factors

drive a reduction in the cost of customer care, often by more than 30%. Global corporations implementing intelligent customer care are seeing impressive results. When customers are offered the opportunity to opt into a digital experience instead of speaking to an agent, more than 30% of customers choose this route and as high as 85% give the experience a “thumbs up.” By intelligently driving customers to digital experiences, these organizations are reducing cost and providing customers with a positive experience and the significant benefit of resolving their issues more quickly.

1.2. Proactive Care

AI-powered proactive care focuses on detecting an issue in near real-time in order to try to solve it before the customer is aware or, alternatively, reach out proactively to inform the customer of the issue and its resolution plan.

Proactive care gathers data directly from both the network and the customer’s home and uses current contextual information to monitor service quality and identify and address immediate issues before the customer notices them. Within the proactive monitoring of the line, a predefined set of parameters is gathered from the customers’ home /devices and enriched with data from other backend/network system data. QPIs and KPIs are compared with predefined thresholds in order to continuously evaluate network performances and proactively identify issues on the line (e.g., CPE Health, VOICE issue, WAN issue, WLAN issue).

When an alarm is raised (e.g. a KPI is over threshold), a worker can proactively perform a network optimization action to try to solve the issue. Automated workers can trigger automatic recovery actions such as automatic reboots or firmware upgrades to avoid trouble ticket openings. CSPs can identify the most critical customers likely to perceive service degradation, proactively notify them of the issue and suggest “work around” options while the issue is being resolved. In instances where the CSP can’t preempt customer impact, front line agents have detailed information at their fingertips to address customers’ questions with accuracy when they call in.

Just as with reactive care, multiple components are needed to provide AI-powered proactive care (Figure 3).

KEY COMPONENTS



Figure 3 - Components of Proactive AI-Powered Care

- An **analytics layer** collects, processes and is the engine for visualization and proactive capabilities
- An **interface layer** displays the collected data for single and aggregated views, use cases, KPIs and alerts
- **Integration** with back-end systems allows for retrieving data to enable the use cases and

trigger recovery actions

- **Proactive use cases** utilize the huge amount of data and the power of the analytics layer

Through proactive care, CSPs can address customer claims with accurate information, prevent trouble tickets from being opened and reduce customer calls – all of which lead to reduced operational expenses, average handle time and overall time to repair. Better experiences lead to less customer churn, improved customer loyalty and build long-term customer relationships. Through proactive care one major European telecommunications company was able to identify 25% of new activated lines that were affected by service degradation, 90% of which were proactively resolved. Trouble tickets were reduced by 10% in the first period after line activation.

1.3. Predictive Care

Predictive care focuses on detecting and preventing potential issues before they occur or, if an issue does occur, accelerate the resolution to minimize impact. The issue could be at the customer or network location and activities to prevent disruption are carried out in the background.

Using a combination of customer and network data, insight based on AI and prediction capabilities help identify factors that are causing current problems or are likely to cause future problems. After a training phase, the machine learning algorithms run over processed data in order to assign a risk score to each line showing the likelihood that line will be affected by a specific type of issue (e.g. instable line, slow connection, etc.). Algorithms identify customers at the highest risk of issue (for a specific issue category), and trigger automatic recovery actions towards external systems, such as automatic reboots or firmware upgrades, to prevent an issue from occurring. Robotic Process Automation is used for preventive incident management, automated ticket resolution and assisted second level troubleshooting. These actions provide a seamless experience using analytics-driven operations, and an AI/ML powered robotics workforce to augment humans in network operations centers.

Components needed to provide AI-powered preventive care include AI-driven assurance, a network operations center and robotic incident management (Figure 4).

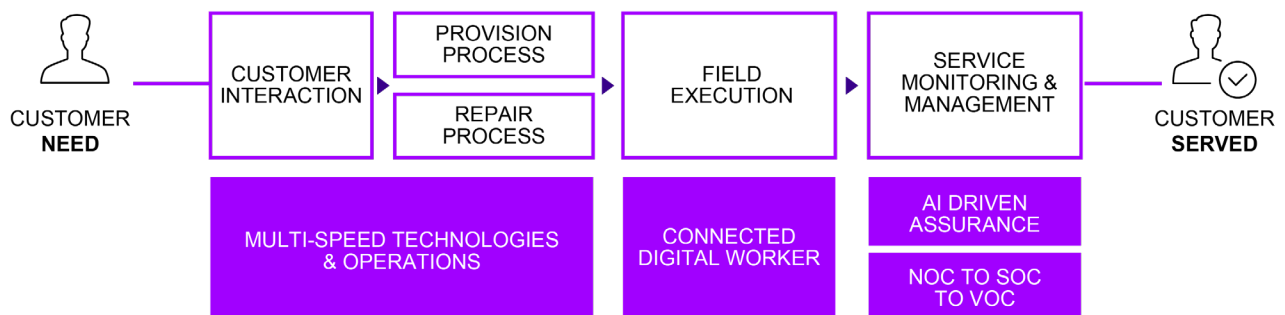


Figure 4 - Components of AI-Powered Preventive Care

- **AI-driven assurance** anchors intelligent data-driven operations around the customer, using a mixed workforce of humans, robots and AI entities to enhance customer experience and operational effectiveness.
- The **Network Operations Center (NOC)** evolves not just to a Service Operations Center (SOC) that facilitates service-centric operations, but all the way to a new digital services operations model that leverages analytics capabilities to capture the Voice of the Customer (VOC).
- **Robotic incident management** employs digital bots that recognize alarms and automatically process tickets until closure, interacting with humans through AI user interfaces.

Benefits of predictive care include increased operational efficiency and reduced operating cost. But importantly, the customer experience impact is substantial. Through predictive capabilities CSPs are resolving issues before they happen that could potentially impact hundreds of thousands of customers at a time. One organization implementing predictive node failure achieved an 18% reduction in level one supporting resources and a 10% reduction in customers calling due to technical issues.

Making the Shift Left

Artificial Intelligence will soon be the foundation for delivering a consistent, engaging customer experience. Right now its use across reactive, proactive and predictive care is differentiated. Soon it will be a strategic necessity. Virtual agents, trained by AI, will be deployed in service delivery. Issues will be identified in near real-time and measures will be initiated to proactively correct and mitigate customer impact. And, service issues will be greatly reduced through intelligent issue prevention.

There are no real dependencies between these three AI-enabled areas and the technological capability is available now to support them all. But, if the CSP makes inroads in reactive care, the savings produced can be used to “shift left” to proactive care. Then savings achieved in proactive care can be used to “shift left” once again and, through predictive care, preempt potential issues before they occur.

Abbreviations

AI	artificial intelligence
CSP	communications services provider
NOC	network operations center
SOC	service operations center
VOC	voice of the customer

Bibliography & References

¹ Gartner, Inc. *Hype cycle for CRM Customer Service and Customer Engagement*, 2017, July 2017.
Gartner, Inc. *Predicts 2018: CRM Customer Service and Customer Engagement*, December 2017.

Software-Defined Service Orchestration for MABR TV Services

A Technical Paper prepared for SCTE•ISBE by

Prabhu Navali

Engineering Manager
Ericsson Media Solutions
43 Nagog Park, Acton MA 01720
Prabhu.navali@ericsson.com

Raj Nair

VP, Technology
Ericsson Media Solutions
43 Nagog Park, Acton MA 01720
raj.nair@ericsson.com

Table of Contents

Title	Page Number
Table of Contents	2
1. Introduction.....	3
2. Software-Defined MABR TV Services Orchestration.....	3
2.1. SCTE Multicast ABR Reference Architecture	3
2.2. Drivers for the converged Multi-Cloud deployment of MABR TV Services	4
2.3. Converged Multi-Cloud deployment of MABR TV Services.....	5
2.4. Need for SDSO for MABR TV Services	5
2.5. Reference Model.....	7
2.5.1. Converged Multi-Cloud Deployment Reference Model	7
2.6. Overview of Multi-Cloud MABR TV Services	8
2.6.1. Operator On-Prem Private Cloud Components	9
2.6.2. Operator Public Cloud Components	10
2.6.3. Operator Edge Cloud Components	10
2.7. Unified Service Orchestration Infrastructure	11
2.8. Software-Defined Service Orchestration.....	12
Conclusion.....	13
Abbreviations	14
Bibliography & References.....	14

List of Figures

Title	Page Number
Figure 1 - SCTE Multicast ABR TV Services Reference Architecture	4
Figure 2 - MABR Services Multi-Cloud Deployment Reference Model	7
Figure 3 - Example Multi-Cloud MABR TV Services Deployment Components.....	8

1. Introduction

Pay TV operators are increasingly moving their existing traditional on-prem only infrastructure-based Multicast-assisted ABR (MABR) TV services to a new next generation distributed deployment architecture – a software defined, media optimized container applications based microservices deployed across multi-cloud (operators public cloud, operators edge cloud and operators on-prem private cloud) infrastructure. This paper explores the composition of such a multi-cloud-based deployment architecture and details software-defined service orchestration of such a distributed platform.

2. Software-Defined MABR TV Services Orchestration

Traditionally, PayTV operators have deployed the managed MABR TV services from operators national and regional data-centers using on-prem components - for video encoding/transcoding, packaging, content protection, user experience, subscriber, multicast data plane and control plane management components, etc. However, recent trends in the industry and standards have made way for these components to be deployed in a multi-cloud environment.

2.1. SCTE Multicast ABR Reference Architecture

Increasingly operators are deploying PayTV platforms in a hybrid deployment model in which some components are deployed on-prem data centers and some components are deployed in public or private cloud. Recognizing this growing trend SCTE WG7 working group has been updating the reference architecture for delivering the MABR TV services.

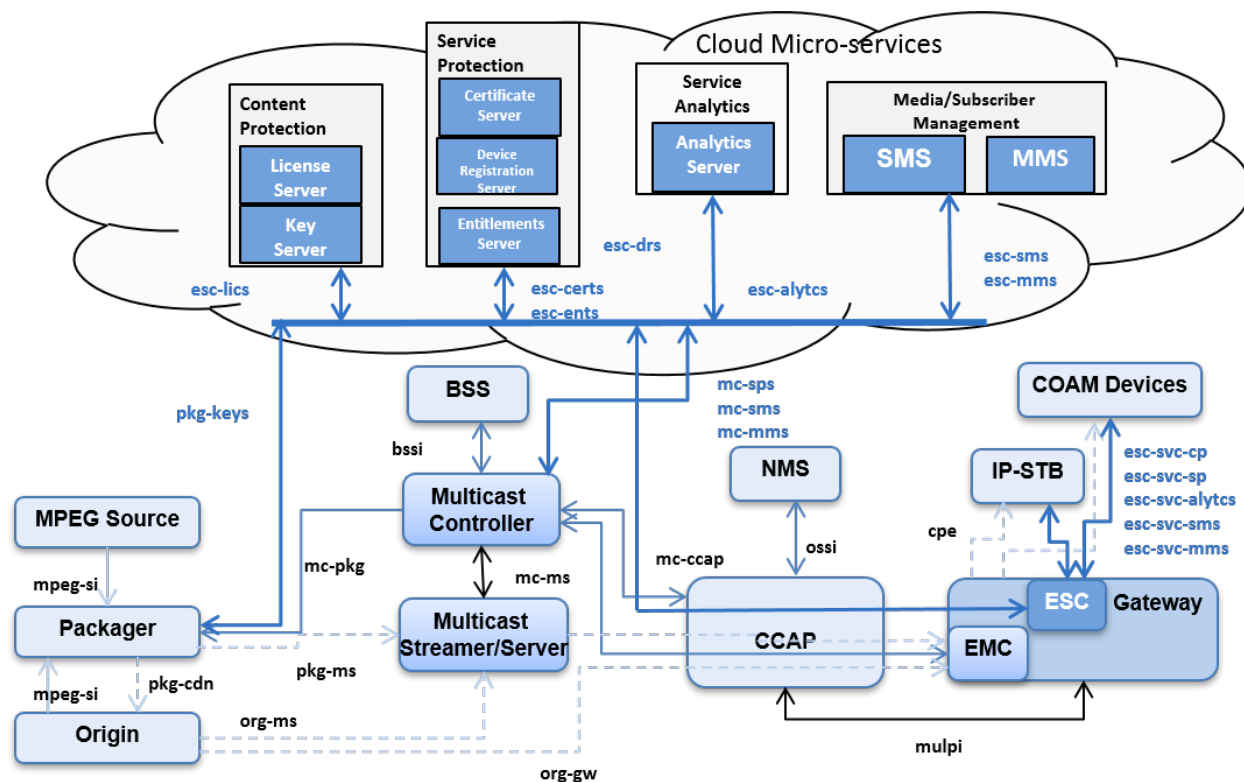


Figure 1 - SCTE Multicast ABR TV Services Reference Architecture

2.2. Drivers for the converged Multi-Cloud deployment of MABR TV Services

Rise of (operator) Public Cloud

Over the years with the rise of cloud computing, MABR TV services are being deployed in a hybrid deployment model - using (operator) public cloud and on-prem components in operator national/regional data-centers. (Operator) Public Cloud computing offers elastic compute/storage access allowing infinite elastic scaling, and utility-oriented usage. However, in order to scale the MABR TV Services workflows to millions of subscribers with good KPI metrics – there are challenges with the Hybrid deployment models. The real-time response and concurrency requirements are driving the cost of the cloud deployments. With ever increasing OTT clients (mobile and OTT STBs) there is a clear need to distribute the applications and workflows closer to the OTT clients on edge networks.

Rise of (Operator) Edge cloud

In order to address the latency and service scaling issues, operators are moving compute and storage resources to the edge network. The industry and standards are responding to this edge cloud evolution with development of architectures and standards [FOG, MEC].

Edge cloud extends the cloud computing to the edge network. Edge cloud facilitates the operation of computing, storage, and network services between OTT end points and MABR cloud services.

Rise of (Operator) On-Prem Private Cloud

In order to improve the efficiency of the compute, storage and network resources operators have been moving towards deploying private cloud infrastructures (like MaaS/OpenStack) in their national and regional data centers.

Rise of Containerized media optimized microservices applications

MABR TV Services are moving towards containerized microservice applications deployed over Container Clusters (like Kubernetes) over a cloud infrastructure.

Rise of DevOps Model for MABR TV Services deployment

Traditionally, for on-prem deployed MABR TV Services, features are developed and rolled out slowly. With the advent of DevOps model, the operators can roll out new features, patches more frequently. In addition, with cloud deployment scaling of the services can be done in short notice. The DevOps model benefits - TBD

The above changes are driving a paradigm shift in the way MABR TV services are being deployed - paving way for a new managed converged multi-cloud deployment of MABR TV Services.

2.3. Converged Multi-Cloud deployment of MABR TV Services

In the new paradigm MABR TV services are deployed over a converged multi-cloud deployment model.

The new deployment model includes MABR TV services deployed over multiple cloud infrastructures. This multi-cloud deployment can be vendor and operator co-owned and co-managed services model.

- Operator Public Cloud
 - Containerized microservices applications
- Operator On-Prem Private Cloud
 - Private cloud infrastructure in national and regional Data Centers
 - Containerized microservices applications
- Operator Edge Cloud
 - Private/public cloud in Edge Network (CO, Edge data centers, POP, MEC, etc.)
 - Containerized microservices applications
 - Intelligent edge caching/proxying microservice applications
- Multicast Core Network, Multicast Access Network, and Multicast in Edge Cloud
- Multicast-Assisted ABR – MABR (Multicast to Unicast conversion at the Edge).
- DevOps Model for MABR TV Services deployment.
- Microservices, Docker containers and Kubernetes clusters and orchestration
- Service discovery and scaling based on the requirements.
 - Service discovery and load-balancing frameworks
- ELK service monitoring and analytics services

2.4. Need for SDSO for MABR TV Services

The evolution of MABR TV Services deployment over a converged multi-cloud deployment model architecture needs microservices containers to be deployed across multiple public and private clouds.

The benefits of such a converged multi-cloud deployment are:

- Enormous scale
- Low latency response and improved KPIs
- Locality preserving
 - Local Ad insertion, personalization of TV services
- QOE management
 - Better user experience
- DevOps model
 - Speed of innovation and delivering new features/services to customers

The success of such a deployment depends on the ability to orchestrate these services across multiple public/private cloud infrastructures. The different clouds (public, on-prem private and edge cloud) may have different cloud infrastructures – for example, AWS, Azure, OpenStack, etc. Containerized applications allow one to deploy the same application across different cloud infrastructure. So, containerization normalizes the application deployment across different cloud infrastructure platforms. However, there is a good need for a software-defined service orchestration architecture to normalize and facilitate the orchestration MABR microservices.

2.5. Reference Model

2.5.1. Converged Multi-Cloud Deployment Reference Model

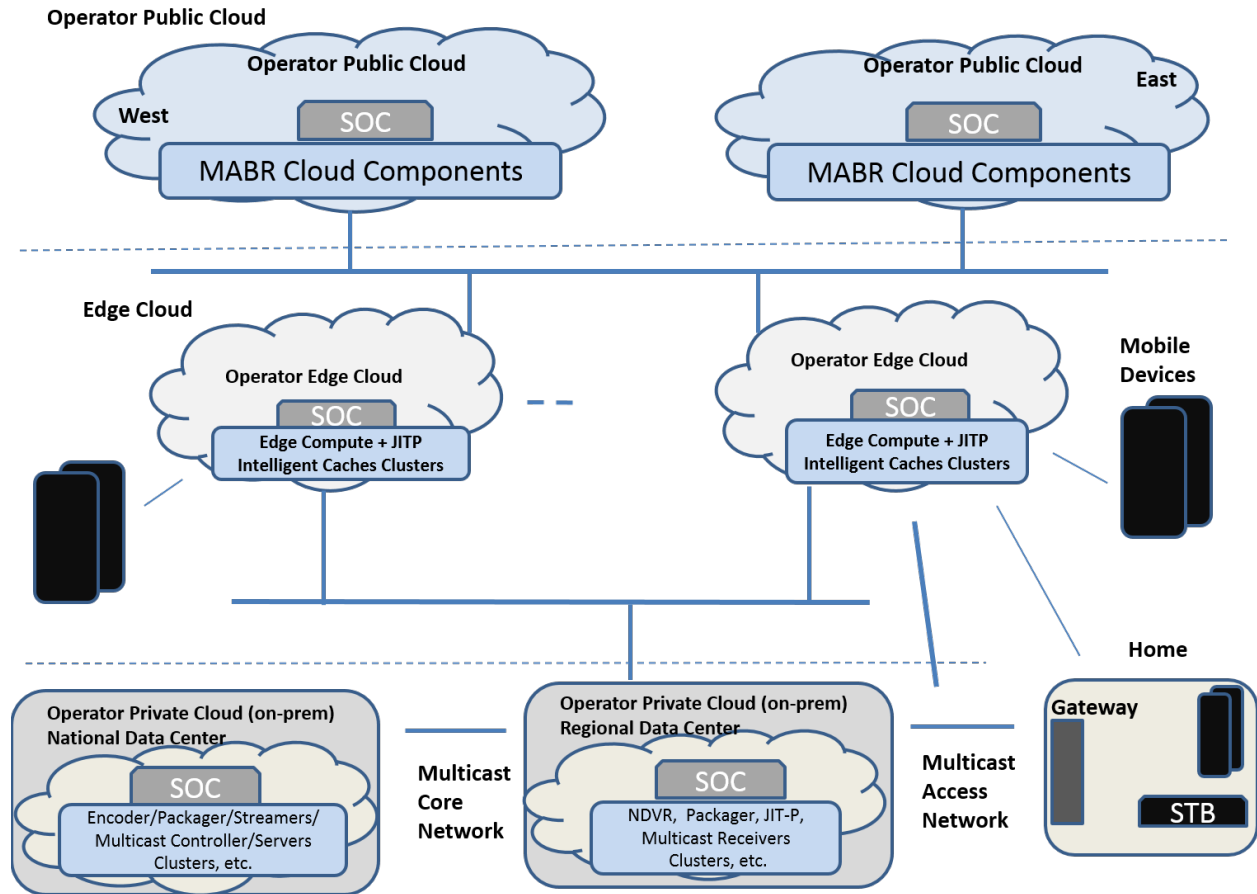


Figure 2 - MABR Services Multi-Cloud Deployment Reference Model

Figure 2 shows a reference deployment model for a distributed MABR TV Services. The MABR TV Services backend microservices components are distributed across public cloud, operator edge cloud and operator on-prem private cloud infrastructure. The services are made up of different layers or nodes that are distributed across different clouds.

The common denominator for these services is that this fabric made of SOC that distributes the resources and services of computation, communication control, storage across all the available devices, systems, storage, clusters, cluster managers across the multi-cloud environment.

The network of software-defined *Service Orchestration Controllers* orchestrates and continuous management of multi-cloud MABR TV Services deployment.

2.6. Overview of Multi-Cloud MABR TV Services

Main functionalities delivered by the Multi-cloud MABR TV services are Live/Linear TV and Tim-Shifted TV services. Figure below shows an example high-level functional block diagram of microservices container clusters deployed across multi-cloud environment to deliver those TV services. Note that this is an extension of the SCTE MABR reference architecture adopted to multi-cloud deployment model.

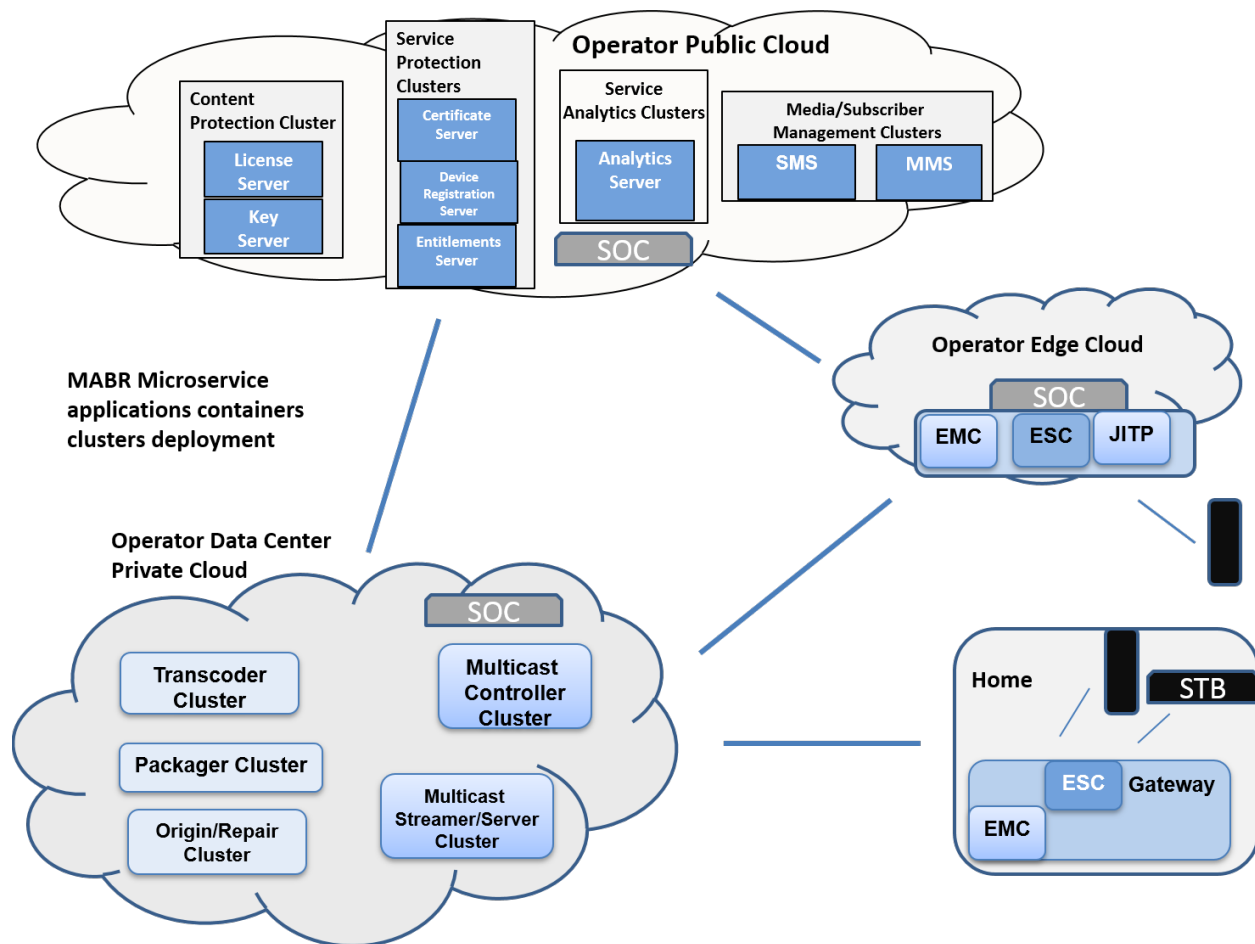


Figure 3 - Example Multi-Cloud MABR TV Services Deployment Components

- Channel Ingestion for Multicast ABR distribution

Involves in ingesting MPEG-TS streams and metadata into the Media Manager and Transcoder/Packager to generate Multi-bitrate multicast ABR MPEG-TS streams or ISOBMFF segments. The MABR streams or segments are encrypted with appropriate Common encryption schemes. The MABR streams are conditioned according to SCTE Adaptive Transport Stream (ATS) with Virtual Segmentation information for the downstream nodes.

- Multicast Data Plane
 - Multicast ABR Streams Distribution

In this low-latency option, the ABR streams are sent over as MPEG-TS multicast streams to the downstream nodes like gateways or edge media routers (EMR)). In addition to carrying the a/v elementary streams, these transport stream may carry associated CMAF/ISOBMFF metadata (no Mdat) in a separate metadata elementary stream. The hybrid Stream+File format is called Common Mezzanine Distribution Format (CMZF). The residential gateway or EMR downstream performs a transformatting function on the received elementary streams to generate HLS or DASH segments for ABR delivery. The ABR segments are delivered using a local HTTP server to OTT clients. The associated manifest is fetched from the origin server. The MABR transport streams can be RTP encapsulated and can use existing IPTV RTP repair mechanisms for repair/recovery and faster channel tune-in.

- Multicast ABR segments Distribution

In this option a file-based multicast protocol like NORM will be used to distribute the ABR segments over multicast. The Transcoder/Packager generates the ISOBMFF segments/manifests and are uploaded to origin server. Multicast server is used to distribute the ABR segments over NORM protocol. The downstream nodes like residential gateways with embedded multicast client (EMC) receive these segments over NORM and delivered using a local HTTP proxy cache to OTT clients. The origin/repair server will be used to perform repair and channel tune-in.

- Multicast Control Plane

The multicast control plane includes multicast controller and its interactions with other nodes in the multicast data plane to control the delivery of channels over multicast. Some of the functions include – content selection, channel map management, reporting and bandwidth control functions.

- MABR Service Control Plane

The service control plane includes backend components that are deployed in the public cloud and an intelligent cache or proxy or function in the residential gateway or in a node in edge cloud.

This control plane provides many backend functions for the MABR TV services delivery – like subscriber management, user experience workflows, media life-cycle management, media databases, time-shifted services, playback control services, content protection services, service protection services, service discovery, load balancing, service access control, license servers, key servers, analytics servers, OSS/BSS integration, EPG ingestion/management, etc.

The residential gateway or edge cloud node may include a embedded service control (ESC) component that includes intelligent caching and proxying functions for faster channel tune-in and to scale the MABR services. The caching and proxying component is used to cache/proxy for functions like - content protection, licenses, keys, service protection services, playback business rules (PBR), media and service entitlements, parental control.

Most of the MABR service components are deployed as media optimized containerized microservice applications in on-prem private cloud, public cloud, operator edge cloud environments and residential gateway component.

2.6.1. Operator On-Prem Private Cloud Components

The operator On-Prem hosts a number of components, mainly focused on the content preparation, transcoding, packaging and multicast distribution. Most of the multicast data plane components are deployed on on-prem private cloud. Most of the components will be deployed as containerized

microservices application clusters. Some of these components clusters can be realized either in edge cloud or public cloud as well.

Some of the example component clusters are:

1. Transcoding containers cluster
2. Packaging (JITP) containers cluster
3. Multicast Streamer containers cluster
4. Multicast Server containers cluster
5. Multicast repair server containers cluster
6. Faster Tune-In server containers cluster
7. Network DVR server containers cluster

2.6.2. Operator Public Cloud Components

A number of MABR Services components are deployed in the public cloud, mainly focused on the Control plane components like - user experience delivery components, Channel Catalog management components, Playback stream control components, Time-Shifted services workflows and data management components, subscriber management components, content protection (DRM) components, EPG ingest and integration components, OSS/BSS integration components, service protection components, and service discovery and delivery components.

Most of the multicast ABR services control plane components are deployed on the Public Cloud. Most of these components will be deployed as containerized microservices application clusters. Some of these components clusters can be realized either in edge cloud or on-prem private cloud as well.

2.6.3. Operator Edge Cloud Components

A number of MABR services components are deployed in the operator Edge Cloud mostly focused on the intelligent caching/proxying components for multicast data/control plane functions and caching components for service control plane functions. Edge cloud components participate in caching/proxying and deliver of ABR segments/manifests, caching/proxying and delivery of service control metadata like – channel map, channel data, metadata, subscriber/channel entitlements, licenses, keys, etc.

The intelligent caching/proxying components include:

- Embedded Multicast Client (EMC) functions
 - Join/Leave Multicast groups to receive multicast data
 - Receive Multicast streams or segments
 - Multicast to Unicast conversion and preparation of Unicast segments for delivery
 - HLS or DASH delivery of manifests/segments
 - Channel Map managements
- Embedded Service Control (ESC) functions
 - Scaling MABR backend Services
 - Faster Tune-In
 - Enable faster channel playback to reduce playback latency - for channel tune-in, for Channel surfing, for channel switching
 - Intelligent caches enable Faster Tune-In by caching the relevant information needed for a channel playback

- Caches information for one or more channels in the Channel Map
- Caches the following to enable Faster Tune-In
- Licenses (keys), program entitlements, entitlements, ratings, channel roll responses
- Service protection

Some of the example component clusters are:

1. EMC multicast receiver containers cluster
2. JITP Edge Packaging containers cluster
3. ESC service control containers cluster
4. Stream/Service personalization/localization server containers cluster
5. EMC/ESC reverse proxy service containers servers

2.7. Unified Service Orchestration Infrastructure

As mentioned in the above sections for a successful delivery of MABR services – microservice applications are deployed as container clusters across multiple cloud infrastructures. The cloud infrastructure and microservice infrastructure services are leveraged to deploy the application clusters.

Cloud infrastructure services

The public and private cloud infrastructure provides services like IaaS/PaaS to orchestrate compute/VM, storage and network resources in a cloud environment. Example cloud infrastructure services are AWS, Azure Cloud, MaaS, OpenStack, VMware ESXi, etc.

Microservice application containers

Container software provides a standardized (normalized) way of developing, deploying microservice applications. Containers can be lightweight and share the host OS resources. Containers can be more portable by packaging all the required packages and OS resources into the container. Containers are secure as the container engines provides resource isolation and access/resource security for the containers. Containers typically run on a container engine (like Docker Engine) on a host (VM or native compute). The containers are portable and can be run on different cloud infrastructures.

Container cluster infrastructure services

Microservice application container clusters are deployed over the cloud infrastructures using container application cluster management and control platforms like OpenShift, Kubernetes, Docker Swarm, etc. These platforms support various types of container software – like docker, Mesos, etc.

Service discovery

Containers and container cluster managers like Kubernetes work with *Service Discovery* frameworks like (Eureka, Consul, Vault etc.). These frameworks focus on providing service registration and discovery services. Containers and cluster managers integrate with *Service Discovery* frameworks to register the containers microservice with related metadata. A *Service Registry* allows a requester (another container application or external applications) to discover the service hosted on different container clusters. The service selection from the service registry is based on the application requirements and constraints. There

could be multiple service registries in a cloud, or one each in every cloud deployment, or a just one global *Service Registry* for the entire MABR services.

Service Access

Service Registry will be used to provide load-balancing across different microservice container clusters. A *Service Load Balancer* (like AWS ELB, NGINX, etc.) will request the available services (frequently updated/fetched) and will provide, update the service access features.

2.8. Software-Defined Service Orchestration

This section (and article) proposes a distributed *Software-Defined Service Orchestration* framework that distributes the MABR microservices throughout the multi-cloud environments. There is a clear need for such a framework given the nature of the microservices configuration, deployment constraints and different cloud infrastructure and microservices environments. The proposed framework allows such a distributed platform to perform orchestration of microservices using available nodes, containers, container clusters, cluster managers, service discovery platforms, storage, and networking resources. The distributed platform is a network of *Service Orchestration Controllers* that facilitates the orchestration tasks. The Distributed *Service Orchestration Controllers* (DSOCs) glue together the service container clusters that are running across multiple clouds (Public Cloud, Edge Cloud, Operator Private Cloud).

Microservice applications specify in the form of *Directives* high level tasks, configuration and requirements for the successful deployment and operation of the service.

These *Directives* are carried in *Service Manifests*. A *Service Manifest* is a document that carries the instructions on how to deploy, delivery, manage, operate a service. The *Service Manifest* is an XML/YAML/JSON based document and contains various *Directives*.

Some of the examples of *Directives* are:

1. Service application configuration data
2. Service graph/chart for deployment
3. Service constraints - CPUs, Bandwidth, Storage, etc.
4. Service cluster configuration data
 - a. Auto scaling metrics, constraints, load-balancing constraints
5. Service access related data
6. Service delivery related data
7. Service monitoring, analytics related data

Service Orchestration Controller (SOC)

A distributed service orchestration platform is a network of *Service Orchestration Controllers* (SOCs) deployed in different cloud platforms. These SOC_s distribute *Service Manifests* and perform the functions based on the *Directives* defined in the manifests. The global topology configuration of controllers allows SOC_s to discover each other. Typically, each SOC is paired up with a local or global *Service Discovery* service through which it discovers and monitors and facilitates orchestration of application containers via container cluster managers (like Kubernetes). The distributed SOC_s manage the life-cycle of the microservices applications containers. The SOC includes a *Service Manifest Controller* which handles functions related to creating, updating, distributing, and receiving *Service Manifests*.

Some of the main functions of *SOC* are:

- Handle submission of a High-Level Service orchestration *Task*
 - Submission of a high-level request triggers the preparation of *Service Manifest* with *Directives* to fulfill the orchestration request
 - Service Discovery
 - This typically involves in the discovery of the service components across multi-cloud environment and determining the existing service components, constraints, infrastructure components, clusters etc. available to make adjustments to *Service Manifest Directives* to fulfill the new orchestration request.
- *Service Manifest Controller* functions
 - creating, updating, distributing, and receiving *Service Manifests*
- Handling *Service Directives*
 - Decomposition of *Directives* into smaller *Orchestration Tasks*
- Communicating *Orchestration Tasks* to Orchestration Managers (like Kubernetes)
- Communicating Orchestration request status, monitoring, etc.

The distributed *SOCs* handle several *Tasks* to setup, teardown or update the microservices containers based on the MABR service business needs. A high-level service orchestration request example like - Live/Linear channels catalog ingestion - will trigger an *Orchestration Task*, with a customized *Service Manifest* with the appropriate *Directives*, to fulfill the request will be created at the public cloud *SOC* and distributed to other *SOCs* in the multi-cloud deployment. The *SOCs* in each cloud environment work with the Cluster Managers to fulfill the *Task* request – which includes service discovery, resource allocation, generating local orchestration charts/manifests to orchestrate the containers in different microservice clusters. The *SOCs* send updates on the progress of the *Task*. The distributed *SOCs* continue work together to handle any *Task* requests for setting up or tearing down or updating the microservices containers deployed across multi-cloud environment.

Conclusion

In this paper we propose a framework for a distributed *Software-Defined Service Orchestration* framework for the orchestration of MABR microservices applications across multi-cloud environments over disparate cloud platforms and normalized container cluster platforms. The MABR Microservices are deployed across Public Cloud, Operator Edge Cloud and Operator Private Cloud infrastructures. The distributed *Software-Defined Service Orchestration* framework provides mechanisms to distribute high-level orchestration requests (*Tasks*) into actionable *Directives* using *Service Manifests* upon which the local, global Cluster Managers, service discovery components and *SOCs* can act on. The *SOCs* in conjunction with local Cluster Managers orchestrate the application microservice containers and manage the life cycle of microservices application containers.

There are industry trends in both moving the microservice applications to Edge Cloud for a multi-cloud deployment model, and software-defined orchestration frameworks to facilitate such a distributed application deployment.

We think there is a need for standardization of such a distributed software-defined MABR microservices orchestration framework.

Abbreviations

ABR	Adaptive bit rate
MABR	Multicast Assisted ABR
NORM	Nack-Oriented Reliable Multicast
SCTE	Society of Cable Telecommunications Engineers
SOC	Service Orchestration Controller
SDSO	Software-Defined Service Orchestration

Bibliography & References

Cablelabs Multicast-assisted ABR Technical Report.

SCTE Multicast ABR Services part -1: Data Plane Workflows Specification – SCTE DVS WG7 (under progress)

SCTE Multicast ABR Services part -1: Service Control Workflows Specification – SCTE DVS WG7 (under progress)

OpenFog Architecture Overview – White Paper, OpenFog Consortium Working Group

ISO/IEC 23009-1 - DASH Specification

RFC 8216 HTTP Live Streaming (HLS) Specification

ISO/IEC 13818-1 MPEG-2 Transport Streams Specification

ISO/IEC 13818-1 Edition 7 PDAM 1 – Carriage of associated CMAF boxes for audio-visual elementary streams in MPEG-2 TS. (under progress)

Solving All Our Problems... Sort of...

Blockchain Integrity, Security, and Reliability for Cable Use Cases

A Technical Paper prepared for SCTE•ISBE by

Steve Goeringer

Principal Security Architect

CableLabs

858 Coal Creek Circle, Louisville CO 80027

s.goeringer@cablelabs.com

Dr. Jason Rupe

Principal Architect, IEEE Blockchain Initiative Co-Chair

CableLabs

858 Coal Creek Circle, Louisville CO 80027

j.rupe@cablelabs.com, jrupe@ieee.org

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
What Good Are Distributed Ledgers and Blockchains, Anyway?	4
1. Authoritative history.....	4
2. Identity management and anonymity	4
3. Event synchronization	5
4. Traffic flow management and message flow.....	5
5. Information reliability	6
Use Case Summary	6
1. New and direct revenue	6
2. Cost optimization.....	7
3. Customer experience	9
4. Reduce ecosystem friction	10
Complex Security and Reliability Design Concerns.....	11
1. Identity, transaction authentication, and transaction authenticity	11
2. Distribution and redundancy	11
3. Network scale and performance considerations	12
4. Governance and code management.....	13
5. Attack Vectors as a reliability problem	13
6. What is the meaning of reliability in the context of blockchain?.....	14
Conclusion.....	14
Abbreviations	15
Bibliography & References.....	15

Introduction

This paper surveys and categorizes blockchain use cases relevant to the cable ecosystem and then discusses key design factors in implementing appropriate blockchains. But first, it explains some basic principles and concepts of blockchains and distributed ledgers. The paper ends with some discussion about important concepts regarding these use cases, many of which are unique to blockchain.

Blockchain networks and distributed ledgers are explained through five important concepts:

- Authoritative history – blockchain networks have the ability to keep track of history with high integrity.
- Identity management and anonymity – blockchains are decoupled from identity management, but may require it.
- Event synchronization – distributed ledgers in a blockchain network can keep track of the order of events.
- Traffic flow measurement and message flow – blockchain networks can be used to enhance networking overall.
- Information reliability – once in the blockchain network, information is very difficult to change; but until it's in there, anything could happen.

The use cases explored in this paper can be organized into these categories:

- New and direct revenue – ways operators can generate new sources of revenue from this technology.
- Cost optimization – techniques to reduce costs of providing services using blockchain.
- Customer experience – methods to enhance the customer experience through use of blockchain.
- Reduce ecosystem friction – applications of blockchain that simplify what is otherwise complex.

To support these use cases, several important security and reliability concerns must be addressed:

- Identity, authentication, and authenticity of transactions and participants.
- Node and ledger distribution and redundancy.
- Network scale and performance considerations.
- Governance and code management.
- Attack vectors as a reliability problem.
- The meaning of reliability in the context of blockchains and distributed ledgers.

Readers are assumed to be familiar with blockchain topics. Background information may be found in [3]. This paper uses both the term blockchain and the phrase distributed ledgers. This use may seem synonymous; however, the authors view blockchain as a cryptographic network technology that produces distributed ledgers. This is an important distinction. Not all use cases require the relatively heavy assurances that blockchains entail but may still benefit from a distributed ledger.

What Good Are Distributed Ledgers and Blockchains, Anyway?

Before discussing use cases, it is useful to revisit the benefits of blockchains. This allows better determination of both how to design a blockchain for a given use case and whether a distributed ledger is necessary for that use case. A comprehensive approach for determining blockchain suitability for a given application is outlined by Scriber in [2].

Fundamentally, blockchains help to achieve security by design by providing highly secure logs of transactions. These logs are highly secure because they are distributed amongst many participants (making them hard to change without participants noticing), because the transactions they include have integrity assurances (signatures built in), and sometimes because they include additional cryptographic protections. However, they are not a replacement for other security controls. Distributed ledgers typically allow us to attest truth, not assert truth.

1. Authoritative history

Blockchains provide layers of integrity, as discussed by Goeringer in [1]. Typically, transactions are signed by clients when submitted to the blockchain. Transactions may include data elements (or anything that can be converted to a data element) or executable code (smart contracts) that may have additional integrity protections. Transactions (and their associated smart contracts) are validated by one or more types of nodes when they receive the transactions. These transactions are compiled into blocks, often using a process called a Merkle tree which really produces a hash of all the transactions included in the block. Blocks themselves are digitally signed by their processors (miners in Bitcoin and Ethereum, validators in HyperLedger). A given block includes the hash of the proceeding block in the blockchain (that's the feature that makes it a chain). The result of all these layers of integrity is a data structure that creates facts. We know that a transaction was submitted and that the contents of the transaction are (nearly) irrefutable. If transactions are signed (for example, using asymmetric keys), we can prove a given client address submitted the transaction.

Integrity seems intuitively useful, but it can be transformational. Many businesses, including multi-service operators, are complex. Network operations and service delivery can include many stakeholders, some independently responsible for their own profits and losses. Moreover, new business models may leverage multiple businesses in service delivery (in a network context, this may include multi-tenant virtualization). Inevitably, each stakeholder will maintain their own service records. Different measurements or even errors in those records may make data reconciliation difficult or even impossible.

A distributed ledger with high transaction integrity can provide an authoritative history of transactions. This can streamline business operations, with particular impact on processes impacting compliance (such as privacy management). But how do we enhance transaction integrity?

2. Identity management and anonymity

Some security pundits have suggested that blockchains remove the need for identity management. Many articles have stated that blockchains create trust. Neither of these statements is true.

Blockchains can be designed to support anonymous transaction submissions, or transactions decoupled from attribution. They can also be designed such that each transaction is submitted only by credentialed clients. For that matter, transactions need not be sourced by or attributed to a person or organization. A

given blockchain platform, in fact, may support either option. The fundamental question is whether a given use case requires the identity of submitters or not. If it does, a process must be applied to issue identity to submitters. This identity may, for instance, be applied to a wallet in the case of a digital currency or security. In this way, blockchains don't eliminate the need for identity management; rather, they consume or rely upon identity management. In contrast, if a given use case doesn't require the identity of submitters to be known (as, for example, in the case of Bitcoin), then identity isn't needed because the use case doesn't require it. The use of blockchains does not protect or remove identity.

In considering what this means for trust, it is important to reconsider the previous section again. Blockchains allow us to create histories of transactions in which submitted transactions can be treated as a fact. Rather than creating trust between entities, blockchains create a data structure in which entities can trust. If applicable to a given use case, this actually removes the need for trust.

However, we must be careful. The fact of the transaction does not mean the submitted transaction itself is accurate. It may have been tampered before it was signed. Or the terms submitted by one party are not actually the terms agreed (offline); the transaction submission did occur, but it is not accurate to what the parties negotiated. The point being that accuracy, authentication, and authorization features must be designed-in using identity management practices related, typically, to public key infrastructure solutions such as cryptographic signing and co-signing. Such systems cannot really be anonymous, of course.

3. Event synchronization

A byproduct of how blockchains work is that the integrity base they enable includes a high confidence ordering of events to a certain degree of granularity. The granularity relates to the length of time it takes to process a block. The result is that transactions included in different blocks can be time ordered with extreme confidence (e.g., ordinality can be treated as fact)¹. For events (transactions) that are encoded within a single block, at least some uncertainty is introduced. Time (hacks) can be added to signed transactions to provide some level of confidence in event timing, but this returns us to traditional challenges in event timing (uncertainty in time distribution protocol, inaccurate or even malicious time codes by clients, etc.).

Moreover, as most blockchain implementations are distributed, high confidence event ordering can be visible to all stakeholders. The result is a highly useful mechanism for synchronization events in multi-stakeholder or similarly complex execution environments. This can be beneficial to managing workflows, tracking fulfillment, and possibly ensuring audit-ability for compliance or audit purposes.

4. Traffic flow management and message flow

Many use cases for distributed ledgers require use of a single kind of transaction for a simple, or at least consistent, purpose. This is the case, for example, with Bitcoin: all Bitcoin does, really, is track the distribution of spendable transactions between parties. However, more elaborate workflows can be developed that provide much more interesting capabilities. This can be done by programming specific behaviors at various kinds of clients. For example, different clients can be coded to process transactions in different ways. Another option, not mutually exclusive, is to use smart contracts. Smart contracts can be implemented in several ways, but the common approach allows conditional execution of transactions based on information provided in the transactions, including the identity of the parties the transaction goes between.

¹ A fork in a chain can record a conflict in this ordinality but should be considered a temporary anomaly as one possible reality will eventually be accepted as fact.

This allows blockchain networks to be means of transport for complex and conditional information between stakeholders. Moreover, complex rules can be applied against transaction transport. Unlike other transmission protocols (such as IP), strong integrity is designed into every transaction, individually and collectively, so very strong traffic and message flow reliability and security can be ensured.

5. Information reliability

Reliable information in the blockchain network will remain reliable, but unreliable information can equally be locked into a blockchain network's version of reality.

If reliable information enters the blockchain and propagates sufficiently, it will remain highly reliable as long as the blockchain is reliable. Once in the blockchain, the information is propagated through the network, and the result is that there is a large number of duplicate records of the transactions. When the nodes of the blockchain network are largely independent in risk and attack, the information in the system is going to be immutable with a high degree of certainty.

Nothing in blockchain architecture assures that information entered into the network is reliable, but there are ways to add some amount of assurance. Contracts can be checked for consistency with other code on the blockchain, and for reliable executability. Sources of information can be authorized and authenticated. Ownership and authority can be checked with information stored on other secured networks if not within the given blockchain network. But without these additional measures, unreliable information can be stored on a blockchain network easily.

There are known attack methods for blockchain networks which can introduce contradictory information, the equivalent of a double-spend in a smart contract implementation. But these attack vectors are very difficult to exploit as long as the system is kept in balance.

Many consider blockchain synonymous with permissionless systems, but that is not the case. From the perspective of information reliability, for a given use case, the full spectrum of permissioned to permissionless systems should be a design consideration [5].

Use Case Summary

Designing blockchains must be approached from the context of the use cases that may benefit from decentralized ledgers. There are literally thousands of blockchain use cases under various phases of development, but only a few relate well to the cable industry. These can be discussed concisely when organized into four categories: new and direct revenue generation, cost optimization, customer experience, and reduction of ecosystem friction. These are briefly discussed in the following subsections. Of course, there is overlap, and a given use case may apply to more than one categorization. Further, there will certainly be use cases that defy these categories. Readers are invited to use their own creativity, and to use the ideas here liberally.

1. New and direct revenue

Significant effort has been focused on whether distributed ledgers provide the basis for operators to enable new services or new markets. Digital currencies or securities may also allow generation of capital through initial coin offerings. Moreover, digital currencies may provide lower cost approaches to bi-directional transaction flows, supporting loyalty and reward programs. Some use case examples:

- Content focused coin offering – One of the most tangible digital assets the cable industry works with is, of course, content in the form of movies, television, and music. Tethering digital assets to

some form of digital currency may provide the basis for new payment models, including capital generation if a new security is generated using an initial coin offering (ICO).

- Games and eSports coin offering – Traditional content is not the only option available to operators. Many have trialed eSport and game related service offerings. Tying an eSport or game to a digital currency provides interesting options, including the opportunity to provide digital assets and new ways of adding non-traditional services into triple- and quad-play bundles.
- Digital goods provenance – Sales, fulfillment, and delivery of digital goods remain very attractive. Some operators offer opportunity to buy rights to digital assets. Use of a distributed ledger provides opportunity for digital goods provenance which may streamline ecosystem operations and provide an improved basis for trust to new entrants into this space. Operators enabling such capabilities may open new revenue opportunities while securing their roles in digital distribution in the future.
- Secure digital media – Related to provenance, providing secure digital media solutions in itself may provide value sufficient for revenue generation. This may be particularly true for user generated content where current methods of ownership assertion are insufficient. Blockchain technologies may also provide an opportunity to disrupt value chains on digital rights management.
- New model for ad revenue – The current advertising technology market is plagued with a variety of fraud and other security problems. It is also very complicated. Distributed ledgers that integrate ad delivery solutions and payment methods may streamline advertising technology while also providing better value to publishers, content owners, marketing firms, and advertisers through improved (but controlled) transparency.
- Multi-party billing – For some markets and market segments, the cost of participating in cable services can be perceived as very high. Providing a blockchain-based billing solution that allows multi-party billing to complex households, various forms of multi-tenant housing, and college campuses may allow much greater participation from those markets while ensuring the operator does get paid for service. This may apply particularly well to wireless access environments.
- Blockchain as a Service – It may not be practical for an operator to address all the potential new revenue that blockchain-based approaches may enable. Fortunately, blockchains can be designed as service platforms, able to support a wide range of transactions. Such blockchains can themselves be offered as a service.

2. Cost optimization

Multi-service operators are complex businesses providing a wide range of services over highly varied and also complex infrastructure. Any given access solution may have multiple stakeholders (CPE operations, access operations, access engineering, product management, security). Any given access solution may integrate perhaps dozens of vendors resulting in a wide range of interoperability and integration challenges. All this complexity inevitably leads to at least some inefficiency which means higher cost per unit served. Application of distributed ledgers may provide new ways to optimize service costs. Some cost optimization examples:

- Virtualization orchestration – Network function virtualization (NFV) is largely about achieving disruptive cost reduction by allowing use of general-purpose server infrastructure rather than “big iron” routing, switching, and CMTS solutions. It is also believed that NFV may support new information and computer technology business partnerships through multi-tenant and even multi-operator solutions. Orchestrating complex service chains may require authoritative history for billing purposes, strong identity management to prevent service theft, and event synchronization.
- Service authentication – Cable services have traditionally been largely premised focused: a given address is subscribed to a given bandwidth and set of features, and that’s that. As we move more

and more to over the top delivery, and households become more complex, more flexibility is necessary; but this has proven complicated. Improved ability to maintain histories between business units, coupled with tools to provide more complex traffic and message flow management, may provide better tools for service authentication.

- Dynamic service creation or provisioning (announce, publish, subscribe) – Traditional service creation and provisioning have been very manual, partially simply to double check all the records and tickets from multiple stake-holders. Blockchains may provide better tools for synchronizing service creation (possibly enabling full automation which has been an elusive goal for decades). Moreover, the distributed ledgers may provide completely new methods of coordinating provisioning activities, enabling much more flexible programmatic service delivery.
- Connectivity negotiation or transaction management – Smart contracts provide new ways to track customer opt-ins for service. Much more granular service agreements may be achieved through pervasive accounting and tracking of user agreements.
- Enhanced content protection – Long- and short-form content are experiencing serious piracy today, with significant impacts on the revenue of both content owners and cable operators. Moreover, ad fraud impacts the profitability of the entire ad tech industry. Even user generated content faces challenges and end users rely basically on the good will of the various services and sites that allow users to share their content. Blockchains provide new ways to assert ownership, track usage, and assert digital rights on content. Enhanced content protection may provide significant cost benefits to all content owners and integrity.
- Provenance – Supply chain integrity remains challenging, and particularly so in the realm of software. Development operations provides methods of achieving live builds and agile service delivery. However, it relies largely on both open and proprietary code dependencies that are hard to track. Blockchains may provide new ways to synchronize software builds, deconflict dependencies, and track both changes to codes and also who made those changes. All with unprecedented integrity.
- Scalable IoT – IoT is resulting in massive deployment of both independent, standalone devices and also intricate autonomous systems that blend IoT sensors and actuators with big data services. The result is explosive growth of managed and unmanaged deployment of devices to homes, businesses, enterprise, campuses, and communities. Manual processes cannot track and manage how all of these components will interact and interoperate. Operators need more dynamic security controls, more flexible on-boarding, adaptive service contracts, and new payment methods to deal with this growth cost effectively. IoT scalability will depend on all five of the benefit areas discussed above.
- Reputation-based authentication – Many large ecosystem operators (e.g., Apple, Google, Amazon, Samsung, etc.) and industry consortia (Open Connectivity Foundation) are working to ensure devices offered within their scope can securely access services. However, visibility between systems is minimal, and not all services and devices are part of these large ecosystems. Many solutions are developing to identify or finger print devices. Distributed ledgers may provide a common resource in home, business, and access networks to record device behaviors and apply trust decisions on authentication and network access. This provides the opportunity to assert reasonable network hygiene and keep costs of managing network security low.
- Media storage consolidation – Currently, most operator contracts for Video on Demand (VoD) services require an individual media copy for every concurrent use or view of the media. So, for example, if an operator wants to provide “Deathly Hallows” on demand to up to 10,000 users at a time, they may have to store up to 10,000 copies of that media on their servers. Distributed ledgers provide the opportunity to reduce the need for trust through better identity management and record keeping. If media owners are made more comfortable by removing the need for trust,

the ability to reduce the number of stored copies of media provide the opportunity for massive cost reduction in storage.

- VoD evolution – Similarly to media storage consolidation, usage rights available to operators can be very restrictive. Again, this is largely due to the need for cost prohibitive trust solutions that simply have not been possible before. Distributed ledgers provide for much more secure and visible transaction management that may provide for much more engaging use experiences while maintaining equity amongst the stake holders (operators, studios, content aggregators, subscribers).
- NFV Management – Orchestration of NFV-based service delivery includes many, many distributed components. Moreover, multiple stakeholders may be engaged (multiple tenants, multiple operators). Event synchronization and tracking using traditional mechanisms may be very difficult. Distributed ledgers may provide much more streamlined orchestration.

3. Customer experience

The benefits of blockchains described previously can provide the basis for streamlined assurance of customer experience. Moreover, the fundamental capabilities of distributed ledgers provide the opportunity to evolve customer experience. This is, of course, challenging, and so the list of examples is correspondingly less than shown for new revenue and cost optimization. Here are four:

- Customer preference tracking – Managing customer preference choices across multiple platforms can be challenging. Moreover, subscribers desire service mobility. And, all their choices are subject to privacy considerations. Distributed ledgers provide the opportunity for streamlined, seamless customer engagement. Also, distributed ledgers provide the opportunity to leverage crypto currency solutions, and so customer preference choices can be more easily coupled to billing.
- Customer loyalty activities – In many markets, the cost of cable-based services can seem very high. Lowering the cost of cable service (both actual and perceived costs) while enabling alternative revenues may be helpful to many subscribers. This can be realized through various customer loyalty activities tracked through distributed ledgers. This can include discounts for ad watching, credits for customer referrals or service recommendations, 3rd party partnerships, and multi-payer households.
- Customer as content provider – User generated content transforms entertainment from storytelling to story sharing – from a passive consumption of presented content to generation and sharing of our own stories. YouTube and Facebook are, of course, the epitome of current user experiences in customer generated content. However, both services make tough compromises in allowing users to control and own the content they share. Distributed ledgers provide the opportunity for users to register and assert ownership rights in ways that have not been possible previous. And, because of the strong ability to secure the records of transactions, operators that enable new content sharing options to users can enable entirely new experiences and control to users in how they share and distribute their content.
- Media sharing – Allowing users to share content amongst themselves has been problematic in many ways. Consequently, license rights on distribution simply have not allowed consumers to share media. The authoritative history provided by a distributed ledger, coupled with strong identity management, may provide the basis to enable media sharing. Event synchronization, coupled with complex transaction flows, can allow very intricate user experiences that improve value and increase engagement among communities of subscribers.

4. Reduce ecosystem friction

The ability to reduce complex transactions to a matter of fact provides a new basis for trust between stake holders. This provides the opportunity to reinvent entire industries. That does sound audacious. However, we have seen fundamental disruptions in transportation (Uber, Lyft) and hospitality (AirBnB). Why shouldn't cable experience similar transformation? Here are five examples:

- **Distributed trust** – Public Key Infrastructure solutions remain one of the most scalable tools to assert identity management across all the evolving ecosystems that comprise the world of information and computer technology, and the related emerging area of IoT. However, PKI is complicated and its use introduces challenging supply chain risks in the identity supply chain. The result is that many companies and organizations are all pursuing development of independent PKI roots (the foundational private key that attests the identity of all the certificates in that ecosystem). Unfortunately, bridging PKI roots is complicated and can introduce additional security risk. Leveraging distributed ledgers to orchestrate certificate issuance may provide a highly secure (reliable, high integrity) means of allowing different ecosystems make trust decisions relying on certificates from other ecosystems.
- **Content distribution convergence** – Several of the ideas above addressed digital transformation of media distribution, usually in the form of video (movies, TV). However, what works for movies might work for books, audio, music, and maybe the evolution of these media to AR and VR. If so, this provides the opportunity for operators to enter other content markets, or to provide new value to those markets.
- **Royalty management and reconciliation** – One of the hardest entertainment industry challenges is ensuring all the contributors to great content get what they are owed. Royalty management and reconciliation have traditionally reduced to rule of thumb-based estimates that may, or may not, have any basis on reality. Distributed ledgers, with or without smart contract capability, may provide cost effective ways to create authoritative histories of distribution and viewing that allow complete transformation of royalty rights management.
- **Customer as content provider** – Cable service providers are in a great position to assure, from the edge through to the core, that a person who creates content can have assured ownership of that content. Using a blockchain network, a customer who creates content can assure the content is encoded into the block, and therefore securing ownership of that content. Equally, they can transact that content in various ways, perhaps transferring ownership. The service provider can provide the blockchain-based solution to secure intellectual property for the customers, thus providing evidence of invention, creation, and ownership. Of course, this is a double-edged sword: a customer who plagiarizes will equally lock the evidence in the same process that protects intellectual property.
- **Media sharing** – In a manner like described above, media can be shared and permissions managed on a blockchain network. The distinction here is that rather than manage all the details of a relationship as needed when the customer becomes the content provider, this use case is simplified so that media can be shared as the user intends, without ownership transfer, contracts, or asset exchange beyond the sharing of initial content.

Complex Security and Reliability Design Concerns

1. Identity, transaction authentication, and transaction authenticity

There is an important subtlety when considering the idea that transactions become facts when recorded on a distributed ledger using blockchains. The fact is that the transaction on the distributed ledger is the transaction that the client submitted. That's it. However, that does not mean that what the client submitted was what they intended to submit, nor whether the transaction submitted was what another client expected or agreed to. Moreover, without some access controls, the transaction submitted to the blockchain network may not actually be what the client sent.

Consequently, specific use cases may need additional security controls added. Two functional areas to consider are identity and authenticity.

Nearly all blockchains use some form of asymmetric key pair to prove ownership of a transaction. A private key is used to sign the transaction; a public key (often included in the transaction or even used as a transaction identifier) is used to decrypt the signature and prove that the transaction is authentic. However, this can be anonymous. If it is important that a given transaction be authenticated and authorized prior to inclusion on a blockchain, identity should (must, really) be issued by an identified authority. It may be possible to use some form of distributed organization to issue identity, but most commonly a PKI certificate authority is used. Then, whenever a client submits a transaction, the transaction will be signed by the client and will include its PKI certificate which is in turn signed by the PKI authorities. This provides a strong basis to attest identity of clients.

How can we assure that what one client submits is what another client has agreed? There are several means, but one is to co-sign the transaction. A signed transaction can be provided by one client to another who then can review and accept the transaction, sign it, and submit it to the blockchain network. Alternatively, both clients can submit transactions, and a validator of some sort can ensure they match prior to approving the transaction for a block. And, of course, some form of smart contract can be used.

Architecture can matter a great deal. For example, it may not be feasible to adequately secure keys for a client on subscriber owned devices. If those keys can be accessed or manipulated, transaction identities cannot be assured, and therefore authentication and authenticity are at risk. So, it may be desirable to use a proxy for the end client (for example, deploying the users' "wallets" on the cloud).

A final note on identity management is warranted. Many architects and solutions providers are attempting to use blockchains as an alternative to strong identity management. Identity must be attestable in some way. This usually requires some type of central authority (such as a certificate authority in a PKI). It may be possible to make some level of trust decisions based on behaviors of clients recorded by peers on a blockchain. However, past behavior is not always indicative of motivations, and therefore may not be indicative of future behaviors. Moreover, it is difficult for a reputation-based system to protect from Sybil attacks [4]. Consequently, it may be more prudent to consider blockchains as consumers of identity rather than proxies for identity.

2. Distribution and redundancy

Distributing the ledger of transactions is a design approach specifically to achieve fault tolerance. The nature of the threat here – that nodes and links can fail or actually be hostile – is a well-defined problem

known as the Byzantine General's Problem [6]. Today, many mechanisms have been crafted, mostly inspired by "Practical Byzantine Fault Tolerance and Proactive Recovery" [7]. Within blockchains, the common algorithmic approach to achieve Byzantine Fault Tolerance (BFT) is to use a consensus protocol. One of the earliest (perhaps the first) formal algorithms for consensus in computer science is Paxos [8].

While these excellent papers provide the formal definitions and approaches to achieving fault tolerance in an uncertain world, their notions can be simply described as "distribute authoritative copies of your transactions widely". The usefulness of the papers is to help understand how to determine how widely distributed and how authoritative is appropriate for a given level of confidence (e.g., security). Common wisdom is that we use a consensus approach designed to achieve at least 51% consensus, and that we need to have a certain minimum number of nodes to achieve tolerance to a certain number of faults ($3f+1$ in [6]).

However, this generalization may have some issues. Consider that a real battlefield has terrain – the ability of a given general to attack or defend or maneuver may be constrained. Further, consider that any given general may not be equal to others in terms of capability or forces. And finally consider that the situation of the terrain and the general likely change over time (for example, because of weather or time of day). In other words, practical BFT must be designed according to the specific conditions in which any given blockchain exists. Applying this idea to familiar concepts of blockchains, some miners (generals) may have higher hashing rates than others and be served by different scales of bandwidth, which in itself may be constrained (by a national firewall, for example). Furthermore, Internet performance varies over time because of global events and natural conditions.

In other words, BFT must be weighted according to the realities of a given blockchain. We may need many more nodes than $3f+1$, or nodes may need to be constrained in some way to achieve a given security result. The closer a blockchain network is to the lower node counts, and the higher the number of faults (failed or malevolent miners/validators), the less likely the integrity of transactions and associated blocks during that period of time. Moreover, the more time it is, the more likely for the blockchain network to come to consensus. This, in turn, may result in uncertainty on the confidence of integrity or validity of a given transaction.

3. Network scale and performance considerations

More nodes in the network doesn't necessarily mean more reliability. For a defined level of consensus, a larger network will take more time than a smaller one, and some transactions are timely, so we need to continue under the assumption that consensus will be achieved, though not guaranteed yet.

As a network grows in scale, propagating a transaction across the network takes more time, and there is an increase in the probability that the blockchain will split. But a well-designed blockchain will handle these situations eventually. That means some amount of time is required to gain high assurance that the transaction becomes fact. That amount of time, for a given amount of certainty (risk), increases with network scale.

But performance may increase from a certain perspective, with the increase in network scale. As the nodes on the network spread farther and wider, access to the network increases, thus reducing the time required to put a transaction onto the blockchain. Performance of the initial step can therefore reduce. The next step is locking the transaction onto a block, as the message propagates. With more nodes, the speed of locking the transaction onto a block should reduce as more nodes compete. The propagation of the information to other nodes on the network should spread at roughly the same speed on a per node basis.

To see this result, consider a model that considers the time to propagate the message, and then to encode onto a block. The time to propagate to a given number of nodes is a function of the network connectivity, connection speed, and processing speed, which should not reduce with network scale (though in some cases it could). The time to encoding onto a block is an order statistics problem, in which the time to first encoding increases with participation. The net result is that in reasonable blockchain network designs, the time to lock a transaction onto the blockchain should reduce with an increase in network size, all other things being equal.

When validating a transaction, how much validation is enough, and how long can an application wait for the validation process? Consider a permissionless network, where participants can join and leave at will. There may be no control as to the membership of subgroups, or their ability to collude. Validation of a transaction among one aligned group of nodes is less assurance than validation by a large group of diverse, non-aligned nodes. In a permissioned network, however, all nodes may be aligned by design, and presumably trusted equally. A single validation may be nearly as good as validation across the entire network in that case.

It may seem that the larger the blockchain network, the greater the chance that a high degree of integrity per transaction can be achieved. This does not necessarily follow. The larger the network, the more likely it is that the network will include bad actors. Moreover, the larger the network, the more messages must be exchanged to achieve BFT (so bandwidth efficiency decreases), and the longer it may take for a given network to come to consensus.

Therefore, if transaction validation time and block confirmation are critical design factors, the size of the network and the associated bandwidth must be designed. In most cases, it will be seen that a smaller blockchain network will result in faster transaction processing times while the security (integrity) of transactions and blocks decreases.

4. Governance and code management

Blockchain networks are very complex systems of hardware, software, and people. Reliability best practices for each element should be followed, but that is not guaranteed. The severe redundancy of distributed blockchain networks is assumed to cover many failure modes, but the tradeoff is not always positive. Because of this complexity, the governance of a blockchain and the processes and practices of managing code become quite important. From a risk management perspective relative to any given use case or service, the body governing a blockchain and the method of how code is managed should actually be considered as supply chain risks and assessed accordingly.

We can consider a specific case in a fork of Ethereum. In June 2016, a major hack was operated against Ethereum that resulted in \$70M of asset losses [9]. To “fix” the results of the hack, the Ethereum community decided (not unanimously) to return all the Ether stolen by changing the smart contract associated transactions. This basically was a willing hack of the immutability of Ethereum by its governance body, with a forced code update that was not accepted by all. Ethereum Classic was born out of the mess.

5. Attack vectors as a reliability problem

While it is beyond the scope of this paper to explore all the known or envisioned attack vectors for blockchain, it is important to note that blockchains must be designed to mitigate known attack vectors, and the list of these vectors is increasing every day. Considering security as a quality of blockchains, security becomes a potential mode of failure for a blockchain. As the attack vectors that apply to a given

blockchain network increase, the blockchain network experiences a form of aging, and as it ages, its rate of failure increases; over time, the blockchain becomes more vulnerable to security failures. Thus, stable blockchains, in relation to security failure modes, experience an increasing hazard rate. Redesign and upgrade are therefore necessary to sustain a blockchain against attack vectors (to maintain its reliability). But introducing new elements to a blockchain expose it to the possibility of infant mortality failure modes. So managing a blockchain's useful life is a complex tradeoff of risk and reliability concerns.

6. What is the meaning of reliability in the context of blockchain?

Any discussion about reliability requires defining the noun involved in the adjective of reliability: reliability of what. A blockchain is by nature reliable in that it is a distributed ledger, copied across multiple locations. Redundancy assures reliability against many failure modes, but not all. And the sheer scale of many blockchain networks assures it will almost never be fully functional, as at least one element may be in failure or disconnected from the network at a given time. Still, the reliability of the blockchain may be far less important than the reliability of the information it contains, or the authentication of the participants, or the reliable responses it gives to translation applications, for example. Large networks are almost always experiencing a failure, yet they reliably support the services and applications that rely on them. The Internet is always experiencing a failure, but services that ride the Internet are generally reliable enough to use.

If a part of a blockchain network is separated from another part, then the network elements can no longer spread information, which is a key function of a blockchain network. If parts of the network can't share information, then the blocks they create will be different, and the chain will split. It is reasonable to assume that eventually the network will rejoin, so that the separation is transient. Eventually, the rejoined blockchain network must prune one branch. For example, in bitcoin, the network eventually will have a branch that is longer, and the shorter will prune. But generally, the mechanism used to decide which is the valid branch in the chain must consider overall reliability and intent; the surviving part must have reliable information, and there should be a recovery mechanism so that reliable information from the other branch is not lost.

Therefore, we can talk about blockchain reliability from multiple frames of reference, and each of them has merit and importance. It is important to consider the reliability of the elements of the blockchain, the overall blockchain itself, the ledgers that are distributed on the blockchain, the information stored in those ledgers, and the applications that rely on it all.

Conclusion

Blockchain networks and the distributed ledgers they maintain have utility in the cable industry, when properly designed and applied. There remains a lot of hype and insufficient clarity around these technologies, so it is important to consider carefully what blockchain networks and distributed ledgers are good for, what they rely on or assume, and what additional work is needed to make them work.

While there will certainly be new emerging use cases in and outside the cable industry in the years to come, it is important to first consider the immediate opportunities. A useful categorization of use cases for operators includes new sources of revenue, ways to use these technologies to reduce costs, applications of these technologies to improve service or enhance the customer experience, and ways to reduce friction or simply make things easier to do.

But to take full advantage of blockchains and distributed ledger technologies, there are several design concerns to address carefully. Predominantly, we believe security and reliability issues are important at

this stage of the lifecycle of blockchain. Instead of avoiding complex problems, these technologies require us to solve some complex problems in security and reliability before we can truly benefit. Fortunately, while complex, the work is reasonable and doable.

Abbreviations

AR	augmented reality
BFT	Byzantine fault tolerance
CMTS	cable modem termination system
CPE	customer premise equipment
ICO	initial coin offering
IP	Internet protocol
NFV	network function virtualization
PKI	public key infrastructure
VoD	video on demand
VR	virtual reality

Bibliography & References

- [1] “A Simple Overview of Blockchains, Why They Are Important to the Cable Industry.” Steve Goeringer. SCTE-ISBE. 2017
- [2] “A Framework for Determining Blockchain Applicability.” Brian A. Scriber, Cablelabs. IEEE. 2018.
- [3] “Mastering Bitcoin,” Andreas M. Antonopoulos, O’Reilly, 2010.
- [4] “Sybil attack,” Wikipedia, downloaded 2018, https://en.wikipedia.org/wiki/Sybil_attack.
- [5] “For Want of a Stronger Chain,” Jason Rupe, IEEE Blockchain Newsletter, July 2018, <https://blockchain.ieee.org/newsletter/july-2018/for-want-of-a-stronger-chain>.
- [6] “The Byzantine Generals Problem”, Leslie Lamport, Robert Shostak, and Marshall Pease, July 1982, ACM Transactions on Programming Languages and Systems, Vol 4, No 3.
- [7] “Practical Byzantine Fault Tolerance and Proactive Recovery”, Miguel Castro and Barbara Liskov, November 2002, ACM Transactions on Computer Systems, Vol. 20, No. 4.
- [8] “The Part-Time Parliament”, Leslie Lamport, May 1998, ACM Transactions on Computer Systems.
- [9] “The DAO, The Hack, The Soft Fork, and the Hard Fork”, CryptoCompare, Online, <https://www.cryptocompare.com/coins/guides/the-dao-the-hack-the-soft-fork-and-the-hard-fork/>, downloaded August 2018.

Supporting The Changing Requirements For Online Gaming

A Technical Paper prepared for SCTE•ISBE by

K. Scott Helms

SVP of Advanced Services

Momentum Telecom

222 Chastain Meadows Court | Suite 100

Kennesaw, GA 30144

404-263-0585

scott.helms@momentumtelecom.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Defining Network Requirements	5
1. Games.....	5
1.1. Battlefield 1 (BF1)	5
1.2. Counter-Strike: Global Offensive (CS:GO)	8
1.3. Destiny 2	10
1.4. Fortnite Battle Royale.....	13
1.5. Minecraft.....	14
1.6. Overwatch	16
1.7. PLAYERUNKNOWN'S BATTLEGROUNDS (PUBG)	19
2. Voice Communications Platforms	21
2.1. Discord	21
2.2. Mumble.....	22
3. Streaming Platforms.....	23
3.1. Twitch	24
3.2. YouTube Gaming	26
4. Remote Rendering – Nvidia GeForce Now.....	27
Conclusion.....	30
Abbreviations	30
Tables of Network Characteristics	32
Bibliography & References.....	33

List of Figures

Title	Page Number
Figure 1 - Battlefield 1 Population.....	6
Figure 2 - Battlefield 1 Network Information Display.....	7
Figure 3 - Battlefield 1 Network Traffic.....	8
Figure 4 - CS:GO Population (PC).....	9
Figure 5 - CS:GO Network Traffic.....	10
Figure 6 - CS:GO In Game Network Information).....	10
Figure 7 - Destiny 2 NAT Type 3 Warning	11
Figure 8 - Destiny 2 Bubble Networking Design	12
Figure 9 - Destiny 2 Networking Traffic.....	13
Figure 10 - Fortnite Network Traffic	14
Figure 11 - Minecraft Active Players	15
Figure 12 - Minecraft Network Traffic.....	16
Figure 13 - Overwatch Population	17
Figure 14 - Overwatch Network Performance.....	18
Figure 15 - Overwatch In Game Network Information	18

Figure 16 - PUBG Global Population	19
Figure 17 - PUBG Versus Fortnite Streaming Viewers.....	20
Figure 18 - PUBG Network Traffic	21
Figure 19 - Discord Network Traffic	22
Figure 20 - Mumble Network Traffic.....	23
Figure 21 - Concurrent Streaming Viewers by Platform	24
Figure 22 - Twitch Network Traffic	25
Figure 23 - Twitch Stream Networking Analytics	26
Figure 24 - YouTube Gaming Network Traffic	27
Figure 25 - GeForce Now Datacenter Locations	28
Figure 26 - GeForce Now Network Analytics.....	29
Figure 27 - GeForce Now Network Traffic (Overwatch)	29
Figure 28 - GeForce Now Network Traffic (PUBG).....	30

List of Tables

Title	Page Number
Table 1 - Network Characteristics, Games	32
Table 2 - Destiny 2 Peer Connections	32
Table 3 - Network Characteristics - Voice Communication	32
Table 4 - Network Characteristics – Streaming Platforms	33
Table 5 - Network Characteristics – Remote Rendering	33

Introduction

Gaming enthusiasts have had a somewhat checkered history of interaction with the service providers that are a critical component of their hobby. On one hand gamers have been early adopters of high speed connectivity, but on the other they have also driven calls, complaints, and had higher requirements than the average end user. We've also seen service providers fail to understand what mattered to gamers when they created packages and marketing materials, and as recent as this year network operators have blamed gamers for driving excess capacity usage. In this paper I will focus on real world usage patterns for online gaming by looking at actual traffic of both games and the supporting software that is commonly used like streaming and voice communications. One of the important shifts in gaming has been the rise of these ancillary programs that gamers use, and these drive very different networking requirements. They also increase the need for consistent performance, and perhaps more impactful in the short term is that they increase the visibility for customers of any issues that might be affecting their traffic. What the testing showed was a dramatic increase in upstream usage and a far higher requirement for low latency and reliable packet delivery. What's even more interesting for operators is that these requirements help cement wired broadband solutions as critical for gamers, and as an industry we should begin thinking about gaming in a similar way that we think about video. Creating relationships inside the gaming ecosystem is clearly in our interests as a way of further fending off encroachment by cellular providers who have aggressively moved into the video space.

Areas of specific measurement and testing were selected for their popularity and for the first time we see the same titles, and their requirements, on most or all of the major gaming platforms PC and consoles.

- Online multi-player games
 - Battlefield 1
 - Counter-Strike: Global Offensive
 - Destiny 2
 - Fortnite
 - Minecraft
 - Overwatch
 - Player Unknown's Battlegrounds
- Voice communication platforms
 - Discord
 - Mumble
- Streaming platforms
 - Twitch (Amazon)
 - YouTube Gaming (Google)
- Remote Rendering – Nvidia GeForce Now
 - Greatly increased downstream requirements (>40 Mbps consistently needed)
 - Upstream requirements similar to “normal” gaming
 - WiFi is a specific point of weakness because any variability in the signal can be seen as frame drops

Defining Network Requirements

Providing broadband service for gaming hasn't been a focus for most service providers, but it's something that we should embrace. The gaming community, especially the action and First-Person Shooter (FPS) genres, require higher performance and more stability for their connections than LTE and later 5G networks will be able easily to provide. I focused on the games selected because they are multi-platform, in fact this list includes some of the most popular games on PC, Xbox, and PlayStation. The networking requirements are assumed to be similar, but testing was done exclusively on PCs for the ability to capture packet performance. Games are sorted alphabetically.

One of the critical takeaways of this research is that customers involved in gaming are likely to have specific information about network performance provided by the games they play themselves as well as the ancillary programs they use: BF1, CS:GO, Twitch, Mumble, and Discord.

1. Games

Gaming traffic, other than file transfers for updates and initial installs is relatively light and generally has sustained data rates measured in Kbps. None of the games in this paper needed more than 7 Mbps of peak download and that was for just a few moments while new graphical files were streamed down, not actual game play.

Measurements were performed using packet captures on the PC running the game client. Data was extracted using Wireshark filtering which was then exported in CSV format for detailed analysis and graphing. Latency and packet loss were introduced in controlled amounts by using Clumsy, a utility written for that purpose. Injected latency and packet loss was done on outbound packets from the PC only. In Fortnite changes in latency were seen as cheating attempts so extended testing with that game was impossible with this methodology. It is also important to understand that the games see and measure latency in different ways so these are all relative measurements. Many games do not provide latency information via an in-game display and in those cases the measurements were derived using ICMP to the specific game server. Bandwidth measurements were done by analyzing the packet captures and should be consistently accurate for all layer 3 (IP) traffic and above. Overhead from the lower levels of the network stack were not included. It is important to know that bandwidth usage can change even in the specific games measured and some activities increase the amount of traffic. This is especially true in Destin 2, given the peer to peer connections and "bubble" architecture of the game.

1.1. Battlefield 1 (BF1)

Battlefield 1 is a first-person shooter that regularly has tens of thousands of players in matches across PC, PS4, and Xbox platforms. The game was released October 21, 2016 but still has a healthy population according to a third-party Battlefield tracker across all three active platforms.

(<https://battlefieldtracker.com/bf1/insights/population?days=30>)




Figure 1 - Battlefield 1 Population

BF1 provides a detailed networking view that measures many of the critical elements related to online gaming. Below you can see the in-game display during latency testing.

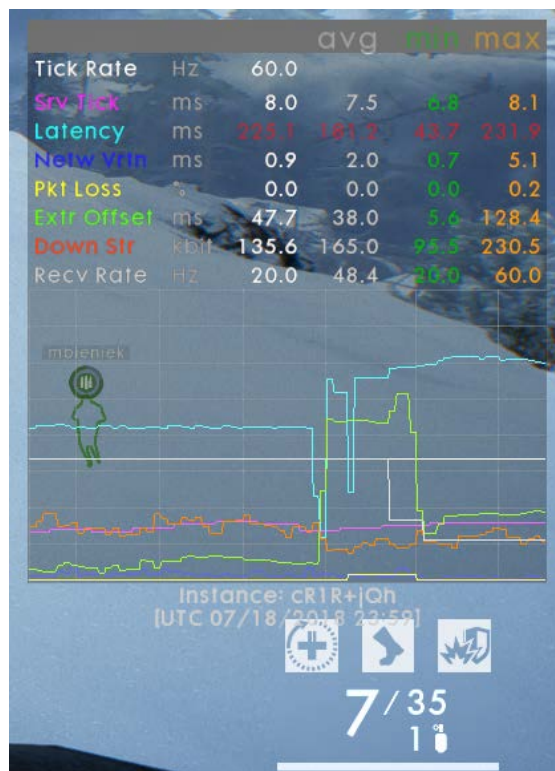


Figure 2 - Battlefield 1 Network Information Display

The display will flag metrics that are outside of acceptable bounds by the game, and in this case we can see that the latency displays in red because it's too high. Latency, even lower than the thresholds I reported, can affect game play in BF1 and most other FPS games because latency affects where a player's shots will land. The impact on accuracy of hit detection is often the first sign of latency. Packet loss has a similar profile but will also often cause issues in movement sometimes called "rubber banding" where a player's in game avatar will appear to teleport back to an earlier position from the player's point of view. Jitter, defined as rapid changes in latency, will have a greater impact on gaming experience than latency of a consistent nature. Often having a moderate amount of latency is better than having lower latency that's frequently fluctuating, even if the average latency is the same or even lower than a connection with consistent packet delay.

The networking traffic is relatively light as you can see from the graph below, but as was just discussed, the game has tight latency and packet loss budgets. The average traffic over several sessions was 222 Kbps down and 101 Kbps up on average with a peak down rate of 510 Kbps and peak up rate of 207 Kbps.




Figure 3 - Battlefield 1 Network Traffic

1.2. Counter-Strike: Global Offensive (CS:GO)

CS:GO is a first-person shooter that was released on August 2012 and it still maintains a highly active player base, though unlike most of the other games on the list almost all them are on PC rather than consoles.




Figure 4 - CS:GO Population (PC)

CS:GO is an older game and has a heavier average download requirement than any of the other games measured. It also has fairly tight latency and packet loss needs. The upstream traffic is also more consistent over gaming sessions with variability as seen in some of the other games.




Figure 5 - CS:GO Network Traffic

CS:GO players tend to be very aware of network performance issues and the game includes a display of several performance metrics that include networking.




Figure 6 - CS:GO In Game Network Information)

1.3. Destiny 2

Destiny 2 is a Massively Multiplayer Online (MMO) FPS game which consists of both Player versus Environment (PvE) and Player versus Player (PvP) play. July 15 saw 469.4K Crucible (PvP) Players and

628.8K PvE Players across all platforms as reported by DestinyTracker, a third-party Destiny 2 statistics site.

Destiny 2 is unique on this list from a networking perspective. It's the only game measured that allows direct network connections between the players on a team and others nearby in the MMO world. This reduces the network traffic that needs to pass through the game's servers so because much of the traffic flows between team mates. Networks that only allow strict NAT produce a warning and cause the game servers to handle all of the traffic flows.

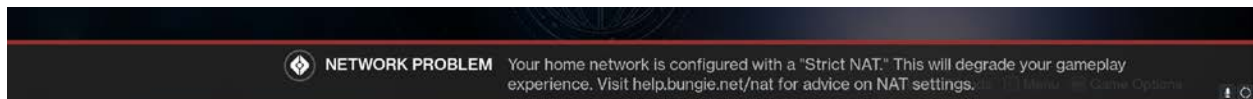


Figure 7 - Destiny 2 NAT Type 3 Warning

This is warning players will see if the network they're playing doesn't allow for direct connections (NAT Type 2 or DMZ).

Destiny 2 doesn't have an in-game display for network performance probably because of the complicated nature of their networking approach which can have direct connections between players on teams as well as just players nearby. Bungie, the developer, describes the approach as "bubbles" where specific numbers of related assets are simulated by the game engine together.




Figure 8 - Destiny 2 Bubble Networking Design

This design in networking means that players who are far away from each other from a networking standpoint will impact each other's performance when they are nearby in the game. Players on the same team in the game are always in the same bubble and packet captures show long term peer to peer connections as a result. This situation has the potential for both intentional and unintentional impacts in the game. PvP matches where each of the players can see the IP addresses of the opposing team is something that seems to be inviting DoS attacks.




Figure 9 - Destiny 2 Networking Traffic

Here you can see the traffic from a session of three teammates. The individual hosts contribute around 50 kbps of traffic in both directions so actions involving large numbers of players like raids, open world public events, or PvP could involve a lot more transfer. Bungie shots for a 6v25 max “bubble” with the 6 being the number of players and 25 being the number of computer generated opponents at any one time. 6v6 PvP would generated around 600 Kbps from peers in the same “bubble” with another 30 Kbps for the game host system.

1.4. Fortnite Battle Royale

Fortnite is a specific genre of action games called a battle royale. Instead of a first-person view point the game seen in third person, where the player is basically looking over the shoulder of his or her in-game avatar. Fortnite is incredibly popular right now with Epic Games claiming more than 125 million active players. They also generated \$318 million in revenue in May 2018. This is remarkable in part because the game itself is free and only generates revenue by selling in-game cosmetic items. Fortnite has surpassed the other big name in the battle royal genre, PUBG, by being more accessible and attracting large numbers of players on many platforms including mobile.

In terms of network requirements Fortnite is very forgiving for an action game, perhaps reflecting their intentions to offer a mobile version early on. The average rates for download traffic were only 32.3 Kbps while the upload traffic was 33.3 Kbps. The max rates were also pretty light, 186.6 Kbps down and 205.9 Kbps up. Latency tolerance is pretty high considering this is an action/shooter title at 175 msec. Packet loss of up to 7% allowed for the game to remain playable, though at that level stuttering could be detected.




Figure 10 - Fortnite Network Traffic

1.5. Minecraft

Minecraft is the only game that was measured that wasn't in the shooter or action category. It was included because of its overwhelming popularity. It appears on PCs, Macs, mobile devices, and most consoles. Since its launch in 18 November 2011 the game has accumulated more than 150 million copies sold and has 74 million people playing every month. Since this is the oldest game that was measured it's remarkable that there was increase of 20% in active players from 2017 to 2018.




Figure 11 - Minecraft Active Players

Minecraft is a building game with several modes including “vanilla” survival, creative which focuses on building and not escaping zombies, and new modes have been created to help teach kids coding and better connect with history and literature. There are more than 2 million users of Minecraft: Education Edition. Minecraft can be played either in single player or multi-player modes with the single player mode not needing network connectivity to function. For multi-player games the networking requirements are fairly modest, though the transfer of custom art can generate a substantial amount of transfer. Testing showed more than 7 Mbps of peak download transfer and the average down rate was 281Kbps. The upload rates were much lower with an average rate of 31.5Kbps and a peak of 116.35Kbps. The latency threshold was around 250 milliseconds before playing online felt really problematic and packet loss as high as 10% could be tolerated for short stretches of time.




Figure 12 - Minecraft Network Traffic

1.6. Overwatch

Overwatch is a multi-platform FPS with a substantial population. It was released in May 24, 2016 and the active population has risen past 40 million across all platforms.




Figure 13 - Overwatch Population

Overwatch shares a similar format to older FPS games but is more forgiving in terms of network requirements. The average download traffic was 249.4Kbps while the upload side averaged at 54.6Kbps. The peak traffic, which on the download side can include streaming of assets, goes quite a bit higher at 3234.2Kbps for short bursts and the upload peaked at 99.63Kbps. Latency tolerance was good for a FPS and even around 140 milliseconds the game felt responsive most of the time. As with other games changes in latency (jitter) have a more negative impact than consistent latency. The game continued to be playable, though I'd recommend sticking to casual or arcade modes, with as high as 6% packet loss.




Figure 14 - Overwatch Network Performance

One interesting item with Overwatch is how much the bandwidth decreases between matches. You can see the gaps in the graph above and the times when the game went into matchmaking mode the traffic dropped substantially.



Figure 15 - Overwatch In Game Network Information

1.7. PLAYERUNKNOWN'S BATTLEGROUNDS (PUBG)

PUBG is another battle royal game that can be played in either first or third person mode. It went into full release on December 20, 2017 but players on PC started playing it during the early access phase which started in March of the same year. It has amassed a tremendous following and only recently was surpassed by Fortnite in terms of active streams.




Figure 16 - PUBG Global Population

PUBG competes directly with Fortnite and as the latter's popularity has increased PUBG has seen its growth slow.




Figure 17 - PUBG Versus Fortnite Streaming Viewers

PUBG has surprisingly low average bit rates. 19Kbps down matched with 34Kbps up is far lower than I expected for this game, however like Fortnite this may have been the result of designs that are mobile friendly. The max rates are also modest with the peak down being 261.68Kbps and the peak up rate being 97.656Kbps. The game tolerates latency moderately well but after about 125 milliseconds of delay the feeling of lag was noticeable. Packet loss tolerance was similar to Overwatch at 6%.




Figure 18 - PUBG Network Traffic

2. Voice Communications Platforms

Voice communications is a key part of most online action and FPS games. In many cases the games provide voice communications and text chat to teammates and others in the game but the popularity of external communications programs, especially Discord, has continued to increase. There are other options for external voice communication and they range from TeamSpeak to Skype in their approaches and focus on gamers. Skype is a very general-purpose voice communication platform, but for small groups it works fine for voice communication. TeamSpeak was once one of the most popular platforms but has faded behind Mumble and especially Discord to the point of now being uncommon. For testing I used a small group of three for each of the voice platforms.

2.1. Discord

Discord is by far the most popular voice communications platform for gamers and it adds a persistent shared chat similar to Slack in terms of functionality, and Discord is free to use. It was released in May 2015 and today has more than 100 million active users and signs up approximately 1.5 million new users a week. Discord is entirely hosted by the developer on Google's cloud infrastructure. Part of the goal behind creating Discord was making a communications platform that was easy for players use and didn't require specialized hosting companies as earlier offerings had. Discord also has mobile apps and that adds to its appeal over some other offerings.

Discord is quite network efficient, even with large numbers of users in channel, however the chat and pictures can add large spikes in traffic. The average download usage was 222Kbps and the average up was 101Kbps with peak download being 6952.91Kbps and peak upload 70.368Kbps. In general, it tolerated latency pretty well with complaints not really occurring until 250 milliseconds of delay. Packet loss was also handled well and Discord tolerated up to 7%. As you can see the spikes were rare, but much larger than the average traffic.




Figure 19 - Discord Network Traffic

2.2. Mumble

Mumble is the distant second in the voice communications genre, but it definitely has its adherents. For one thing Mumble and its server-side component Murmur are fully open source. It's also built around a very low latency and low bit rate codec which gives it a lower bit rate than Discord, often much lower as well as less lag in voice communications. It's impossible to give hard numbers around Mumble users because there's no central place from which to get statistics. Gamers who want to use Mumble either contract with a hosting company to deliver it as a service or install and run Murmur from a server they already have. I don't think Mumble will ever go completely away, but the ease of use of Discord has already made a substantial change in usage.

Mumble is very network friendly and on average only needs 36.5Kbps down and 32.4Kbps up. The peak seen during testing was 123.32Kbps down and 113.464Kbps up. It's a little less latency tolerant than Discord with audio problems showing up around 200 milliseconds of latency and is also less tolerant of packet loss with 5% causing issues. It's worth noting that Mumble appears to have a firm footing with small groups of competitive players where the latency and sound quality are more important than the ease of use from Discord. Make sure to keep in mind the scaling of the bandwidth axis if you compare the graphs of Mumble and Discord traffic.




Figure 20 - Mumble Network Traffic

3. Streaming Platforms

Streaming of games is a relatively new phenomenon and reflects the shift away from traditional forms of media content. Streaming is generally done while the streamer is playing a game and providing commentary at the same time. Many streamers will show their face via webcams in addition to their voice commentary. The number of people consuming streaming, almost all around gaming, is staggering.




Figure 21 - Concurrent Streaming Viewers by Platform

For the purposes of this paper I focused on the top two platforms, Twitch (owned by Amazon) and YouTube (owned by Google). Streaming is also where gamers begin to substantially diverge from the normal networking requirements of low latency and low jitter but also low bit rates. It also dramatically increases upstream usage. An important note, most of the streaming services look very similar from a networking standpoint. The differences in upload speeds are largely around settings in the streaming client for video bit rate.

3.1. Twitch

Twitch is by far the most popular streaming platform right now. It was launched as a service in June 2011 as a spinoff of a general-purpose streaming site and has since far eclipsed its progenitor Justin.tv. Amazon acquired Twitch in 2014 for \$970 million. In terms of network requirements Twitch needs very stable connections and consistent latency especially on the upstream side. Average down speeds were 124.4Kbps while average upload speeds were 5,564.8Kbps. Peaks were also impressive with peak upload speeds of 10,884.5Kbps and peak down of 271.76Kbps. This is a very asymmetrical usage pattern, but in the opposite direction of what the broadband industry has been building for in many cases. Packet loss of greater than 1% made streaming almost impossible.




Figure 22 - Twitch Network Traffic

Twitch also provides streamers with detailed networking analytics. If a streamer is using your service and experiencing issues you will likely get detailed information around their problem.

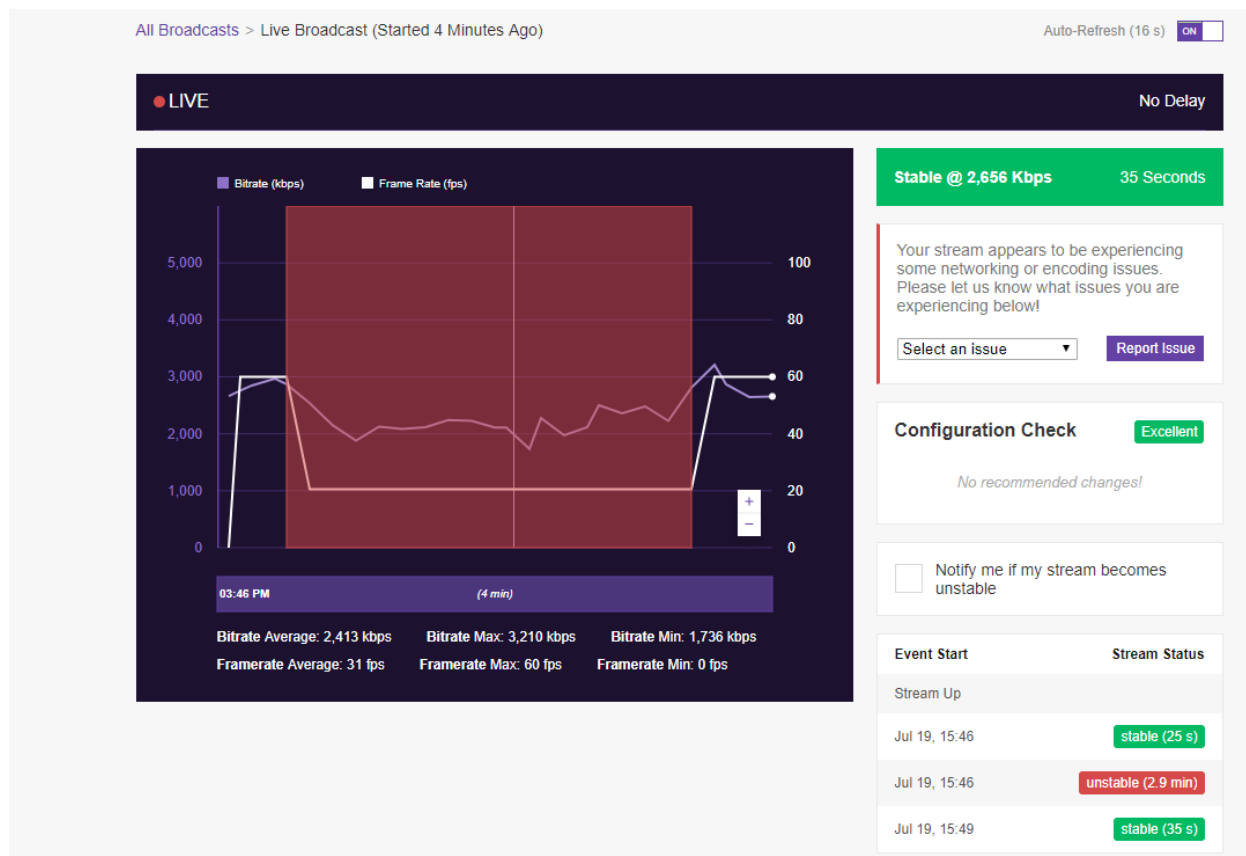


Figure 23 - Twitch Stream Networking Analytics

3.2. YouTube Gaming

YouTube Gaming offers a very similar experience to Twitch from a streaming standpoint, but it also captures all of the videos automatically for replay. The networking requirements are nearly identical. Average download speeds were 119Kbps, average upload speed 3,350Kbps (note this is lower because all the testing was done at the lower video bit rate), while the peaks were 175.22Kbps down and 8,225,87Kbps up.




Figure 24 - YouTube Gaming Network Traffic

4. Remote Rendering – Nvidia GeForce Now

Remote rendering is an idea that's been around for a while, but as of yet no company has been able to make it a commercially viable offering. Much of the challenge in the past had more to do with the lack of very highspeed broadband. That looks to be changing with some major companies entering the market including Nvidia and Google. The key ideas behind remote rendering is turn gaming into more of a service similar to Netflix for gaming. One challenge for gaming is that you periodically need to upgrade your hardware whether it be a PC, a consoler, and increasingly this applies to mobile devices. If a remote server is doing the rendering then the local device is simply displaying a video stream and relaying the control information from the player to the remote server. This makes it possible to play very high end games on PCs and devices without dedicated video cards or powerful processors. Nvidia clearly sees the move to gaming as a service as part of their strategy. Nvidia GeForce Now is currently in beta testing. Figure 25 shows an excerpt from a recent Nvidia presentation.

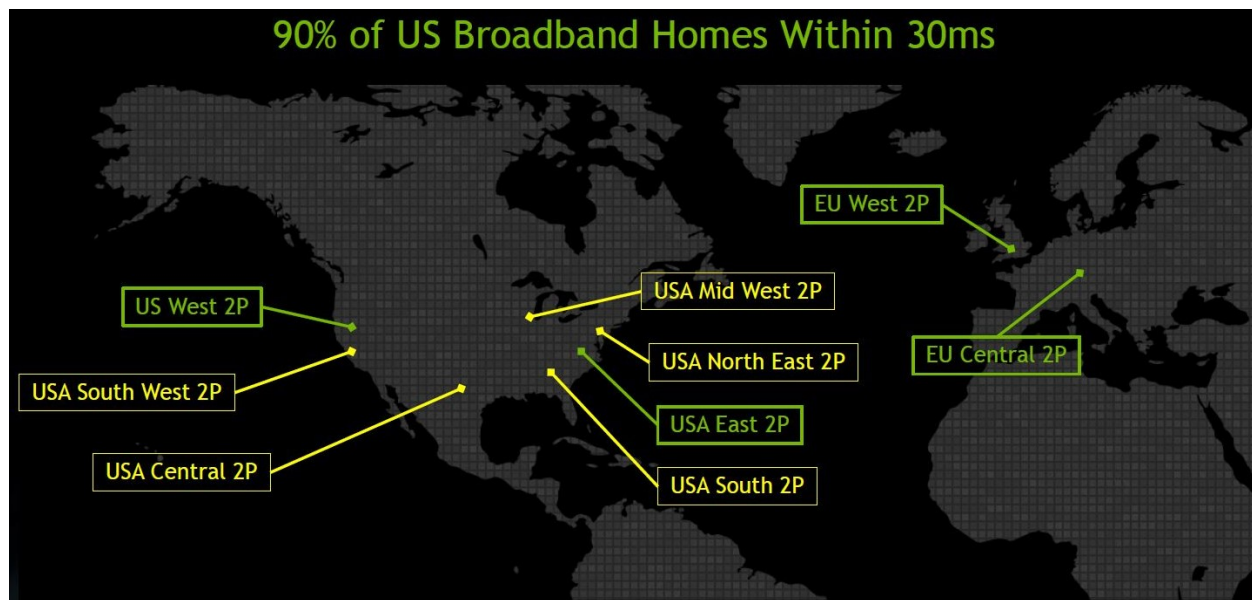


Figure 25 - GeForce Now Datacenter Locations

The speed and latency requirements are substantial. From Nvidia's support FAQ we can see what speeds support what visual qualities.

- 10 Megabits per second – Required broadband connection speed
- 20 Megabits per second – Recommended for 720p 60 FPS quality
- 50 Megabits per second – Recommended for 1080p 60 FPS quality

In testing my 50mbps package was not able to sustain 1080p gaming sessions. The amount of sustained transfer for a gaming session is basically double that of a Netflix ultra HD stream. Testing did not show substantial difference between games, but that could change in the future. The system also provides users with feedback on the quality of their networking connection.

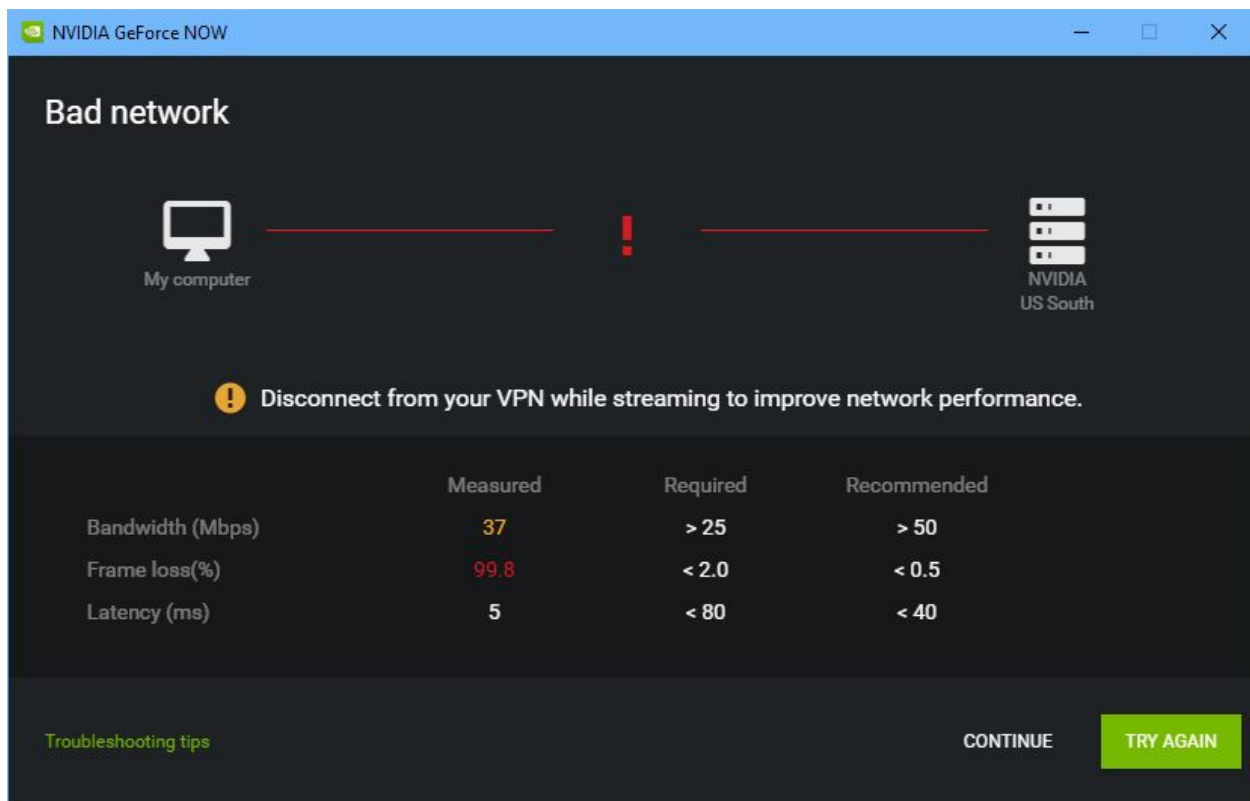


Figure 26 - GeForce Now Network Analytics




Figure 27 - GeForce Now Network Traffic (Overwatch)




Figure 28 - GeForce Now Network Traffic (PUBG)

Conclusion

Gaming has not been recognized as a key service for customers in the way that streaming video has. Network engineers need to recognize and plan their capacity around the current and near future requirements for gaming and that includes much higher bit rates, especially on the upstream, and optimizing for lower latency and packet loss. It might also include direct peering or other arrangements with gaming providers and CDNs to maximize performance and control costs. Today it's not uncommon for consumers to consider how well an MSO supports Netflix and other OTT video providers in their decision making. I believe that gaming has the same ability to influence decisions if we build our services correctly. Customer support also needs to be educated in supporting the games and auxiliary applications. Most service providers today are comfortable handling questions around why Netflix is slow and we should get comfortable handling the questions around gaming performance and Twitch streaming.

Abbreviations

BF1	Battlefield 1
bps	bits per second
CS:GO	Counter-Strike: Global Offensive

DMZ	Demilitarized Zone, portion of the network without firewall protection
DoS	Denial of Service, a type of attack that knocks out access but doesn't compromise a system
FPS	First-Person Shooter
Kilobits	1000 bits
Latency	Time delay created by packets transiting the physical network
Mbps	mega-bits per second
MMOFPS	Massively Multiplayer Online First-Person Shooter
NAT	Network Address Translation
PUBG	PLAYERUNKNOWN'S BATTLEGROUNDS
PVE	Player Versus Environment
PVP	Player Versus Player

Speeds are in kilobits per second (Kbps) unless otherwise noted. The thresholds are based on personal observation of when the impairment becomes clearly noticeable to the user. Given that there is no empirical method to discern acceptable performance in gaming this opinion and other users in other network conditions may see unacceptable performance at lower thresholds. I have tried to create reasonable values for these measures and I will be providing updates to this information over time.

Tables of Network Characteristics

Table 1 - Network Characteristics, Games

Name	Average Down Kbps	Average Up Kbps	Peak Down	Peak Up	Latency Threshold	Packet Loss Threshold	Method
BF1	222.00	101.00	510	207	130	3%	Game Display
CSGO	343.00	70.00	581.12	105.688	100	4%	Game Display
Destiny 2	32.00	27.90	123.32	113.464	150	6%	Derived from local ping plus injected
Fortnite	32.33	33.33	188.648	205.992	175	7%	Derived from local ping plus injected
Minecraft	281.00	31.50	7093.84	116.352	250	10%	Derived from local ping plus injected
Overwatch	249.40	54.60	3234.2	99.632	140	6%	Game Display
PUBG	19.00	34.00	261.68	97.656	125	6%	Derived from local ping plus injected

Table 2 - Destiny 2 Peer Connections

Average Peer Down	Average Peer Up	Peak Peer Down	Peak Peer Up
50.25	52.25	141.088	114.88
Average Peer Down	Average Peer Up	Peak Peer Down	Peak Peer Up

Table 3 - Network Characteristics - Voice Communication

Name	Average Down Kbps	Average Up Kbps	Peak Down	Peak Up	Latency Threshold	Packet Loss Threshold
Discord	222.00	101.00	6952.91	70.368	250	7%
Mumble	36.5	32.4	123.32	113.464	200	5%

Table 4 - Network Characteristics – Streaming Platforms

Name	Average Down Kbps	Average Up Kbps	Peak Down	Peak Up	Latency Threshold	Packet Loss Threshold
Twitch	124.4	5,564.8	271.76	10,884.5	300	1%
YouTube Gaming	119	3,350	175.22	8,225.87	300.00	1%

Note that speeds in streaming is mostly determined by the audio and video bitrate, especially the video rate.

Table 5 - Network Characteristics – Remote Rendering

Name	Average Down Mbps	Average Up Mbps	Peak Down Mbps	Peak Up Mbps	Latency Threshold	Packet Loss Threshold
GeForce Now	21	0.341	34.26	0.5	40	<1%

Bibliography & References

Battlefield Tracker -- <https://battlefieldtracker.com/bf1/insights/population>

Bungie Presentation on Destiny 2 Networking -- http://twvideo01.ubm-us.net/o1/vault/gdc2015/presentations/Truman_Justin_Shared_World_Shooter.pdf

Clumsy -- <https://jagt.github.io/clumsy/>

Statista -- <https://www.statista.com/>

SUPERDATA -- <https://www.superdataresearch.com/us-digital-games-market/>

The Benefits Of Leveraging Multi-Vendor Orchestration To Achieve True Service Agility

A Technical Paper prepared for SCTE•ISBE by

Arvinder S Anand

Director

Ericsson

1 Ericsson Drive Piscataway NJ

201-470-3197

arvinder.anand@ericsson.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	3
1. Key Business Challenges	4
2. Key Operational Challenges	4
3. Primary Drivers For NFV	4
4. NFV Deployment Drivers	5
5. Top Barriers For NFV Deployment.....	5
6. Key Benefits Of Multi-Vendor Orchestration	6
7. Architectural Principles For Multi-Vendor Orchestration:	6
8. Multi-Vendor Orchestration Value Proposition	7
9. Key differentiators of Multi-Vendor Orchestration	7
Simplicity, Flexibility, and Automation	7
Comprehensive, Built-in Features	7
Vendor Independent and Cross-domain Orchestration	8
Secured Management of Distributed Cloud	8
Policy-driven Resource Handling	8
10. Transformation Journey	9
Conclusion.....	10
Abbreviations	10

List of Figures

Title	Page Number
Figure 1 - Approach to Multi-Vendor Orchestration	3
Figure 2 - NFV Deployment Drivers	5
Figure 3 - Top Barriers For NFV Deployment	5
Figure 4 - Transformation Journey	9
Figure 5 - Multi-Vendor Orchestration.....	10

List of Tables

Title	Page Number
Table 1 - Multi-Vendor Orchestration Value Proposition	7

Introduction

Operators need an easy but comprehensive solution that transitions their operations from providing traditional services to offering a hybrid of physical and virtual services. In this environment, innovative services are not only delivered to customers on demand but must be able to respond to surrounding dynamics in real time. Customer provisioning must move from weeks to minutes and Service Level Agreements must be measured and enforced.

As Operators transform to meet these needs, enabling automated orchestration of available resources at scale is foundational. To accomplish this, operators are shifting their focus from implementation of Operation Support Systems “stacks” made of vertical fulfillment or assurance silos with heavy reliance on System Integration services, to a simplified pre-integrated horizontal architecture that can be readily configured to support specific service offerings including those that leverage VNFs. This need goes well beyond what ETSI Management & Orchestration (MANO) Network Function Virtualization (NFV) Orchestrators and NFV Managers can offer.

Multi-Vendor Orchestration answers the call through a flexible and modular service orchestration solution that fully automates the multiple layers of complex processes for service creation, delivery and assurance. It provides rapid validation of VNFs, onboarding of new services, resource management, service design and configuration, and closed-loop policy-based service assurance for Service Level Agreement compliance. It also supports capacity management to provide the right level of resources in real time.

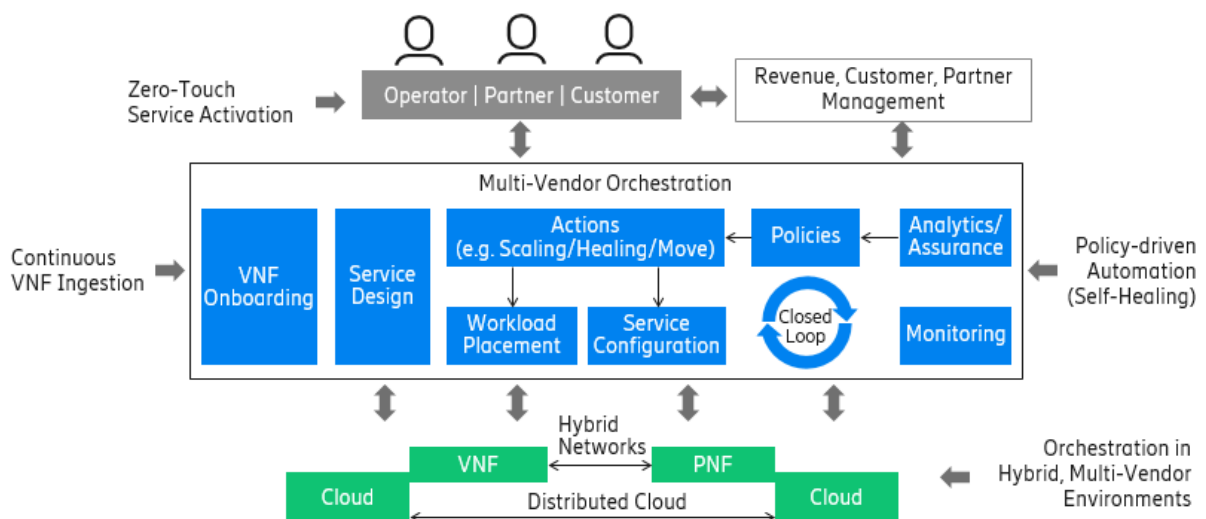


Figure 1 - Approach to Multi-Vendor Orchestration

Content

Operators are under increasing pressure to deliver services faster and more flexibly at the lowest cost possible. This may have created interest in adopting cloud architectures and network functions virtualization (NFV). However, managing virtual applications and resources in the cloud requires a structured, disciplined approach.

- Speed Time to Market

- Zero-Touch Automation
- Visualize End to End Services
- Maximize Resource Usage
- Benefit from Hybrid Orchestration
- Gracefully Transition Operations
- Eliminate Vendor Lock-In
- Increase Business Agility

1. Key Business Challenges

- Minimizing up-front investments – finding the right level of leverage to keep risk at comfortable levels.
- Minimizing operational (recurrent) costs – adopting cloud technologies while keeping outsourcing, energy and data center expenses at bay.
- Maximize resource utilization – having the kind of operational visibility to recover stranded resources, identify underutilized resources, and to determine the impact of the level of resource utilization on service quality.
- Create compelling service offerings – introducing new services that fully exploit the elasticity of virtual resources.
- Become more agile without compromising service quality – acquiring operational speed and flexibility while maintaining consistent levels of service quality.
- Differentiate from over-the-top competitors – identifying the kind of role the network should play to dramatically improve the customer experience.
- Guarantee security – devising the right mix of security features to deal with regulations and the shared nature of cloud implementations.

2. Key Operational Challenges

- Manual configuration and troubleshooting individual network segments for each End to End Network Services
- Manual and paper hand-offs among different isolated network segments
- Lack of End to End Network Services automation
- Network service / device modeling based on open standards is nonexistent.
- Current assurance processes are mostly manual for Virtual Network Functions
- Many different incompatible EMS systems between silo's, CLI's, scripts, templates, cookbooks
- Lack of Automated testing framework

3. Primary Drivers For NFV

- Service Agility resulting in quicker time to revenue. Operators can quickly add, drop and change the services and applications they offer by using Software Defined Networking control software and Network Function Virtualization on virtual machines or containers on commercial servers
- Operation Efficiency to provide a global view of the network for provisioning multi-vendor network and multiple layers. The fine-grained control offered by Software Defined Networking Control Software to enable carriers to utilize network equipment better, thereby minimizing the amount of equipment they need and reducing capex costs.

Operators must be agile to deliver services on a global scale in an era when speed and governance are essential. Stimulated by the advances in IT networking and driven by the migration to NFV, services are

becoming more and more cloud-based. In this environment, innovative services are not only delivered to customers on demand but must be able to respond to surrounding dynamics in real time. Customer provisioning must move from weeks to minutes and Service Level Agreements must be monitored and enforced. As operators transform to meet these needs, enabling automated orchestration of available resources at scale is foundational.

4. NFV Deployment Drivers




Figure 2 - NFV Deployment Drivers

5. Top Barriers For NFV Deployment

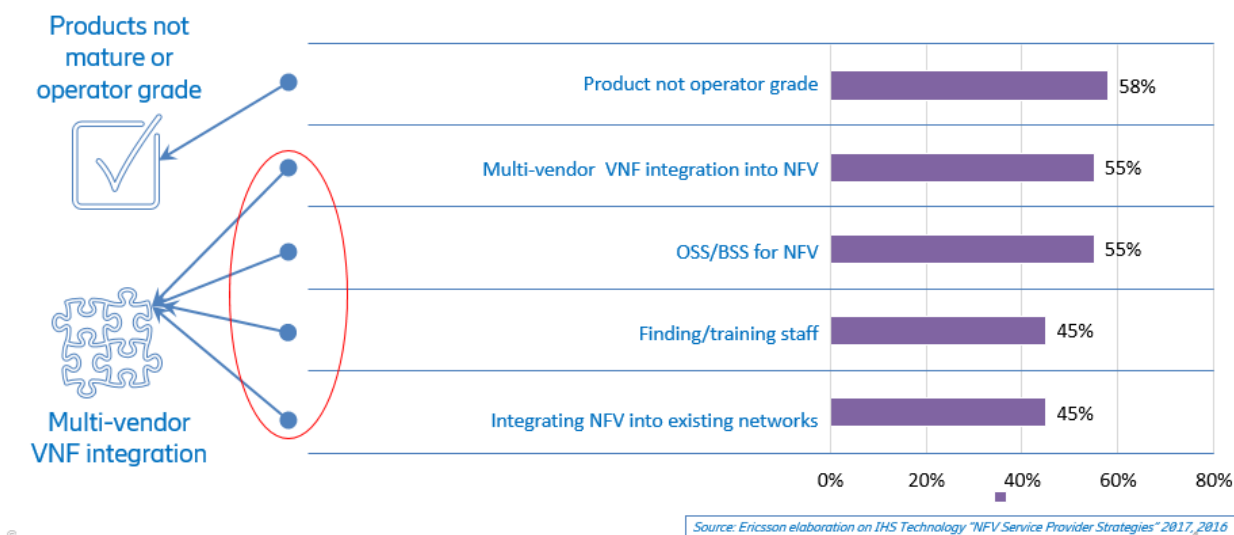


Figure 3 - Top Barriers For NFV Deployment

6. Key Benefits Of Multi-Vendor Orchestration

Operators should consider Multi-Vendor Orchestration if they want to:

- Speed time to market. Quickly create and package new service bundles that can be offered to consumers and businesses via self-service.
- Achieve zero-touch automation. Enable closed-loop orchestration for flawless provisioning, configuration and assurance. The solution has been proven to reduce customer service delivery time from days to minutes.
- Visualize end to end services. Let system users and end customers alike see across multiple domains, including the cloud domain.
- Maximize resource usage. Load balance cloud-based assets for better data center utilization.
- Benefit from hybrid orchestration. Seamlessly orchestrate across physical and virtual network domains in a consolidated solution for lower TCO.
- Gracefully transition operations. Evolve naturally from Ericsson or other vendors' systems to Dynamic Orchestration.

7. Architectural Principles For Multi-Vendor Orchestration:

Model & Catalogue Driven: Network Service introduction should be catalog and Model driven.

- Use of open standards: End-to-End Network service modelled with TOSCA
- Programmability: Transaction-safe, model-driven APIs, Publish/subscribe approach for real-time synchronization
- Transactional: Ensures consistent state, automatic recovery from failed configurations
- Consistent State: Mapping between service and devices in real-time
- Modularity: Components must be reusable, modular, loosely coupled and self-contained

8. Multi-Vendor Orchestration Value Proposition

Table 1 - Multi-Vendor Orchestration Value Proposition

		Multi-Vendor Orchestration					
		Optimize Utilization & Quality	Faster Innovation	Minimize Risks	Smooth Integration	Multi-vendor	Enforce Security
Business Challenges	Minimize up-front investment	X		X	X		
	Minimize operational costs	X		X	X	X	
	Maximize resource utilization	X	X	X		X	
	Create compelling services	X	X		X	X	X
	Become more agile without compromising quality	X	X		X	X	
	Differentiate from OTT	X	X		X		X
	Guarantee security			X	X		X

9. Key differentiators of Multi-Vendor Orchestration

- Simplicity, flexibility, and automation
- Comprehensive, built-in features
- Vendor independent and cross-domain orchestration
- Secured management of distributed clouds
- Policy-driven resource handling

Simplicity, Flexibility, and Automation

Multi-Vendor Orchestration combines the simplicity and flexibility of IT activities with the scale of telecom operations to configure, coordinate and manage VNFs and associated services across highly distributed cloud environments. This is made possible by a closed-loop orchestration that flexibly adapts to the changing environment. With it service providers can manage the lifecycle of services and resources as the underlying VNF capabilities evolve. Its comprehensive workflow automation engine executes both predefined and user-defined workflows. By having a flexible catalog driving its workflow engine, Multi-Vendor Orchestration can enforce the consistent execution of workflows within and across domains to expedite the rollout of new products, services and VNFs.

Comprehensive, Built-in Features

Multi-Vendor Orchestration comes with a complete set of configuration/activation, fault management, performance, accounting and security features for the end-to-end operation of cloud platforms.

Multi-Vendor Orchestration allows service providers to reduce their OPEX significantly in managing large scale, distributed Clouds by providing a cohesive platform that reduces the operational fragmentation of current solutions.

Vendor Independent and Cross-domain Orchestration

Multi-Vendor Orchestration can work with any domain managers to coordinate virtual applications and their resources in hybrid cloud environments across virtualized and physical domains. In addition, Multi-Vendor Orchestration includes open APIs, workflow design tools and a software development kit to facilitate the integration with third-party infrastructure and systems.

Secured Management of Distributed Cloud

Multi-Vendor Orchestration is designed to configure, coordinate and manage applications, services and their underlying virtual and physical infrastructure in highly-distributed cloud environments connected over one or more networks.

Multi-Vendor Orchestration supports virtual data centers (VDCs) and virtual applications (vApps). A VDC logically groups distributed compute, storage and networking resources across data centers and geographical boundaries. The deployments of vApps within a VDC inherently utilize the distributed resources across multiple physical data centers, including the network(s) connecting them. As a result, distributed cloud environments require a more granular level of security than what traditional cloud management solutions can offer. Multi-Vendor Orchestration can handle multi-tenancy by partitioning the data at every level. In addition, Multi-Vendor Orchestration includes audit features specifically designed to monitor and enforce security policies throughout distributed cloud environments. It also provides encryption with key management to control the access to the handling of physical and virtual resources.

Policy-driven Resource Handling

Given the levels of automation and dynamicity typical of cloud environments, resource handling becomes more of a challenge, when it comes to:

- optimizing resource and workload allocation while meeting required quality of service.
- reducing stranded capacity and containing the virtual machine sprawl typical of self-service provisioning.
- forecasting resource utilization because of expected business growth or configuration changes.

Multi-Vendor Orchestration addresses these operational challenges with policy-driven resource/workload management and what-if analyses to identify the most viable scenarios.

10. Transformation Journey

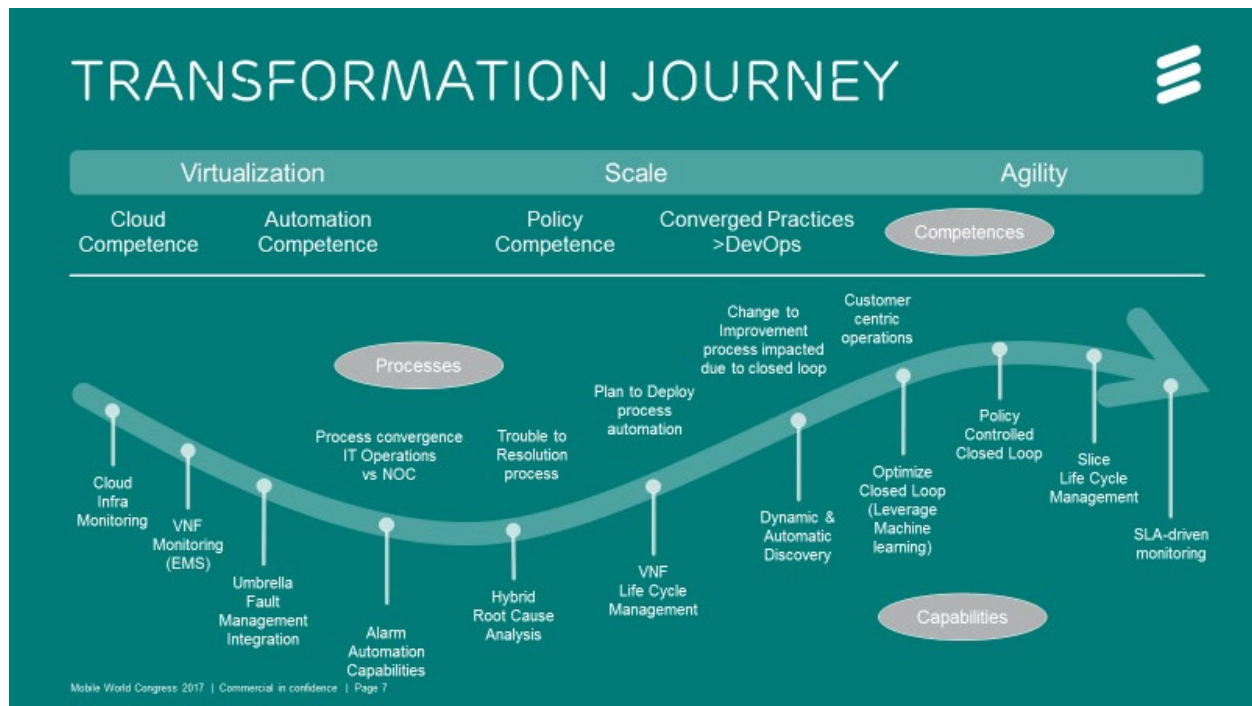


Figure 4 - Transformation Journey

Conclusion

To Achieve the Service Agility, Orchestration is the KEY.




Figure 5 - Multi-Vendor Orchestration

Abbreviations

NFV	Network Function Virtualization
MANO	Management and Orchestration
ETSI	European Telecommunications Standards Institute

The Emerging Impact and Use Cases of Blockchain Technology in the Era of HFC Connected People and Things

A Technical Paper prepared for SCTE•ISBE by

Sandeep Katiyar
Senior Consultant
Nokia Bell Labs Consulting
Bldg. 9A, 7th Floor, DLF Cybercity
Gurugram, India-122002
sandeep.katiyar@bell-labs-consulting.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Background	4
Blockchain Impact and Benefits on the HFC Network and Services	6
Blockchain Design for SDN Implementation	8
Blockchain Design for IOT Implementation.....	10
Example HFC enablement for Blockchain	12
Conclusion.....	12
Abbreviations	13
Acknowledgments	13
Bibliography & References.....	13

List of Figures

Title	Page Number
Figure 1 - Nodes & Links in Centralized vs. Decentralized vs. Distributed Environment	3
Figure 2 - Example Blockchain Showing a List of Cryptographically Verifiable Records	4
Figure 3 - MSO SDN Architecture.....	5
Figure 4 - Example MSO Blockchain Architecture.....	7
Figure 5 - Security Factors Concerned with SDN Implementations	9
Figure 6 - Hash Block Chain & Merkel Tree	10
Figure 7 - HFC IOT Blockchain	11

Introduction

The number of connected devices in the future is expected to reach into the billions with the advent of the IoT, 5G, and the continued proliferation of smart devices. The extensive footprint of Multiple System Operator (MSO) Hybrid-Fiber Coax (HFC) networks will play an important role in rapidly expanding the connectivity of these devices across the globe. While virtualization and Software Defined Networking (SDN) reduce the network architecture complexity and provide a better way of processing and routing data, the security of such architectures to support these billions of devices, data integrity, and content privacy are still under question and will remain a key concern in upcoming years. Blockchain is emerging as a new way to address such security concerns through decentralizing the security construct and letting each connected device fundamentally become part of an overall security architecture.

Given the constant source of interest due to its decentralized secure way of transferring value or information with help of smart contracts or major industry sectors (i.e., telecom, banking, education, health-care, government, etc.), organizations are evaluating the ways in which Blockchain can be adopted in their areas of influence. The first ones are mostly financial organizations where highly-secure transactions play an important role. The core premise of Blockchain is to distribute the whole aspect of application or operation, where the operator¹ provides a simple convenient way to organize, manage and provide services to its subscribers. Figure 1 illustrates the concept of decentralization, where the left side shows a star topology with centralized authority for the network nodes, and the right side showing network nodes with a decentralized and a distributed configuration. Further, topology, network complexity and its applicability pave a path for the type of Blockchain implementation an operator decides to implement: Private (Enterprise), Semi Public or Public. The reason for opting for a Private Blockchain in communication networks is that mostly all the operations inside a network will greatly impact the transactions between network nodes and will have lesser amount of interaction with an end subscriber until the enterprises use cases like SDWAN gather. That is broadly, until Public Blockchain implementations for uses cases such 3rd party storage services and content services are offered by the operator and need to make the 3rd party provider and subscriber part of the Blockchain.




Figure 1 - Nodes & Links in Centralized vs. Decentralized vs. Distributed Environment

This paper discusses the impacts and use cases of Blockchain technology in cable architectures and broadly outlines the mechanisms on how Blockchain distributed ledgers and trust mechanism can help protect the integrity of the network.

¹ The one who owns the Blockchain.

Background

This section provides a basic overview of the key Blockchain and SDN concepts and technologies.

A. Blockchain

A Blockchain is a distributed database consisting of a continuously growing set of records which are referred to as “blocks”, each record is list of transactions and each transaction is signed. Cryptography is used to link blocks together through signing (one way hash functions that are encrypted with a key). Blocks include a hash of the previous block, with proof of work (or similar proof, i.e., Proof of Stake, time, etc.) that help verify the integrity of a transaction, transactional data, and a timestamp. This makes an interconnection between the blocks, thus creating a chain of blocks or a Blockchain. Altering any data in a block retroactively cannot be done unless all subsequent blocks are altered. Any unauthorized alteration of a block or transaction is easily identifiable as corrupted.




Figure 2 - Example Blockchain Showing a List of Cryptographically Verifiable Records

As previously mentioned, the three classical main pillars that Blockchain stands-on are, Transaction, block mining and distributed ledger, Transaction among network nodes is a transfer of modified configuration or a status change value that is broadcast to the network and is collected into blocks. Block mining is the process of adding earlier transaction records to the ledger of past transactions or Blockchain. Distributed ledger is a database which is shared and synchronized across the network. The Blockchain eco-system is enabled by these pillars to provide a holistic trust-based approach to secure transactions and the networks, largely reducing the barrier between trusted and untrusted aspects of networks. Translating these pillars to Telco network requirements the main pillars are transaction integrity (e.g., transactions are signed using asymmetric key cryptography), blocks are compiled that include a hash (may or may not be a signature) of the previous block, and the resulting blockchain is distributed amongst a sufficiently large network.

While Blockchain adopts the decentralized concept among peers for transparent information transmission, some characteristics of Blockchain, such as using all network entities to distribute Blockchain, might not be fully applicable to telecom networks. Blockchain might be helpful in securing SDN, cloud storage, virtualization, IoT and billing in MSO networks.

B. Software Defined Networking (SDN)

SDN decouples the network control plane and forwarding plane functions and enables network control to be programmable. The underlying network infrastructure is abstracted to both applications and network services. Control, decoupled from hardware, is implemented in software within SDN controllers. To illustrate this concept the following is provided. In traditional non-software defined networks, a data packet arriving at a switch or router is forwarded to a destination based on decisions made in firmware.. However, in SDN, such packet forwarding decisions are made by SDN controllers. For each packet entering a switch or router node, the SDN forwarding plane decides what to do with the packet. The SDN controller defines the flows which denotes the data itself. A set of packets transferring from source to destination (or set of endpoints) is characterized by a flow. Internet Protocol (IP) address, TCP/UDP port pairs, Virtual Local Area Network (VLAN) endpoints, layer three tunnel endpoints and input ports, etc. define the endpoints in SDN. The forwarding action that the SDN devices take into decision is determined by one set of rules which apply to all packets belonging to that flow. A flow is unidirectional in that packets streaming between those same two endpoints in the inverse direction could each constitute a separate flow. The Open Flow protocol helps to periodically collect information from network devices concerning their status along with commands involving how to handle traffic. Furthermore, an Orchestrator provides overall management of the different domains of a network from an end-to-end perspective. Refer to Figure 3.

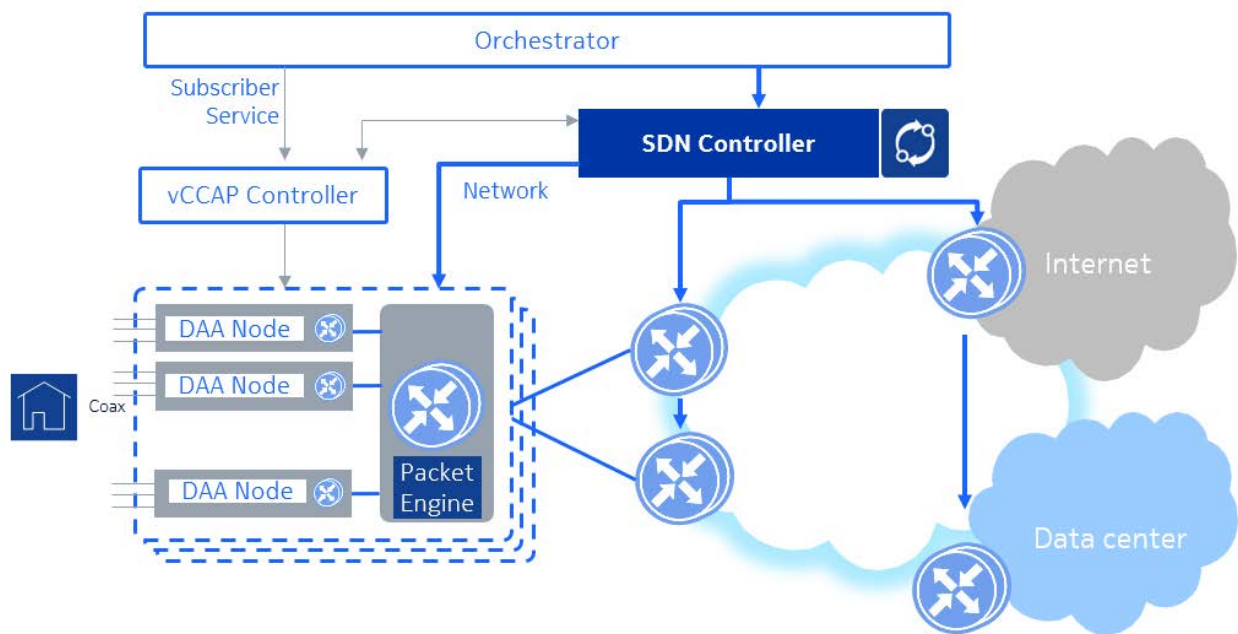


Figure 3 - MSO SDN Architecture

MSOs will move from a centralized Integrated Converged Cable Access Platform (I-CCAP) to a virtualized Distributed Access Architecture (vDAA) over time. Various Data Over Cable System

Interface Specification (DOCSIS®) functions will be disaggregated and distributed in a vDAA. In a Remote PHY (Physical) vDAA architecture, the MAC (Media Access Control) layer will be centralized in a cable hub/data center and the PHY layer will be distributed to a vDAA node located deep within the HFC Network. In a Remote MAC/PHY vDAA architecture, both the MAC layer and the PHY layer will be located at the vDAA node. SDN will enable greater flexibility and programmability of this emerging architecture. One such example is automatically configuring and administering complex DOCSIS profiles based on the network condition, with further possibility to have an end to end process by integrating together with Transport-SDN which automatically optimizes the transport network, based on varying characteristic in networks .

Blockchain Impact and Benefits on the HFC Network and Services

Blockchain technology may benefit future HFC architecture and services provided by an MSO. Future HFC architectures may enable decentralized business models. This may seem an odd assertion for access networks. However, consider that emerging HFC networks, particularly with new DOCSIS 3.1 technology standard, bring massive bandwidth businesses. In the WAN market place, we've seen dramatic adoption of virtualization technologies. Enterprises in terms of size and locations are getting more dependent on connected resources every day to tackle the customer expectations and dynamic management of WAN bandwidth and resources is becoming critical to their everchanging game for their business networks. For example, if we look at the growth of virtual operators to Slicing concept in the Mobility world we can relate the same here with our HFC networks, given that complete Virtualization and cloudification of the networks is ongoing and achievable, where once its enabled a HFC owner that has deployed Slicing could offer an tenant a dedicated slice of its network for carrying traffic, but further this tenant is also enabled to deploy its own VNF's (virtual network functions), helping both the operator and virtual operator achieve better flexibility, reliability & savings.

As HFC scales, we should anticipate similar technologies and business solutions. This leads to defining a new plane in our access infrastructures, the business plane, which together with the existing management, control and data planes, will provide secure and automated service enablement. In such architectures where an MSO interacts with other operators or content providers, it will be useful to manage incentives using consensus processes, such as proof of work, to support a variety of distributed network functions. This results in an access platform for new business services. Some examples are shown below.

- Media Content services: Leasing content from the content provider.
- Storage as a Service: Leasing storage to another enterprise.
- Platform as a Service: Leasing solution stack to enterprise.
- Infrastructure as a Service: Leasing necessary hardware, storage etc. to enterprises.
- SD-WAN: Overlay connectivity option for enterprise using SDN or similar concepts.

Apart from the business services, which will have a major impact/influence of how the client is enabled and managed, the HFC network will be impacted due to ongoing virtualization of the headend, IOT infrastructure creation, and SDN enablement. The last two are more practical and have larger influence in terms of Blockchain implementation.

A main aspect of Blockchain on the network side which makes it suitable for HFC networks is the provisioning and feedback mechanism performed in a better way. Though we have 2-way feedback mechanisms with SNMP, NETCONF, etc., considering the way SDN works (i.e., depending on network status), changes triggered by SDN controller might impact multiple nodes to ensure the integrity of a configuration across the addressed nodes.

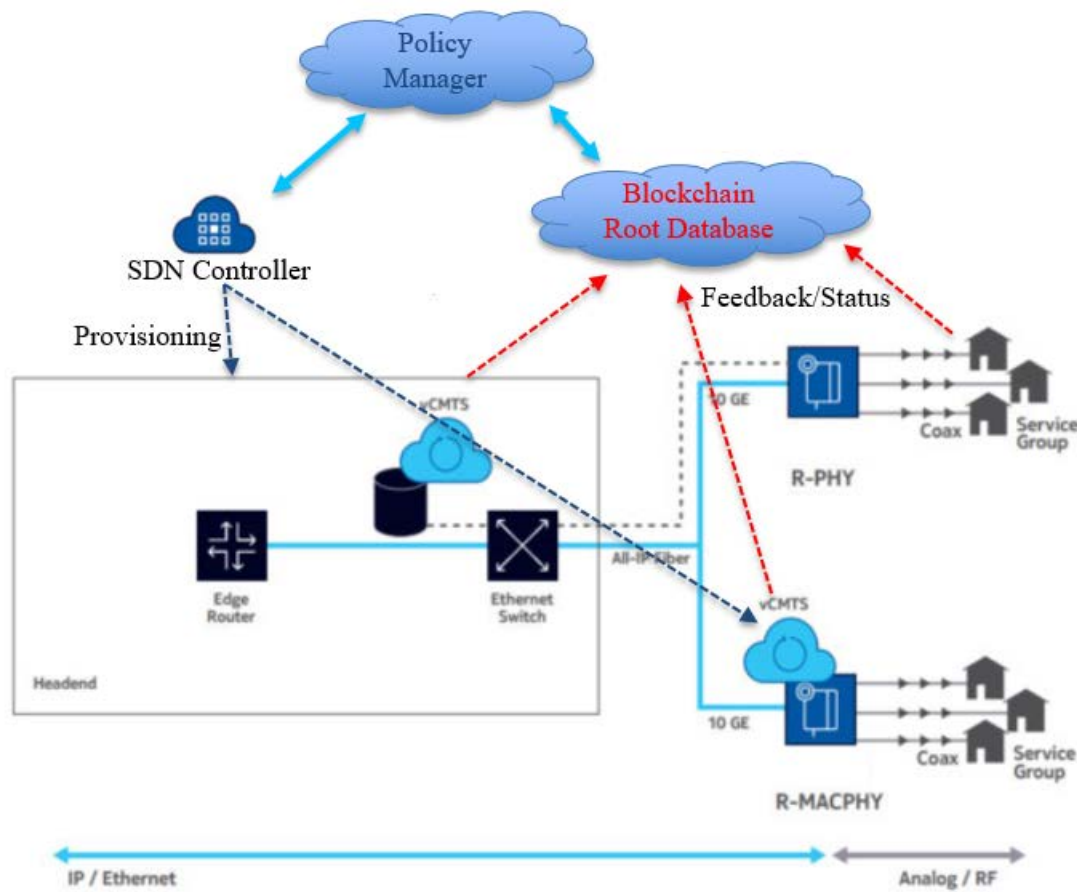


Figure 4 - Example MSO Blockchain Architecture

Similarly, in case of the business services, the resource configuration for the enterprise or subscriber can be provisioned based on the feedback/status principles of Blockchain to link heterogeneous resources with heterogeneous needs in a digitally enabled world. This in turn means that Blockchain provides a more secure and transparent way to manage the network and services benefiting the network owner and the end user.

As depicted in Figure 4 above, the architecture brings two new nodes into the network, the Policy Manager and the Blockchain Root Database. The Policy Manager sends direction to the SDN Controller or SDWAN Controller or Access Controller, to act on the given requirement i.e., setup circuit between NodeA and NodeB with so and so bandwidth, it will help program automated behaviors in a network to coordinate between the required hardware and software elements to support different applications and services. The Policy Manager role can be integrated with the Orchestrator (Figure 3), and whenever any change in configuration is triggered across a network between multiple nodes, the Policy Manager queries the Root Database for the last hash value to confirm the integrity and authorizes an action. The change it will have in comparison to normal flow is that, here after successful execution of the command, the

triggered changes and participating nodes generate a top hash between them and that gets stored in root database for a particular time t and is queried each time before executing a new command to ensure the integrity of network, the generally, The Blockchain Root Database contains the hash values for referred transactions at particular time intervals using a Merkle tree. This concept is further explained in the next section.

The high-level benefits that Blockchain provides to an MSO are:

- 1) Decentralized Network data: The HFC network data is stored off-chain in a distributed way facilitating multi-party trust, and a participating node i.e., router, vDAA can easily find the storage address through the Blockchain i.e., quicker response, minimal impact during failure of one or more nodes.
- 2) No Centralized Trust based mechanism: The access to HFC data is controlled by the majority of the Blockchain enabled network entities (modems, RPDs, CMTSS that can originate or relay transactions as part of a blockchain network), without any intervention from single trusted source i.e., rather than single source of trust mechanism, whole network works in sync mode as trust enabler.
- 3) Traceability and Accountability: Activities such as accessing and modifying the HFC network configuration data, can be recorded by the Blockchain. No malicious attempts can go undetected.
- 4) New Business Services: Considering Multi-tenant reality over period of time, BC will help operators in ensuring their network integrity and ease of operations.

Blockchain Design for SDN Implementation

As SDN is steadily being deployed, various security attack vectors could also penetrate SDN implementations, examples of known ones are listed below:

- Malicious SDN applications.
- Malicious controller creating entries in the flow tables of the network elements, thus gaining complete control of the network.
- Malicious network element, or a hacker posing as an administrator.
- Unauthorized access to an SDN framework.
- Unauthorized configuration, network or topology change.
- Attack on SDN function causing service disruption.

While the Open Networking Foundation (ONF) provides several recommendations including ONF TR-511, TR-529 and TR-530 for securing SDN, Blockchain can help strengthen it further and mitigate the threat created by SDN's centralization of control - a key security concern with the SDN model. SDN's centralized control model raises security concerns as described below, given it can become a single point of security attacks that can have disastrous results. Blockchain can improve security and mitigate threats using hash-values

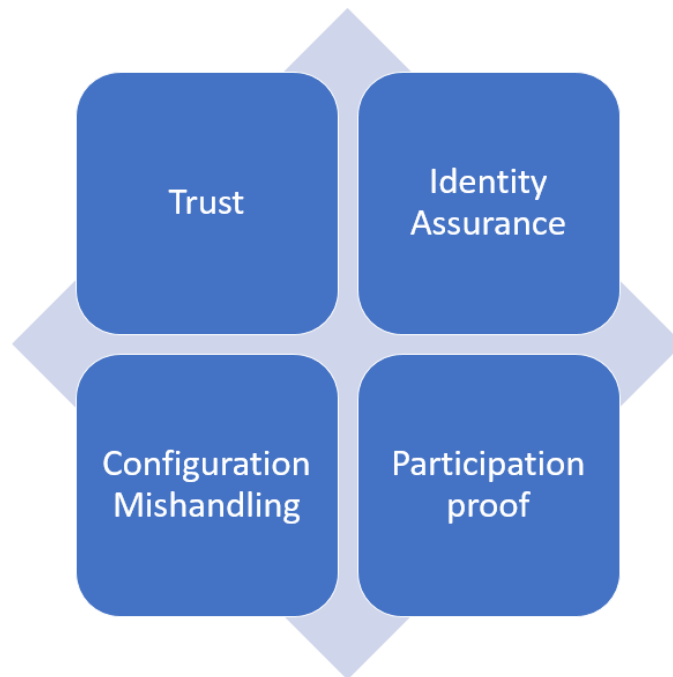


Figure 5 - Security Factors Concerned with SDN Implementations

and binary trees, irrespective of where an asset is located or where data is stored. Every node (vDAA router, switch, etc.) and its related configuration, operator, or via automatic function can be tagged, tracked and located with real-time verification independent of trusted administrators. Blockchain provides a system wherein the need for trust can be eliminated.

Using a Merkle tree implementation as shown in Figure 6, a distributed binary tree can be periodically generated using hash-values of data generated in the MSO network. Two input values, along with required parameters, are concatenated and run through a hash function. This process is iterated, resulting in a single root hash value. Shared secrets are still used for authenticating clients during the signature validation process, keys are not needed every time for the signature verification itself. The integrity of the signatures is protected using hash functions thus reducing the repetition of key exchange and need of the signature verification each time.

Finally, the root hash is calculated and stored in top database or root database which consists the hash output of top chain for any changes propagated by policy manager at given time t and is broadcasted to all other nodes (vDAA node, router, switch, etc.). For every hash value included in the tree, there is a unique hash-chain, or series of hash-values that allows the root hash-value to be recreated. This hash chain is returned and stored as the signature. A signature identifies the node or configuration through the hash tree, from the node's own hash value, up to the root database consisting of a complete transaction record. With access to the root database, anyone (nodes), anywhere (any part of the network), can receive data and verify the signature, which includes indications of time, identity and integrity. The process is unique and one of concepts in which the integrity of the transaction i.e., the change propagated to the network nodes gets covered with hashing and in case of subsequent changes in the network where the policy manager queries the node and root database gets verified with policy manager private key and Nodes public keys in turn providing the non-repudiation.




Figure 6 - Hash Block Chain & Merkel Tree

By integrating a Private Blockchain model, the components of an MSO SDN network will be able to sign and verify data as it moves between components. This provides the MSO with a data integrity infrastructure in which data can be verified in near-real time. The data residing in the configuration storage is regularly monitored for its consistency and verified against the associated signatures that were created upon the creation of the configuration data. The design of this chain can be implemented using the hierarchical model used in the current generation of networks, where the hashes at different layers of network are aggregated and processed at each layer and the top database (Database holding computed hash values of the chain relevant to specific portion of network for a time intervals) is settled and stored at the MSO datacenter.

Dealing with multiple services, deploying SDN, network slicing, and virtualization makes this ideal for MSOs.

Blockchain Design for IOT Implementation

Unlike SDN, in the IOT Implementation of Blockchain, subscriber CPE at the originating point of the Blockchain can be used to manage the IOT device configuration, maintain integrity of the received sensor data, possibly enable micro payments (payments triggered for using any specific content, or storage) based on need. The utmost concern that had been shown by operators is how to leverage the Blockchain as service, With IOT, the constraints are much more varied and misunderstood. Blockchain can be enabled as a service on the end IOT platform running over the HFC networks.

While implementation can be achieved with the previously explained SDN model, a key question arises, where should the Blockchain be hosted? Hosting the Blockchain directly on IoT devices is not possible due to resource constraints, although this requires implementation and a need to deploy better computational resources in the network. These resources will cater to the need of the IOT application host. The host collects a consistent set of IOT sensor data for analytics, which preferably can be

implemented at various boundaries of network layers as well as implementation of cloud datacenter for this data hosting.

An IOT application provider will look for opportunities to create/store/transfer digital assets inside any practical MSO network, adding value for the provider, and potentially providing a new business stream for the MSO.

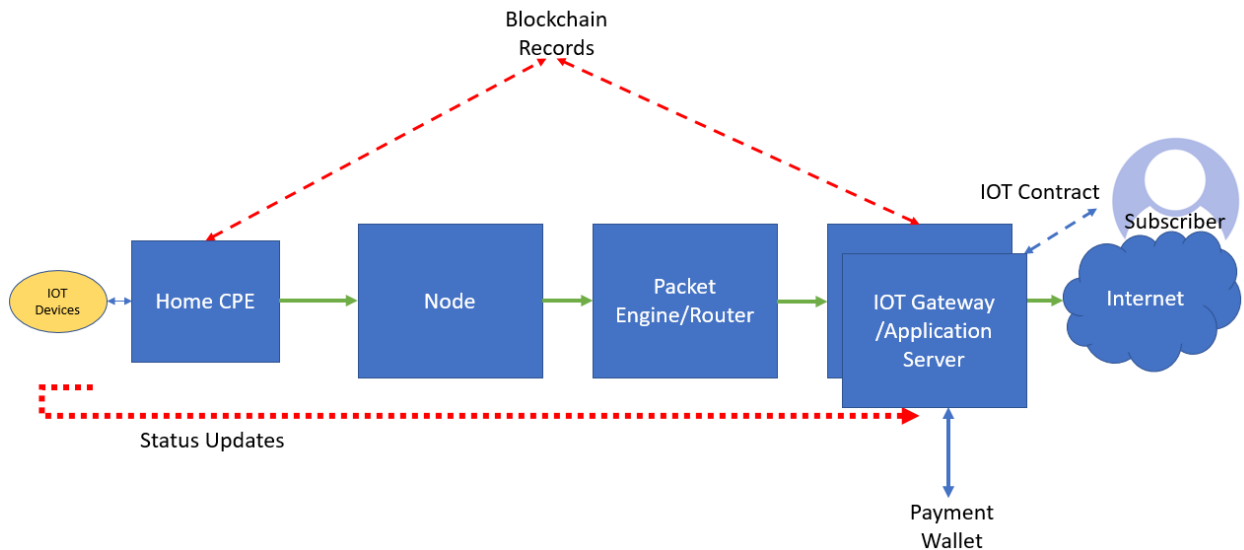


Figure 7 - HFC IOT Blockchain

Possible High-Level Transactions or Use Cases over the Blockchain enabled HFC network includes:

- Smart contracts (a digital contract which facilitates, verifies or enforce the negotiation or performance of a contract between two parties).
- IOT provisioning.
- IOT management.
- On demand service/update sharing.
- End user security.

Some traditional MSO use cases include:

- Financial settlements (e.g., peering and transit) for backhaul services MSOs provide to MNOs.
- Settlement of MVNO minutes - MSO and MVNO.
- Financial settlements with vendor supply chain.
- Consumer credential mobility – i.e., a Cable customer typically is given free WiFi access in the MSO footprint) requiring log-in and authentication using the MSO credentials. Blockchain can help enable easy and automatic authentication as well as enable mobility between two MSO WiFi networks if they choose to enable such a feature.
- Data Integrity - Of long-term databases, providing protection of subscriber databases from hacking.
- Asset Management – parts and spares inventory tracking and management.

Example HFC enablement for Blockchain

The following table contains a high-level roadmap and impact of Blockchain convergence with HFC networks:

Implications/Period	Phase-1	Phase-2	Phase-3	Phase-4
Key Developments/Enablers	Studying and Leveraging SDN in HFC Networks from operations point of view	SDWAN and IOT Offerings	Key developments and Solution design for Sustained introduction of Blockchain in HFC Networks to secure internal network assets	Introduction of Blockchain as a service Integration with existing ecosystem
Outcome	E2E Management Ease of operations Better fault handling	New revenue potential Surveying the slicing needs for HFC	Improved network security Better compliance to the integrity of services provided and data used for storage and analysis Cost efficiencies due to shared resources & processing	New revenue streams Predictive subscriber needs

Conclusion

This paper explored the combination of Blockchains, SDN and IOT over MSO networks which can be beneficial and help secure the networks in more pragmatic and simplistic way.

While Blockchains provide resilient, distributed peer-to-peer systems together with SDN, IOT and content transactions, it helps in automating workflows with new methods and flow, achieving trust, with significant cost and time savings in the process. The Root Database provided by Blockchain in an SDN network used for a network integrity check makes the transactions in the SDN environment more robust against untrusted members outside the Blockchain. Similarly, together with IOT, Blockchain can be utilized as another service enabler.

Continued research and new implementation models will bring about new business models in the security domain.

Abbreviations

CPE	customer premises equipment
DOCSIS	Data over cable system interface specification
HFC	hybrid fiber-coax
IOT	Internet of Things
ISBE	International Society of Broadband Experts
MAC	media access control
MNO	mobile network operator
MSO	multi service operator
MVNO	mobile virtual network operator
ONF	Open Network Foundation
PHY	physical layer
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined networking
SD-WAN	software defined wide area network
vDAA	virtualized distributed access architecture

Acknowledgments

- Martin Glapa, Partner & Bell Labs Fellow, Bell Labs Consulting, USA.
- R.J Vale, Principal, Bell Labs Consulting, USA.
- Ben Tang, Principal, Bell Labs DMTS, Bell Labs Consulting, USA

Bibliography & References

A Simple Overview of Blockchains, Why They Are Important to the Cable Industry. Steve Goeringer. SCTE-ISBE. 2017 <https://www.nctatechnicalpapers.com/Paper/2017/2017-a-simple-overview-of-Blockchains-why-they-are-important-to-the-cable-industry/download>

Decentralized access control mechanism with temporal dimension based on Blockchain. Mayssa Jemel and Ahmed Serhrouchni. The Fourteenth IEEE International Conference on e-Business Engineering. 2017

Blockchain in Telcos: A tool for openness. Dimitris Mavrakis. 2017

Introduction to Blockchain. Javier Antich Romaguera, Juniper. MPLS+SDN+NFV. 2018

The Role of SDN in Broadband Networks. Hassan Habibi Gharakheili. 2016

Blockchain basics: A non-technical introduction in 25 steps. Daniel Drescher. 2017

Merkle Signature Schemes, Merkle Trees and Their Cryptanalysis. Georg Becker. Seminararbeit Ruhr-Universit at Bochum. 2008

Demystifying Bitcoin and Blockchain. Ganesh Kondal. 2016

A Framework for Determining Blockchain Applicability. Brian A. Scriber, Cabelabs. IEEE. 2018

A Distributed Blockchains-Based Secure SDN Architecture for IoT Networks. Pradip Kumar Sharma; Saurabh Singh; Young-Sik Jeong ; Jong Hyuk Park. 2017

An Innovative Security Architecture for Low Cost Low Power IoT Devices Based on Secure Elements. Urien, P. IEEE CCNC. 2018.

Bitcoin and Beyond: A Technical Survey on Decentralized Digital Currencies. F. Tschorsch and B. Scheuermann. IEEE Commun.Surveys & Tutorials. 2016.

The Blockchain as a Software Connector. X. Xu et al. IEEE/IFIP Conf. Software Architecture. 2016.

Comparing Blockchain Implementations. Zane Hintzman. SCTE-ISBE. 2017

The Future of Fixed Access:

A Techno-Economic Comparison of Wired and Wireless Options to Help MSO Decision Process

A Technical Paper prepared for SCTE•ISBE by

Jean-Philippe Joseph
Principal, Fixed Access
Bell Labs Consulting, Nokia
600 Mountain Avenue, New Providence, NJ 07974
+1 908 679 5798
Jean-Philippe.Joseph@bell-labs-consulting.com

Amit Mukhopadhyay
Partner, Wireless Networks
Bell Labs Consulting, Nokia
600 Mountain Avenue, Rm 2A-402, New Providence, NJ 07974
+1 908 956 4744
Amit.Mukhopadhyay@bell-labs-consulting.com

Ashok Rudrapatna
Principal, Wireless Access
Bell Labs Consulting, Nokia
600 Mountain Avenue, New Providence, NJ 07974
+1 908 432 3445
Ashok.Rudrapatna@bell-labs-consulting.com

Carlos Urrutia-Valdés
Principal, Wireless Networks
Bell Labs Consulting, Nokia
600 Mountain Avenue, New Providence, NJ 07974
+1 908 679 5671
Carlos.Urrutia-Valdes@bell-labs-consulting.com

Tom Van Caenegem
Senior Consultant, Fixed Access
Bell Labs Consulting, Nokia
Copernicuslaan 50, Building 2018 Antwerp, Belgium
+32 3240 7617
Tom.VanCaenegem@bell-labs-consulting.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Background	4
Service Targets	6
1. High-bandwidth service targets.....	6
2. Low-bandwidth service targets.....	7
Technology options	7
1. Copper using G.fast/x digital subscriber line (xDSL)	7
2. Hybrid Fiber-Coax (HFC) using data over cable interface specification (DOCSIS).....	8
3. Fiber using Passive Optical Network (PON)	8
4. Fixed Wireless Access (FWA) based on 5G (mmWave and mid band)	8
5. FWA based on WiGig.....	8
FWA performance modeling	9
1. mmWave system on a utility pole.....	9
2. mmWave system on a macro tower	10
3. Mid band system on a macro tower	10
4. WiGig on utility pole.....	11
Cost modeling for different technologies.....	11
1. Metropolitan deployments	11
1.1. Overview	11
1.2. Modeling assumptions	13
1.3. Analysis Results	13
1.3.1. Baseline	13
1.3.2. Sensitivities	14
2. Rural deployments	17
2.1. Overview	17
2.1.1. VDSL2 with fiber backhaul (Cu+Fiber)	18
2.1.2. VDSL2 with copper + fiber backhaul (Cu+Cu+Fiber)	18
2.1.3. VDSL2 with copper + microwave backhaul (Cu+Cu+MW)	18
2.1.4. Fixed Wireless Access (FWA)	19
2.1.5. Fiber to the home from the central office (CO-FTTH).....	19
2.2. Modeling assumptions	19
2.3. Analysis results	20
2.3.1. Impact of household density	20
2.3.2. Impact of service peak bit rate	21
2.3.3. Impact of fiber install cost	22
2.3.4. 4G FWA versus 5G FWA.....	23
Conclusions.....	24
Abbreviations	24
Bibliography & References.....	26

List of Figures

Title	Page Number
Figure 1 - Rate and range of fixed access echnologies.....	5
Figure 2 - Fixed Wireless Access business growth [The Carmel Group]	5
Figure 3 - Throughput vs. ISD with pole mounted small cells in 28 GHZ with 400 MHz BW	9
Figure 4 - Capacity vs. ISD for 28GHz macro deployment with 800 MHz spectrum	10
Figure 5 - Capacity vs. ISD with 4 GHz band and 100 MHz spectrum.....	11
Figure 6 - Metropolitan areas deployment solution options	12
Figure 7 - Metropolitan zones of advantage for gigabit access	14
Figure 8 - TCO breakdown for each solution.....	15
Figure 9 - Zones of Advantage, with trenching (left) and without trenching (right)	16
Figure 10: Rural scenarior deployment solutions considered	18
Figure 11 - Rural deployment CapEx per HHC as function of HH density	20
Figure 12 - Rural deployment CapEx per HHC as function of peak rate	21
Figure 13 - Rural deployment CapEx breakdown per HHC.....	22
Figure 14 - Rural deployment impact of fiber install cost.....	22
Figure 15 - Rural access ZoA with FWA LTE (left) and FWA 5G (right)	23

List of Tables

Title	Page Number
Table 1 - VR Bandwidth Requirements.....	6
Table 2 - TV Resolution and Throughput.....	7
Table 3 - Rural deployment assumptions that vary with HH density	19

Introduction

Recent major advances in centimeter/millimeter wave, massive Multiple In Multiple Out (MIMO) antennas, beam forming, hybrid radio technologies, and new systems such as 5th Generation (5G) wireless, Wireless Gigabit (WiGig) have accelerated Fixed Wireless Access (FWA) solutions to become an alternative to wired solutions such as Hybrid-Fiber Coax (HFC), fiber, and copper for providing ultra-broadband access to residences and small to medium-sized businesses. Innovative spectrum solutions that include unlicensed and shared regime, besides the licensed spectrum, further enhance the attractiveness of FWA. These advances have facilitated both gigabit per-second service for high-end subscribers in metropolitan areas as well as tens of megabits per-second peak service in lower housing density and rural areas at competitive costs. Prior to these advances, FWA was a viable solution only in providing lower data rate services in certain niche markets, such as rural and in developing countries.

However, FWA, even with the recent advances, will not be a solution of choice for all end-user requirements, nor in all usage scenarios. This paper provides a techno-economic analysis comparing different FWA and wired technologies including HFC, fiber, and copper under different deployment scenarios to establish the relative “Zones of Advantages” for each solution. It identifies the optimal technology of choice for a given deployment, considering factors such as throughput requirements, household densities, morphology, deployment conditions, take rates, capital and operational expenses (CapEx and OpEx).

Background

Fixed access continues to be a key business segment for Communications Service Providers (CSPs) and it is expected to remain relevant in the foreseeable future, given the predicted growth in demand for ultra-high bandwidth services like 12K and 16K or volumetric Television and Virtual Reality (VR) with full head and body movement. Such services may generate multi-gigabit per second (Gbps) throughput per user. On the other end of the spectrum, there is a substantial population in the country that live in low housing density areas where it is challenging to provide a few megabits per second (Mbps) connectivity economically, even though much higher demands exist.

Historically, higher-bandwidth fixed services in metropolitan and suburban areas with higher population densities have been provided through fixed access technologies like copper, HFC or FTTH. In rural areas with lower population densities, the primary vehicle for providing low-bandwidth fixed access service has been copper and wireless technologies (both terrestrial as well as satellite).

Recently, two fundamental developments are shifting the dynamics of the solutions. First, improvement in trenching technologies has driven the cost of fiber deployment much lower; this has positively impacted the techno-economics of all the fixed access technologies, thus enabling higher throughput at lower cost points. Secondly, major developments in wireless technologies like Long Term Evolution Advanced (LTE-Adv), 5G and WiGig have made it technically feasible to deliver multi-gigabit per second services over a fully wireless connection as illustrated in Figure 1.




Figure 1 - Rate and range of fixed access echnologies

FWA is expected to grow significantly over the next few years, as predicted by The Carmel Group [1] and depicted in Figure 2.




Figure 2 - Fixed Wireless Access business growth [The Carmel Group]

However, over-all cost and performance of different fixed access solutions vary tremendously and establishing which technology suits best under which circumstances, i.e., a zones of advantage (ZoA), is a challenge. To compare solutions in a meaningful way, we refer to solutions for two different target services – higher-bandwidth services for metropolitan and suburban housing densities, and lower bandwidth services for rural housing densities.

The rest of the paper is structured in the following way. The reference target service rates are established first, followed by a description of various technology options to deliver target services. We then present key performance modeling results for emerging wireless technologies. Finally, we present a cost modeling analysis and associated results for two deployment options in metropolitan and rural areas

respectively. Since the challenge is greater in lower housing density areas, this paper provides more details for the rural deployment. We bring the analysis together in the final section to draw over-all conclusions.

Service Targets

Before delving into solution architecture details service targets need to be established. Two distinct service criteria are defined – peak throughput and sustained throughput. Peak throughput is the highest instantaneous bandwidth capacity demand per household. Sustained throughput is the average bandwidth capacity demand during the busy hour per household. In [2] such sustained throughput requirements have been calculated from a bottom-up perspective. The average sustained throughput per household is expected to be a few megabits per second, even though peak throughput may be several hundred megabits or even multiple gigabits per second.

Throughput is often (mistakenly) associated with technology advancements or advertised speeds. For example, one may associate a 1 Gbps service over a gigabit HFC network with the actual usage of the customer, which is far from reality. In a typical HFC network, even though a subscriber may be able to burst at 1 Gbps occasionally, the engineered capacity for sustained throughput of all users on the system may only be a few hundred Mbps.

For our discussions, we consider true user requirements for peak and sustained throughputs as those parameters driving a solution's techno-economic feasibility. In addition, planning for fixed access networks should consider subscriber needs for the next five to ten years as technology investments typically have long payback periods.

1. High-bandwidth service targets

For high-bandwidth services, we consider delivery of VR applications with full eye, head and body movement. As Table 1 below shows, the throughput requirement can vary widely, depending upon various parameters.

Table 1 - VR Bandwidth Requirements

	Field of View	Eye and Head movement			Eye, head and body movement		
Parameters	Basic	Basic	Codec improv.	Higher refresh	Basic	Codec improv	Acuity, refresh imp.
Visual Acuity (Arc-minute)	1	1	1	1	1	1	0.7
FOV (HXV degrees)	30X30	120X150	120X150	120X150	180X360	180X360	180X360
Pixels	3.2M	64.8M	64.8M	64.8M	233.2M	233.3M	476.1M
Frame rate (per sec)	60	60	60	120	60	60	120
Pixels/sec	194.4M	3.9G	3.9G	7.8G	14G	14G	57.1G
Coding (bits/pixel)	0.125	0.125	0.08	0.08	0.125	0.08	0.08
Throughput (Mbps)	24.3	486	311.04	622.08	1,749.60	1,119.74	4,570.38

We consider a peak user throughput requirement of 1 Gbps for our modeling purposes. In reality, a fraction of users will be active during the busy hour and such individual users will likely use the service for a fraction of the busy hour period.

For sustained throughput requirements, we will refer to [2] and consider user requirements at the 99th percentile as these high-bandwidths will most likely be targeted in metropolitan areas where competition will be fierce. The target sustained throughput requirements are 70 Mbps per user for designing a network five to ten years down the road.

2. Low-bandwidth service targets

For low-bandwidth service targets we assume that the peak throughput demand will be driven by a mix of limited VR applications and high-end TV resolutions as shown in Table 2.

Table 2 - TV Resolution and Throughput

Parameters	FHD(2K)	UHD(4K)	8K	12K
Display	1920X1080	3840X2160	7640X4320	12288X6912
Pixels	2.1M	8.2M	33.0M	84.9M
Refresh rate (per second)	30	60	120	120
Raw bit rate (per second)	62.2M	497.7M	3.98G	10.2G
Bits per pixel**	0.08	0.04	0.03	0.018
Throughput (Mbps)	4.98	19.91	119.4	183.5
YouTube 24/30fps (Mbps)	4.5			
Netflix 24fps (Mbps)	4.8*	~15.0 (24fps)		
Broadcast 1080p (Mbps)	6 to 8			

We assume a 100 Mbps peak throughput for the low bandwidth future service target, which is substantially higher than the few 10's of Mbps service available today in sparsely populated areas. For sustained throughput requirements we again refer to [2] but limited to the 75th percentile (rather than the 99th percentile for the high-bandwidth services since most likely these bandwidths will be targeted to rural areas where competition is less) and set a service target of 25 Mbps.

Technology options

Different technology options for delivering fixed access services are explored. Some options are evolutions of older technologies to increase performance, some options are in the bleeding edge of evolution.

1. Copper using G.fast/x digital subscriber line (xDSL)

G.fast [3, 4] technology enables delivery of gigabit speeds over copper loops for distances up to 100m using Orthogonal Frequency Division Multiplexing (OFDM) in a 212 megahertz (MHz) frequency spectrum. This requires deploying fiber deep to distribution point units (DPUs) to keep copper loop length under 100m.

Different versions of DSL, Asymmetric Digital Subscriber Line (ADSL), Very high bit rate (VDSL), etc., commonly known as xDSL can be used [5, 6] to deliver lower bandwidth services over shorter copper loops.

2. Hybrid Fiber-Coax (HFC) using data over cable interface specification (DOCSIS)

DOCSIS provides asymmetrical high-speed data services on HFC networks. Multiple versions exist, with DOCSIS 3.1 [7] deployments underway to provide Gbps downstream bandwidths and multi-hundred Mbps upstream bandwidths in conjunction with a migration to deeper fiber and smaller fiber nodes. The move to Distributed Access Architectures (DAA) will push fiber even deeper as nodes get closer to users and enable higher per subscriber sustained bandwidths. Full-Duplex (FDX) DOCSIS [8, 9] will bring multi-Gbps symmetrical service bandwidth in the future.

3. Fiber using Passive Optical Network (PON)

Most FTTH solutions are based on PON architectures where the point-to-multipoint (P2MP) outside fiber plant has no active elements and the capacity is typically shared across up to 64 or 128 subscribers via a tree and branch. Multiple PON technologies are available and can co-exist on the same PON. Time Division Multiplexing (TDM)-PON, which includes GPON [10], GE-PON, XG-PON [11], 10G EPON, uses a passive splitter to connect multiple users or optical network units (ONU) to one optical line termination (OLT) port. Wavelength Division Multiplexing (WDM)-PON uses a passive wavelength router that enables a logical one-to-one channel mapping between the ONU and the corresponding OLT port. Time-Wavelength Division Multiplexing TWDM)-PON(e.g., NG-PON2) uses a splitter but combines multiple (typically 4) TDM wavelength pairs on a given fiber. As a result, TWDM-PON achieves even higher throughputs (e.g., symmetric 40Gbps).

4. Fixed Wireless Access (FWA) based on 5G (mmWave and mid band)

5G technology is the latest evolution from a large family of mobile wide area systems, predecessors of which include 2nd generation (2G), 3rd generation (3G), and 4th generation (4G)/LTE [12]. 5G [13, 14] employs key technologies to enable very high service targets compared to any previous technology:

- Wider carriers: Up to 1 GHz of spectrum can be combined in a channel to deliver higher bandwidth.
- Higher frequency operating bands: since most lower bands (below ~2.5 gigahertz (GHz)) have already been exploited, 5G is expected to be deployed around 3.5 GHz spectrum, often referred to as “mid band”, for wide area coverage. Much larger spectrum is likely to be freed up in centimeter (cm) or millimeter (mm) wave bandwidths (e.g., 24, 26, 28, 39 GHz or even higher).
- Massive multiple input, multiple output (mMIMO): Antenna dimensions are inversely proportional to frequency. Thus, in the higher cm and mm wave frequencies, antenna elements become quite small. This size reduction can be exploited to enable mMIMO arrays that help overcome higher path and penetration losses. They also enable advanced beam forming and higher order MIMO techniques. These techniques help improve coverage and capacity.

5. FWA based on WiGig

WiGig is a technology standardized by IEEE as 802.11ad [15], and is often referred to as “Gigabit Wi-Fi”. WiGig’s first applications have been focused on in-home use as a wireless replacement for a high-definition multimedia interface (HDMI), delivering up to 8 Gbps. It also has traction for use as a “last mile” broadband access solution, with first product availability in 2018, referred to as wireless PON (WPON).

WPON relies on a line-of-sight operation between the customer premise equipment (CPE) and the access point (AP), due to its operation in the 60 GHz frequency spectrum where radio signals fade very quickly with increased distance between CPE and AP. This is especially true if they are not on direct Line of Sight (LoS). WiGig leverages mMIMO technology and through phased array antennas, narrow beams can be formed pointing towards the CPEs, thus improving the signal quality.

WiGig FWA (also known as WPON) products come equipped with a wireless relay capability feature, where APs can connect with each other via self-backhaul. Either dedicated WiGig channels are employed for the backhaul relay link or the channel(s) is/are shared across the drop links (to the connected CPEs) and the other AP/distribution nodes using time division multiplexing (TDM). This wireless relay function is a crucial element for reducing total cost of ownership (TCO) across all pole-based FWA solutions as it enables savings on civil works for fiber backhaul installation.

FWA performance modeling

Key FWA performance modeling aspects are described in this section. Modeling results vary tremendously, based on morphology, LoS, AP height, CPE location and target sustained and peak throughputs [17]. Since these variabilities are less associated with copper, HFC or FTTH and the performance of those technologies is much more predictable their modeling is skipped in the interest of brevity.

1. mmWave system on a utility pole

The first example is a utility pole mounted AP using 400 MHz of spectrum in the 28 GHz band delivering services to an externally home mounted antenna in vLoS (vegetation LoS) conditions. Figure 3 shows the cumulative distribution function (CDF) for varying inter-site distance (ISD) between the poles. With 228m ISD, 80% of users can be served with peak throughput of 1 Gbps or higher.




Figure 3 - Throughput vs. ISD with pole mounted small cells in 28 GHz with 400 MHz BW

2. mmWave system on a macro tower

Figure 4 illustrates coverage for a 28 GHz system when the service is provided from a macro tower (25m height) in a suburban environment. Similar results are available for dense urban and urban morphologies as well but are left out of this paper for the sake of practicality. At 0.5 km ISD, 22 users can be simultaneously served with 100 Mbps service, i.e., sector throughput of 2.2 Gbps and a sector throughput of >4 Gbps is achievable with an ISD of 0.2 km or less. Under typical RF conditions, this would imply a peak throughput of 4Gbps or higher is achievable (under lightly loaded conditions) for 80% households within the coverage area with a 228m ISD.




Figure 4 - Capacity vs. ISD for 28GHz macro deployment with 800 MHz spectrum

3. Mid band system on a macro tower

Figure 5 illustrates performance simulation results for a 4 GHz system in a rural environment. Since greater coverage is achievable at 4 GHz, the results are more applicable to lower household density areas. The amount of spectrum, peak and sustained bandwidth available is less than other models. At 10 km ISD, 15 users can be simultaneously served with 25 Mbps service, i.e., sector throughput of 375 Mbps and a sector throughput of >600 Mbps is achievable with an ISD of 6 km or less. Under typical RF conditions, this would imply a peak throughput of 600 Mbps or higher is achievable (under lightly loaded conditions) for 80% households within the coverage area with 10 km ISD.




Figure 5 - Capacity vs. ISD with 4 GHz band and 100 MHz spectrum

4. WiGig on utility pole

Typical WiGig performance, using a 1x2.15 GHz channel, will provide 1 Gbps at a distance up to 100m from the AP, even under heavy rain conditions, if the AP and the CPE are within LoS. These results have been obtained from field measurements. Higher WiGig capacities will be enabled through carrier aggregation (e.g., 2 x 2.15 GHz channels) specified in the next version IEEE 802.11ay standard. Since WiGig is deployed in unlicensed spectrum, some unpredictable performance degradation may occur over time if multiple operators start using the same spectrum in the same area.

Cost modeling for different technologies

The cost of delivering service targets in metropolitan and rural morphologies is determined in this section. Since the service targets and housing densities are vastly different for the two morphologies, the analysis focuses on technologies that are relevant for the given morphology. The selection of architectures and technologies is based on deployment practices as well as Bell Labs Consulting's experience of modeling different technologies.

1. Metropolitan deployments

1.1. Overview

The TCO of 4 different solutions enabling Gigabit access speeds are compared for metropolitan areas.

A street model for the metropolitan TCO is used with the following assumptions: all deployments are built and owned by the operator (no rental fees except for pole rental cost for FWA and WiGig deployments), housing is equally spaced and distributed along both street sides. We also assume that at the street entrance a fiber and (if required) a power feed point-of-presence (PoP) is available. The TCO of each solution is compared to serve one gigabit access speed (peak throughput) to each household, excluding the cost to bring and install the (feeder) fiber (or powering feed) to near the street entrance.

This provides for a comparative TCO analysis among Gigabit access solutions [18, 19] - all requiring a proximate fiber PoP - but where only the “last mile” is considered “in scope”. Further, having a fiber PoP at the street entrance may represent a common scenario where only the main streets are initially provisioned with one or multiple fiber cables (e.g., buried underneath the sidewalk). These fiber strands enable FTTH services or selective FTT-Building (FTTB) service solution in the main street (first). The presence of fiber at the entrance of the (side) streets crossing the main street, is then also the “enabler” for gigabit access delivery in these side streets, which for our analysis will be based on any of the following gigabit access-capable technologies: G.fast, DAA/DOCSIS, FWA based on WiGig and FTTH.

Figure 6 illustrates the deployment practice for each of the four considered technologies:

1. For G.fast, DPU nodes are deployed in pitches (small holes) on the sidewalks leveraging the fiber umbilical to the fiber distribution point in the PoP located near the entrance of the street. The drop-side of the DPU leverages existing copper loops to connect to the customer locations.
2. For HFC/DOCSIS, remote Distributed Access Architecture (DAA) nodes are deployed and leverage existing coax drops and taps to connect customer locations.
3. For FTTH, distribution fiber and splitters are installed along both street sides to pass 100% of the customer locations. Drop fibers are used from splitters to connect subscriber premises.
4. For WiGig FWA, fiber is extended to one or multiple AP locations. The AP can be deployed on one (or both) side of the street to achieve LoS. Subscriber premises are connected via a wireless link extending from the AP to an outdoor CPE.

5G-based FWA has been purposely left out of this metropolitan area analysis. It is expected that the 5G infrastructure deployed will be shared between mobile and FWA applications and only a part of the infrastructure cost will have to be apportioned to FWA for realistic cost comparisons. The degree of partitioning between fixed and mobile applications is still an open discussion in the industry and will vary greatly from operator to operator and market to market.

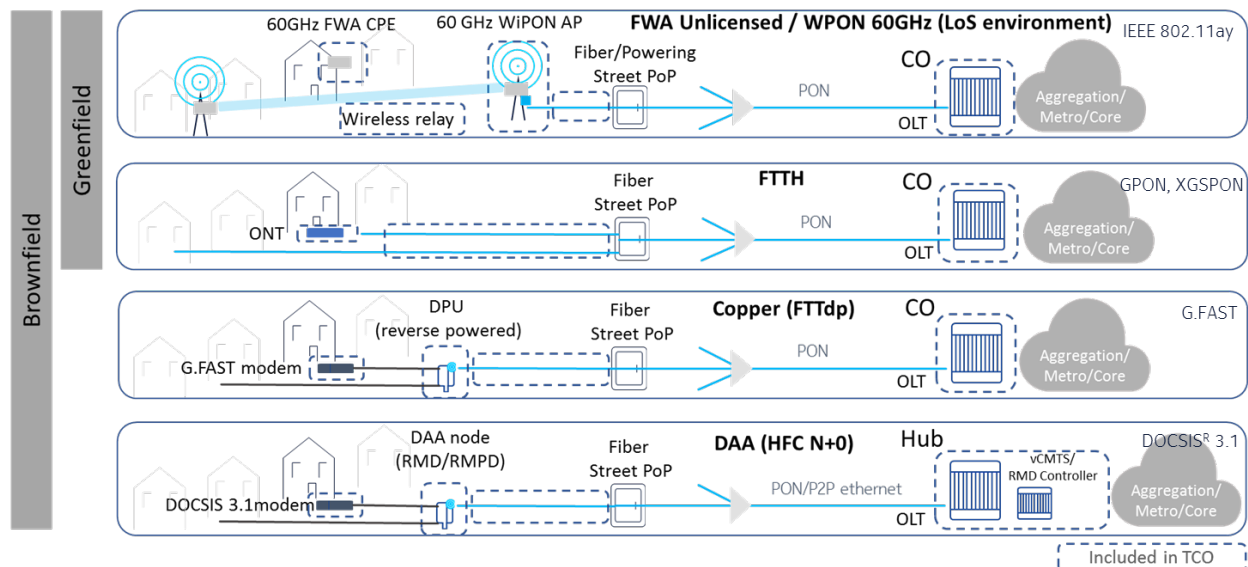


Figure 6 - Metropolitan areas deployment solution options

1.2. Modeling assumptions

The main assumptions are as follows:

1. 10-year TCO analysis covering capital expenditure (CapEx) and operating expenses (OpEx)
 - CapEx includes equipment and installation costs. All active equipment will be swapped and replaced by new generation equipment near year 10. Thus, equipment and installation costs are included twice -except for CPE installation.
 - OpEx covers annual recurring expenses such as outside plant maintenance, failing equipment repair/replacement (estimated at 2 or 4% failure rate depending on equipment type), equipment vendor support (i.e., licensing, maintenance contracts etc., set at 5% of the active equipment CapEx), powering expenses. For FWA (e.g., WiGig) a pole rental cost (\$20/pole/year) is assumed. No depreciation is applied.
2. Household density (expressed as number of households per km²) and service take rate are the two main TCO analysis parameters. If sensitivity analysis is shown for other parameters, a density of 3000 households/km² (HH/km²) and/or a take rate of 40% are assumed.
3. Baseline values for other key parameters are:
 - Distribution fiber underground installation cost of \$40/m.
 - Street length can considerably impact the TCO. An 800m street length is assumed.
 - FTTH drop installation cost is fixed at \$300, for both SDU (sub-urban) and multi dwelling unit (MDU) (dense urban) scenario¹.
 - CPE/optical network unit (ONU) installation cost (year 1) is \$150 for both FWA WiGig and FTTH. For the FWA solution 50% of the CPE will be self-installed.
 - DPU/DAA and WiGig AP installation costs are respectively \$400 and \$1,000. It is assumed that WiGig utilizes wireless relay. For DPU node we assume 8 ports. The rather high cost for the AP installation includes the powering supply implementation cost. This can be realized for example by leveraging the AC main power available at a utility pole requiring technical support from the utility/pole provider. The baseline scenario assumes availability of a local power supply.

1.3. Analysis Results

1.3.1. Baseline

The TCO results for the baseline scenario are presented in Figure 7. The top table shows the “lowest cost” solution for each density/take rate combination, the bottom table the associated TCO (\$) per connected household, and the middle table the relative TCO difference with the “second lowest cost” solution. The top table clearly shows the zones of advantage for WiGig FWA, DAA and FTTH. For example, WiGig is the lowest cost solution for HH density of 500 HH/km², with a TCO ranging from \$2.2K to \$6.7K per household connected (HHC). However, as the HH density increases, WiGig only remains the least cost solution for low take rate areas, while other solutions such as DAA and especially FTTH become most economical for higher take rate and higher density areas.

¹ Vertical riser fiber implementation may impact costs.

lowest TCO solution		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	WPON	WPON	WPON	WPON	WPON
	20%	WPON	WPON	WPON	DAA	DAA
	30%	WPON	WPON	DAA	DAA	DAA
	40%	WPON	DAA	DAA	FTTH	FTTH
	50%	WPON	DAA	FTTH	FTTH	FTTH

Relative TCO difference		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	67%	50%	42%	30%	23%
	20%	50%	30%	7%	4%	16%
	30%	44%	6%	16%	21%	26%
	40%	32%	7%	22%	17%	14%
	50%	21%	18%	16%	12%	5%

10 year TCO/HHC (\$)		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	6,763	3,755	2,827	2,452	2,218
	20%	3,979	2,368	2,029	1,762	1,524
	30%	2,930	2,033	1,551	1,283	1,124
	40%	2,483	1,702	1,244	1,337	1,231
	50%	2,215	1,427	1,351	1,207	1,122

Figure 7 - Metropolitan zones of advantage for gigabit access

1.3.2. Sensitivities

Figure 8 shows a TCO breakdown for each solution for the considered take rate and household density. For illustrative purposes, we have chosen a housing density of 3,000 households/km2, which borders suburban and urban housing densities; similar charts can be presented for all the housing densities shown in Figure 7.

Copper/coax solutions have a balanced contribution of network equipment cost, civil works/fiber install cost, CPE cost and OpEx to the 10-year TCO. However, the FTTH and FWA WiGig solutions show a different cost break down: FTTH civil works (e.g., fiber install, including the fiber drop) accounts for 75% of TCO (with very low OpEx and equipment cost contributions). FWA is the reverse, fiber install only contributes 15% to the TCO, with very high equipment cost contribution - driven by the CPE cost included twice due to equipment replacement cycle in year 10 and relative high OpEx.

The lower total fiber install cost for the WiGig FWA solution is because the APs (5 required for the modelled scenario) can, to a certain extent, rely on the wireless relay/backhaul capability. This backhaul link has limited capacity/reach and for high load (high take rate and/or higher sustained bit rate per connected HH) as well as high-density scenarios, fiber must be pulled deeper into the street, as the wireless backhaul link can no longer carry the traffic load of all connected HHs in the street as in a low load scenario.




Figure 8 - TCO breakdown for each solution

Full advantage of the wireless relay capability for FWA is only possible when no power cable must be installed for the APs, the assumption for the baseline scenario.

Figure 9 shows the ZoA view for two other scenarios. The differences with the baseline scenario are:

- Scenario 1 (left): FWA WiGig where dedicated powering cable must be installed (\$40/m install cost). We maintain the high AP installation cost accounting for a remote DC supply implementation requiring DC up convertor in the CO and a down convertor per AP.
- Scenario 2 (right): same as scenario 1, but no trenching is needed for fiber/power cable install (use existing ducts or aerial cabling) resulting in a much lower install cost for fiber/power cable of \$8/meter, applied to all solutions.

For scenario 1, we have basically the same ZoA result for FWA WiGig, DAA and FTTH as considered earlier. For Scenario 2 (duct or aerial cabling deployment), FTTH becomes the preferred solution across all density and take rate combinations, with a TCO of less than \$800/HHC for take rates beyond 20% and densities of 4000 HHP/km2 or more.

lowest TCO solution		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	WPON	WPON	WPON	WPON	DAA
	20%	WPON	WPON	DAA	DAA	DAA
	30%	WPON	DAA	DAA	DAA	DAA
	40%	WPON	DAA	DAA	FTTH	FTTH
	50%	DAA	DAA	FTTH	FTTH	FTTH

lowest TCO solution		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	FTTH	FTTH	FTTH	FTTH	FTTH
	20%	FTTH	FTTH	FTTH	FTTH	FTTH
	30%	FTTH	FTTH	FTTH	FTTH	FTTH
	40%	FTTH	FTTH	FTTH	FTTH	FTTH
	50%	FTTH	FTTH	FTTH	FTTH	FTTH

Relative TCO difference		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	15%	10%	5%	1%	2%
	20%	11%	1%	8%	13%	16%
	30%	6%	8%	16%	21%	26%
	40%	2%	13%	22%	20%	14%
	50%	2%	18%	16%	12%	5%

Relative TCO difference		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	248%	185%	157%	137%	115%
	20%	190%	128%	81%	57%	42%
	30%	162%	79%	43%	25%	14%
	40%	133%	51%	21%	59%	47%
	50%	105%	33%	43%	40%	29%

10 year TCO/HHC (\$)		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	9,848	5,132	3,797	3,168	2,724
	20%	5,408	3,048	2,164	1,762	1,524
	30%	3,981	2,162	1,551	1,283	1,124
	40%	3,206	1,702	1,244	1,337	1,231
	50%	2,688	1,427	1,351	1,207	1,122

10 year TCO/HHC (\$)		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	1,755	1,187	1,022	949	906
	20%	1,221	937	854	818	796
	30%	1,043	853	798	774	760
	40%	954	811	770	752	741
	50%	900	786	753	739	730

Figure 9 - Zones of Advantage, with trenching (left) and without trenching (right)

Finally, it must be noted that FTTdp/G.FAST did not show up as lowest cost solution in previous ZoA views. However, when considering urban areas where no good Line of Sight conditions occur (too much foliage, ruling out WPON as potential solution), FTTdp has also its sweet spot for a CSP (i.e., we leave out the DAA solution as well for the ZoA result), but it does require in general a take rate of at least 40% to beat FTTH for underground fiber deployment scenario. DPU nodes with higher port count (e.g. 16 or 32, with DPU installed in the MDU buildings' basements) allowing to share the DPU node cost across more connected HHs, can also lower FTTdp TCO for high density/high take rate area. In the ZoA view of Figure 9, no HW equipment replacement cycle was considered in the 10 year period.

lowest TCO solution		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	FTTH	FTTH	FTTH	FTTH	FTTH
	20%	FTTH	FTTH	FTTH	FTTH	FTTH
	30%	FTTH	FTTdp	FTTH	FTTH	FTTH
	40%	FTTH	FTTdp	FTTdp	FTTdp	FTTdp
	50%	FTTdp	FTTdp	FTTdp	FTTdp	FTTdp

Relative TCO difference		HH density (/km2)				
		500	2000	4000	6000	8000
Take rate	10%	20%	15%	28%	35%	42%
	20%	12%	5%	13%	17%	20%
	30%	7%	1%	5%	6%	8%
	40%	3%	7%	3%	2%	2%
	50%	0%	12%	10%	10%	10%

Figure 10 - Zones of Advantage

[key assumptions: \$40/m fiber/power cable installation cost (underground deployment), no equipment replacement cycle, no LoS (WPON not considered), Telco view (DAA not considered)]

2. Rural deployments

2.1. Overview

In rural areas characterized by low population density, rolling out gigabit access services will, in general, be cost prohibitive [20, 21]. Competition is often lacking and there is no incentive for a service provider to be best-in-class. However, providing high service speeds can result in higher average revenue per user (ARPU) for the service provider as it enables triple play packages. Regulation and digital agendas may impose a certain requirement, for example, a 50 Mbps peak service rate in exchange for government funding. In this rural deployment analysis, we consider 100 Mbps as downstream peak rate and 25 Mbps as maximum sustained speed per connected household.

To compare the cost of different solutions enabling up to 100 Mbps service speeds in rural area, we built a model where the present mode of operation (PMO) constitutes a low-speed service deployment over a completely passive copper outside plant. The service provider offers 10 Mbps internet access service based on legacy ADSL2 technology from one or multiple CO locations via a twisted copper pair to each subscriber's premise. This service provider now wants to provide higher (peak) service rates up to 100 Mbps. The solutions enabling these higher Future Mode of Operation (FMO) speeds included in the model are listed below and shown in Figure 10. Items in boxes with solid outline are active equipment included in the model; boxes with dotted outline imply that new passive equipment (i.e., fiber) costs are also included. Since copper exists in the deployments, no additional cost is modeled.

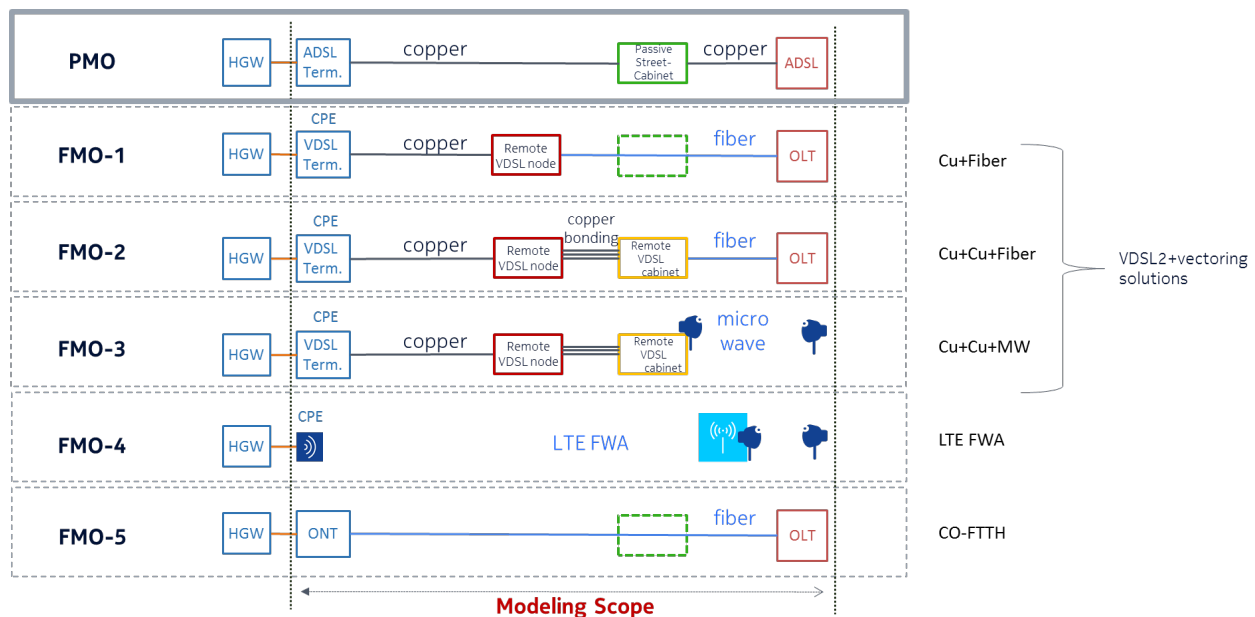


Figure 10: Rural scenario deployment solutions considered

2.1.1. VDSL2 with fiber backhaul (Cu+Fiber)

Higher spectrum (compared to ADSL2) and crosstalk interference mitigation enabled by vectoring technology, enable speeds of 100 Mbps and beyond to be achieved with VDSL2 solutions for loop lengths under 400m. Optimally, active equipment is placed near existing street cabinets containing the copper loop distribution frames. This requires bringing fiber to those locations for backhauling traffic to the CO. The new remotely positioned equipment must also be power-fed which could be based on remote DC powering (from the CO) or local AC powering.

2.1.2. VDSL2 with copper + fiber backhaul (Cu+Cu+Fiber)

To reduce the amount of fiber installation labor, existing copper loops can be leveraged for backhauling the traffic from VDSL2 nodes at remote locations, to a VDSL2 node positioned near a street cabinet location closer to the CO. To provide sufficient backhaul capacity, 8 pairs of copper loop are bonded and VDSL2 with vectoring is applied on the bonded pairs. Depending on the exact physical topology of the copper outside plant and the capabilities and sizes (e.g., port counts) of the VDSL2 (remote node) equipment, a star connectivity or a daisy chained copper bonding backhaul connectivity overlay for connecting the different remote nodes with one another can be used.

2.1.3. VDSL2 with copper + microwave backhaul (Cu+Cu+MW)

Another variant considered for the VDSL2 solution based on copper bonding backhaul where the fiber BH between the node that is aggregating all copper-backhailed traffic and the CO is replaced by a microwave (MW) link. This may be a suitable solution when fiber deployment is very costly (e.g. requiring high cost trenching). A MW link may typically enable up to 1.5 Gbps capacity across several kms of distance.

2.1.4. Fixed Wireless Access (FWA)

Both 4G and 5G FWA solutions were considered, where the CPE antenna is placed outdoor:

- LTE (1.8 Ghz, 2 x 20 MHz spectrum, 4x2 MIMO with 16dBi gain CPE antenna, available today
- 5G (3.5/4 GHz band, 100 MHz spectrum, 8x4x2 MU-MIMO, with 10dBi gain CPE antenna placed at rooftop (6m height). This solution is expected to become available in 2019.

The macro base stations are equipped with a 3-sector radio and associated baseband processing, with an initial ISD that depends on the household density of the area. If capacity is not sufficient, then new sites are added (site densification). The cost of adding these new sites is high since it entails new tower installations, along with radio and baseband processing equipment/installation, but also backhaul must be accommodated, which is assumed to be a microwave link. As this investment co-purposes both mobile and FWA services, a cost sharing split of 20%-80% and equivalent capacity sharing split where the larger portion is allocated to FWA are assumed. It is expected that 5G usage for mobile applications will take some time to reach significant penetration in rural areas.

2.1.5. Fiber to the home from the central office (CO-FTTH)

FTTH, based on GPON technology, with the OLT located in a central location is assumed. FTTH may make sense as in some areas where government funding is available. Such funding for rural deployment can make the FTTH business case more attractive.

2.2. Modeling assumptions

For the rural model, CapEx investments required for each solution are compared. This model is not based on a street simulation as applied for the urban/metropolitan environment, but rather on statistical derivations considering that the rural area is vast and any dependency on street patterns or specific clustering of houses may be questionable.

The main assumptions include:

- Street cabinet density(PMO), initial (PMO) and minimum (FMO) FWA base station inter-site distance (ISD) and average FTTH drop cost are all dependent on the household density of the considered area.
- Copper loop length distribution is Rayleigh distribution model based
- Non-VDSL2 solutions require a truck roll when subscribers opt in for the enhanced service: a customer premise visit for installing the FWA CPE, and a double truck roll for FTTH: 1) one for drop implementation, and 2) one for ONU installation and service activation.
- Equipment and implementation costs are based on United States (US) market broadband deployment benchmarks, with a cost estimate for the 5G FWA solution (since it is not yet commercially available).

Table 3 - Rural deployment assumptions that vary with HH density

HH density (/km2)	10	50	100	500
Street cabinet density (/km2)	0.5	1	2	2
Initial ISD (km)	10	5	4	2
Minimum ISD (km)	2	1	1	1

HH density (/km ²)	10	50	100	500
FTTH drop cost (\$)	800	400	250	150

The CapEx of the different solutions is compared on the following key parameters and associated values: household density (10, 50, 100 and 500 HH/km²), service take rate (30, 40, 50 and 60%), service peak bit rate (10, 25, 50 and 100 Mbps), sustained bit rate per connected household (3, 6, 12 and 25 Mbps, capped by the service peak bit rate), and fiber installation cost per meter for the distribution/feeder sections (20, 30, 40 or 50 \$/m). Unless mentioned otherwise, the copper backhauled VDSL2 solution is deployed in a daisy chain mode.

Note that the fiber may be installed in existing ducts, on poles or it may require civil works for underground deployment. The fiber installation cost/m can reflect any or a mix of these installations.

The baseline values for the key deployment parameters were chosen to be 50 HH/km², 50% service take rate, 50 Mbps peak service bit rate, 6 Mbps sustained bit rate per HHC and a fiber install cost at \$30/m.

2.3. Analysis results

Results of the rural modeling exercise are based on the main parameters as discussed above. Note that in Figure 11 through Figure 14, the FWA solution is based on LTE 1.8 GHz (40 MHz spectrum).

2.3.1. Impact of household density




Figure 11 - Rural deployment CapEx per HHC as function of HH density

All wireline solutions have lower cost for higher density. For baseline values the VDSL solution with copper + fiber backhauling (Cu+Cu+Fiber) is always lowest cost for mid-to high HH densities (> 50 HH/km²). For mid densities, the other VDSL solutions follow closely. For very low density (10 HH/km²), VDSL2 nodes have very low filling and combined with the longer fiber distances from the CO to the remote nodes, street cabinet density decreases with lower HH density, resulting in very high cost per connected subscriber. In such low-density areas, FWA becomes the lowest cost solution. For the very

higher (500 HHs/km²) density, FTTH becomes lower cost than LTE FWA. This is because of the very high FWA base station site densification that is required to meet the given service requirements in a high household density area. The charts reveal that cost points per connected HH for 50 HH/km², range between \$2000 and \$4500 when relying on the existing copper outside plant for the 50 Mbps peak bit rate service offering.

2.3.2. Impact of service peak bit rate




Figure 12 - Rural deployment CapEx per HHC as function of peak rate

Whereas the CO-FTTH and FWA LTE solutions' TCO have no dependency on the service peak bit rate (note that this peak bit rate is not necessarily guaranteed), the VDSL2 solutions' TCO show a high dependency on the service peak rate, starting from 25 Mbps. For 100 Mbps peak rate, the cost/HHC quadruples for fiber backhauled VDSL2, and doubles for the VDSL2 solutions relying on the copper bonding, which is explained by the requirement for higher VDSL2 node density. Note that for 100 Mbps peak service bit rate, the FWA LTE solution TCO almost approaches the VDSL2 solution with copper and fiber BH for the considered 50 HH/km² density.

Figure 13 shows the solutions' CapEx break down of active network equipment versus passive OSP investments.

LTE-based FWA is clearly not a sustainable solution for given service scenario and assumptions, where equipment cost (radio, baseband processing and MW Backhaul) and civil works costs for the new sites equally contribute to the CapEx as driven by cell densification.

It clearly shows that for FTTH, the fact that no active equipment is needed in the OSP cannot be taken advantage of because of the high civil works cost contribution. For Cu+Cu+MW, no civil works are required, but with CapEx based 100% on equipment/installation cost, a significant OpEx contributor can be expected.

The difference in CapEx between daisy chain and star (ST) topology for the VDSL2 copper BH solution is also shown. Star topology is higher cost than daisy chain - mainly driven by the higher equipment cost - but still its total cost is lower than the alternatives. For the given key parameter values, 7 active nodes

must be deployed - 6 are copper backhauled to the node positioned near an existing street cabinet- to ensure the high peak rate offer (100 Mbps), but each node on average services only 6 subscribers. This means for a remote node with 48 ports, 34 ports are still left unused in star topology.




Figure 13 - Rural deployment CapEx breakdown per HHC

2.3.3. Impact of fiber install cost




Figure 14 - Rural deployment impact of fiber install cost

Figure 14 shows the impact of fiber install cost. FTTH has the strongest dependency on the fiber install cost, followed by VDSL2 with fiber backhaul and the VDSL2 solution that depends both on copper and fiber backhauling.

2.3.4. 4G FWA versus 5G FWA

As shown in previous section, the LTE-based FWA in most cases has relative high cost, and we now compare the wireline solutions with the 5G FWA solution. Similar to the metropolitan section, the matrices shown in Figure 15 present the Zone-of-Advantage (ZoA) for different HH densities and different sustained bit rates as follows: (top to bottom) 1) the lowest cost solution; 2) its cost/HHC; 3) the CapEx difference with the 2nd lowest cost solution; and 4) the 2nd lowest cost solution. On the left side of the ZoA matrix the FWA solution is based on LTE (1.8 GHz, 20 MHz FDD), on the right the FWA is based on 5G NR technology (3.7-4.2 GHz spectrum and 100 MHz TDD available).

Lowest CAPEX solution		HH density (/km2)			
		10	50	100	500
sust bit rate/HHC (Mbps)	3	FWA	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Cu+Fiber
	6	FWA	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Cu+Fiber
	12	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Cu+Fiber
	25	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Fiber

Lowest CAPEX solution		HH density (/km2)			
		10	50	100	500
sust bit rate/HHC (Mbps)	3	FWA	FWA	FWA	Cu+Cu+Fiber
	6	FWA	FWA	FWA	Cu+Cu+Fiber
	12	FWA	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Cu+Fiber
	25	FWA	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Fiber

CAPEX/HHC (\$)		HH density (/km2)			
		10	50	100	500
sust bit rate/HHC	3	2307	1957	1308	434
	6	4167	1961	1310	458
	12	7494	1971	1315	500
	25	7540	1986	1381	538

CAPEX/HHC (\$)		HH density (/km2)			
		10	50	100	500
sust bit rate/HHC	3	954	810	849	434
	6	1444	1235	1235	458
	12	2319	1971	1315	500
	25	4148	1986	1381	538

Relative CAPEX diff. Lowest- 2nd lowest		HH density (/km2)			
		10	50	100	500
sust bit rate/HHC (Mbps)	3	224%	21%	36%	21%
	6	79%	81%	36%	15%
	12	33%	80%	36%	6%
	25	32%	78%	30%	524%

Relative CAPEX diff. Lowest- 2nd lowest		HH density (/km2)			
		10	50	100	500
sust bit rate/HHC (Mbps)	3	682%	142%	54%	21%
	6	418%	59%	6%	15%
	12	223%	3%	36%	6%
	25	82%	78%	30%	524%

2nd lowest CAPEX solution		HH density (/km2)			
		10	50	100	500
sust bit rate/HHC (Mbps)	3	Cu+Cu+Fiber	FWA	Cu+Fiber	Cu+Fiber
	6	Cu+Cu+Fiber	Cu+Cu+MW	Cu+Fiber	Cu+Fiber
	12	Cu+Cu+MW	Cu+Cu+MW	Cu+Fiber	Cu+Fiber
	25	Cu+Cu+MW	Cu+Cu+MW	Cu+Fiber	CO-FTTH

2nd lowest CAPEX solution		HH density (/km2)			
		10	50	100	500
sust bit rate/HHC (Mbps)	3	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Fiber
	6	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Cu+Fiber	Cu+Fiber
	12	Cu+Cu+Fiber	FWA	Cu+Fiber	Cu+Fiber
	25	Cu+Cu+Fiber	Cu+Cu+MW	Cu+Fiber	CO-FTTH

Figure 15 - Rural access ZoA with FWA LTE (left) and FWA 5G (right)

LTE-based FWA is the lowest cost solution in the (left) ZoA only for very low density and very low sustained rate /HHC at a cost level above \$2300/HHC (for 50% take rate). However, in the 5G era, FWA has lowest TCO for HH densities up to 100HH/km2 and sustained bit rate levels up to 25 Mbps, at cost points between \$1,000 and \$1,500/HHC. The difference in CapEx for the 5G FWA solution relative to the 2nd lowest cost solution (VDSL2 with copper bonding) is a minimum 60% for densities up to 50 HH/km2.

For higher densities and/or higher sustained rates, the copper + fiber BH based VDSL2 (Cu+Cu+fiber) solutions remain the lowest cost for the modeled 50 Mbps (or higher) peak bit rate service.

For the considered baseline values for the main parameters- 5G FWA solutions can indeed become a cost-effective alternative for (rural) broadband access deployments enabling, e.g., high definition (HD) video service packaging, once these solutions become available for mass deployment.

Conclusions

MSOs should review these conclusions from both an opportunity as well as a threat perspective. As the demand for high-bandwidth services increases, advancements in HFC architecture continue to give them a performance advantage over copper. If competition comes in the form of fiber, a deep fiber deployment supporting HFC keeps MSOs ready for fiber drop deployment quickly, whenever needed.

Competition is also expected in the form of FWA, primarily from wireless operators who will attempt to increase ARPU with quad-play. These deployments will likely be in licensed spectrum [17,18], and the wireless operators will be able to leverage the same infrastructure for both mobile and fixed services. But FWA, especially in unlicensed band, e.g., WPON, also provides opportunities for MSOs to expand into new territories without large investment in drop costs.

As we saw in this paper, the technologies and their ZoA are very different between the high-bandwidth and low-bandwidth services. For high-bandwidth services, FTTH comes in as advantageous whenever aerial deployments are feasible. While this is generally practical in suburban neighborhoods, local regulations may prohibit aerial deployment in urban and dense urban environments. FWA with 5G in centimeter-wave as well as WPON with millimeter-wave could become viable options, especially in good LoS conditions. Copper and HFC could also be advantageous under certain housing densities and take rates.

For low-bandwidth service, different VDSL options are often advantageous. In rural areas with low housing densities where HFC is not an incumbent technology, wireless options with LTE-Adv as well as 5G in mid-band can be viable options [20, 21] depending upon service targets and housing densities.

In the next few years, multi-Gbps technologies are expected to become more common, starting likely with enterprise access. Advancements in both wired and wireless technologies will render these deployments cost-competitive. At the same time, demand for high bandwidth services will continue to grow and the ZoA observed in the paper will continue to evolve. As cost points mature and technological advancements continue, the analyses will have to be revisited.

Abbreviations

2G	2 nd generation wireless
3G	3 rd generation wireless
4G	4 th generation wireless
5G	5 th generation wireless
AP	access point
ARPU	average revenue per user
BH	backhaul
CapEx	capital expenditure
CCAP	converged cable access platform
cm	centimeter
CDF	cumulative distribution function
CO	central office
CPE	customer premise equipment
CSP	communications service provider

DAA	distributed access architecture
DC	direct current
DOCSIS	data over cable interface specification
DPU	distribution point units
DSL	digital subscriber line
eMBB	enhanced massive broadband
FTTB	fiber to the building
FTTdp	fiber to the distribution point
FTTH	fiber to the home
FTT-SMB	fiber to the small/medium business
FMO	future mode of operation
FOV (HXV)	Field of vision (horizontal x vertical)
FWA	fixed wireless access
Gbps	gigabit per second
GHz	gigahertz
GPON	gigabit passive optical network
HDMI	high-definition multimedia interface
HFC	hybrid fiber-coax
HH	household
HHC	household connected
HHP	household passed
ISD	inter-site distance
LoS	line of sight
LTE-Adv	Long Term Evolution Advanced
MDU	multi dwelling unit
MHz	megahertz
mm	millimeter
mMTC	massive machine type communications
mMIMO	massive multiple input, multiple output
MSO	multiple system operator
MW	microwave
OFDM	orthogonal frequency division multiplexing
OLT	optical line termination
ONU	optical network unit
OpEx	operating expenses
OSP	outside plant
P2MP	point-to-multipoint
PMO	present mode of operation
PoP	point-of-presence
RF	radio frequency
RMD	remote MAC/PHY device
RPD	remote PHY device
RxD	remote x device
SDU	single dwelling unit
SG	service group
TCO	total cost of ownership
TDM	time division multiplexing
TWDM	time-wavelength division multiplexing

URLLC	ultra-reliable ultra-low latency
US	United States
VR	virtual reality
WiGiG	wireless gigabit
WPON	wireless PON
ZoA	zones of advantage

Bibliography & References

1. The Carmel Group: Broadband Wireless Access Providers Prepare to Soar with Fixed Wireless THE BWA INDUSTRY REPORT: 2017
2. J. Wellen, P. Kapauan and A. Mukhopadhyay: Sustained Throughput Requirements for Future Residential Broadband Service, 2017 Fall Technical Forum, SCTE-ISBE, NCTA, Cable-Labs.
3. G.9700: Fast access to subscriber terminals (G.fast) - Power spectral density specification". ITU-T. 2014-12-19. Retrieved 2015-02-03.
4. G.9701: Fast access to subscriber terminals (G.fast) - Physical layer specification. ITU-T. 2014-12-18. Retrieved 2015-02-03.
5. G.fast broadband standard approved and on the market. ITU-T. 2014-12-05. Retrieved 2015-02-03.
6. Oksman et al., "The ITU-T's new G.fast standard brings DSL into the Gigabit era," IEEE Commun. Mag., vol. 54, no. 3, pp. 118–126, Mar. 2016.
7. Cablelabs, Data-Over-Cable Service Interface Specifications - CM-SP-PHYv3.1 specifications
8. Cablelabs, Data-Over-Cable Service Interface Specifications - Modular Headend Architecture v2 technical report - CM-TR-MHA-V2
9. Cablelabs, Data-Over-Cable Service Interface Specifications, DCA Distributed CCAP Architectures Overview Technical Report - CM-TR-DCA-V01
10. Gigabit-capable passive optical networks (GPON): General characteristics, ITU-T G.984
11. 10-Gigabit-capable passive optical networks (XG-PON) systems: ITU-T G.987
12. 3GPP TS 36.300 Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description, Stage 2 – Release 15
13. 3GPP TS 23.501 System Architecture for the 5G System Release 15
14. 3GPP TS 38.300 NR: Overall description; Stage-2 Release 15
15. IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi
16. 3GPP TS 22.261, "Service requirements for next generation new services and markets", Release 15
17. "Mobilizing 5G NR Millimeter Wave: Network Coverage Simulation Studies for Global Cities", Qualcomm Technologies, Inc., Oct. 2017, <https://www.qualcomm.com/media/documents/files/white-paper-5g-nr-millimeter-wave-network-coverage-simulation.pdf>
18. "5G White Paper", NGMN Alliance, Feb. 2015, https://www.ngmn.org/fileadmin/ngmn/content/downloads/Technical/2015/NGMN_5G_White_Paper_V1_0.pdf
19. 5G Americas, Wireless Technology Evolution Towards 5G: 3GPP release 13 to release 15 and beyond, February 2017
20. Nokia White paper, Broadband Transformation for 21st Century Digital Rural Society, 2016.
21. Energy efficient 5G Deployment in Rural Areas, A. Karlsson, O. Al-Saadeh, A. Gusarov, R V R Challa, S. Tombazy, and K W Sung, 2016 IEEE 12th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)

The New Home as a Hotspot: Wi-Fi Meet CBRS LTE and Meet Your Long Range Brother LoRa

A Technical Paper prepared for SCTE•ISBE by

J.R. Flesch

Director, Advanced Technology
ARRIS International plc
3871 Lakewood Drive
Suwanee, Georgia 30024
+1 678 473 8340
jr.flesch@arris.com

Charles Cheevers

CTO/CPE
ARRIS International plc
3871 Lakewood Drive
Suwanee, Georgia 30024
+1 678 473 8507
Charles.Cheevers@arris.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Content.....	4
1. Snapshotting the Wireless Home.....	4
2. 3.5 GHz CBRS	4
2.1. CBRS/LTE Home Inside/Out Signal AP Reach Analysis.....	7
2.1.1. Bungalow Data and Voice Services.....	10
2.1.2. Average Home Data and Voice Services	13
2.1.3. Large Home (Mansion) Data and Voice Services.....	16
2.2. Tabular Summary of Inside/Out 3.5 GHz CBRS AP Reach	19
2.3. 3.5 GHz CBRS Neighborhood Roaming Via HaaT (Home as a Tower).....	21
2.4. Roaming User Considerations	26
2.5. 1W HaaT Simulations	27
2.6. CBRS/LTE and IoT	29
2.7. Timing Considerations Across LTE and Cable Domains.....	30
3. LoRa.....	30
3.1. Base Stations and CEDs.....	32
3.2. Link Specifics	33
3.3. Link Budget and Service Throw (Range)	34
3.4. Geolocation of Clients as Calculated Benefit.....	37
3.5. The Issues of Scale.....	39
4. LoRa Compared Vs LTE Narrow Band IoT Options	41
Conclusions.....	42
Abbreviations	43
Bibliography & References.....	44

List of Figures

Title	Page Number
Figure 1 - 3.5 GHz CBRS Shipborne Radar Coastal Exclusion Zones	5
Figure 2 - CBRS Priority Tier Membership Distribution	6
Figure 3 - CBSD Category A and B Power Signature Limits.....	7
Figure 4 - 3.5 GHz Propagation Study.....	8
Figure 5 - Schema of Inside/Out Propagation Studies for 3.5 GHz CBRS.....	9
Figure 6 - Defining the Roaming (Home Exterior) Proposition for CBRS/LTE	9
Figure 7 - Sample Bungalows of Various External Construction	10
Figure 8 - Inside/Out Bungalow Service Reach @ 25/5 Mbps	11
Figure 9 - Inside/Out Bungalow Service Reach @ 10/2 Mbps	11
Figure 10 - Inside/Out Bungalow Service Reach for Voice Call	12
Figure 11 - Sample Average Homes of Various External Construction.....	13
Figure 12 - Inside/Out Average Home Service Reach @ 25/5 Mbps	14
Figure 13 - Inside/Out Average Home Service Reach @ 10/2 Mbps	14

Figure 14 - Inside/Out Average Home Service Reach for Voice Call	15
Figure 15 - Sample Mansions of Various External Construction	16
Figure 16 - Inside/Out Mansion Service Reach @ 25/5 Mbps.....	17
Figure 17 - Inside/Out Mansion Service Reach @ 10/2 Mbps.....	17
Figure 18 - Inside/Out Mansion Service Reach for Voice Call.....	18
Figure 19 - 25/5 Mbps Service Radii Performance.....	19
Figure 20 - 10/2 Mbps Service Radii Performance.....	20
Figure 21 - Voice Call Service Radii Performance.....	21
Figure 22 - Inside/Out Coverage for 25/5 Service Assuming Wood Sided Homes (1 x 6 Homes)	22
Figure 23 - Inside/Out Coverage for 25/5 Service Assuming Brick Sided Homes (1 x 3 Homes)	23
Figure 24 - Inside/Out Coverage for 25/5 Service Assuming Stone Homes (Yellow) and the Need to Augment with 4W Strand Mount APs (Red) (1 x 1 Home + 6 x 4W POPs)	24
Figure 25 - 1W HaaT Coverage for 25/5 Service (1 x 10-16 Homes)	25
Figure 26 - 4W HaaT Coverage for 25/5 Service (1 x 20-40 Homes)	26
Figure 27 - 1 W HaaT Simulation Vs Calculation; ~ 2:1 Service Radius Improvement in Sim.....	27
Figure 28 - 1 W HaaT Simulation, Bitrate Variance in Coverage Map	28
Figure 29 - Service Mounting Potential for Outdoor 4W CBRS Mesh	29
Figure 30 - Schematic of LoRa Edge Network.....	31
Figure 31 - LoRa US ISM 900 MHz Band Occupation	31
Figure 32 - LoRa US Spreading Factor Implications to Noise Margin and Bitrate	32
Figure 33 - LoRa Tradeoff of Sustained Bitrate for Better Range.....	32
Figure 34 - LoRa Client Mix: Sensors and Actuators.....	33
Figure 35 - LoRa US ISM 900 MHz Link Parameters.....	34
Figure 36 - Measured Sensitivity of Commercial LoRa Product	35
Figure 37 - Sample LoRa Suburban Mobile Inside/Out Connectivity @ 1300' Radius	36
Figure 38 - Sample LoRa Rural Mobile Inside/Out Lossless Connectivity @ 49 Km (!) Radius	37
Figure 39 - LoRa TDoA (Time-Difference of Arrival) and RSSI Geolocation Accuracy.....	38
Figure 40 - LoRa Timestamp Error Contribution to Geolocation Error	39
Figure 41 - LoRa Message Dwell Time Implications to Payload and SF.....	40
Figure 42 - LoRa Congestive Behavior Due To Closed Loop ACK of ALOHA Upstream Messages.....	41
Figure 43 - US LoRa Comparison to LTE as LPWAN	42

Introduction

We have given customers a valuable resource – the home Wi-Fi hotspot – a well understood ‘inside out’ service. Is there an opportunity to use the Home for additional inside out services? We now can potentially add to Wi-Fi with CBRS LTE and LoRA services – leveraging the connection to the DOCSIS or Fiber network to provide in home and outside services. With the inclusion of a home cell containing CBRS and LTE, the service provider can build an inside out network targeting the emerging CBRS capable smart phone and NB-IOT devices.

This paper reviews the home architecture required to add CBRS and LoRA home cells to complement existing Wi-Fi hotspots and the software solutions to manage them. The paper further discusses the potential for adding LoRa to the home, as an inside out LoRa edge network, and how to build a comprehensive NB-IoT solution. The RF decisions around the deployment of this cell will also be discussed – 1 per home or 1 per X homes for more efficient initial coverage.

Content

1. Snapshotting the Wireless Home

Home Wi-Fi may be thought of as a cable bandwidth enfranchisement technology which binds wireless cable-native or ‘Bring your own’ CPE used in the home (and immediately on its periphery) to the wireline cable network for backhaul. Emerging management layers seek to assign available Wi-Fi spectrum to this (potentially dynamically mounted) client palette in a manner which load balances air time and wireless channel usage based on anticipated consumption rates, service priorities, user priorities and monitored spectral availability with the goal of maximizing availability and throughput of these user devices and the services they represent. Home Wi-Fi as yet only unevenly accommodates 2.4 GHz ISM band sharing with non-Wi-Fi, NB-IoT radio traffic (as defined by the Bluetooth/BLE family of devices, Thread 802.15.4, and the various Zigbee flavors) – begging additional remediation gambits like band co-existence (TDM) semaphore schemes and (perhaps, in future) explicit FDM (parsing) of the 2.4 GHz band to assign narrowband data conduits through the Wi-Fi clutter for IoT devices to establish competing inband links (dependent upon their radio/MAC technology and service propositions).

While these areas of investigation are being paced in a demanding, immediate market sense by inflating home Wi-Fi bitrate demands and new IoT vertical businesses with critical latency expectations, additional wireless capability (and service hardening – the addition of battery power to bridge mains loss and a wireless backhaul option for loss of wireline) can be found in out-of-Wi-Fi-bands overlay architectures represented by 3.5 GHz CBRS/LTE and 900 MHz LoRa. With respect to the former, the presence of wideband channels also facilitates the deployment of data or voice services in addition to the payloads associated with IoT – the lone motivation in the case of LoRa.

2. 3.5 GHz CBRS

As regards CBRS, the FCC’s creation of this 3.5 GHz spectrum opportunity in 2016 mandates no particular services to be mounted or MAC to be used (though for unlicensed enthusiasts of fine-granularity scheduling and low-latency connectivity either a private LTE network or a new MAC offering called MultiFire were possible; and while not pursued in the literature, one could have conceivably employed 802.11ac in the band at that time as well.) The band’s location between the Wi-Fi 2.4 GHz and 5 GHz bands also promises at least a derivative understanding of its propagation characteristics (with respect to known Wi-Fi art) – and this, in regard to both in-home and inside-out strand-mount AP reach.

As a bookmark, CBRS was broadly envisaged by the FCC to be exploited as a private LTE/TDD technology consisting of fifteen 10 MHz wide channels contiguously arrayed from 3.55 GHz to 3.70 GHz whose spectral access was to be dynamically managed by an entity called the Spectrum Allocation System (SAS). SAS arbitrates requests for bandwidth from potential users and refers these to an executive policy which determines if the request comes from an Incumbent, Priority License Access (PAL) license holder or a member of the General Authorized Access tier. Incumbents (largely shipborne radar, though some fixed satellite and wireless ISP accounts are represented) are given pre-emptive priority. That is, even if services are running on a channel to which they request access, such services are forced to idle themselves.

The FCC mitigated the impact of incumbent exclusion zone requirements by relaxing the radar keep-out footprint in acknowledgment of CBRS' reduced radiated power impact, as shown below – the early cut is in yellow and the final boundary is in blue.

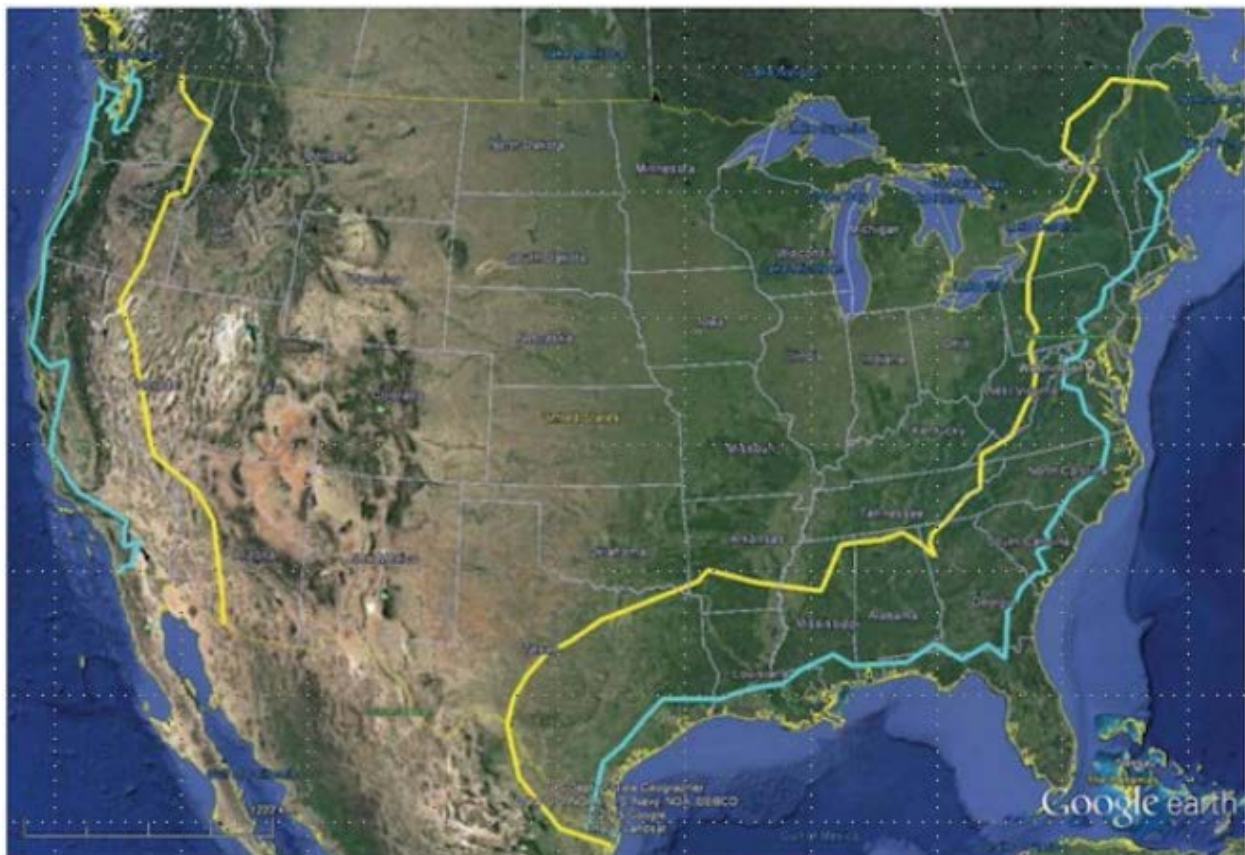


Figure 1 - 3.5 GHz CBRS Shipborne Radar Coastal Exclusion Zones

PAL accounts receive the next use preference and in fact are the highest priority users in most inland use scenarios. They are guaranteed access to 70 MHz of the 150 MHz CBRS spectrum. The final tier (GAA) represents the lowest priority unlicensed users who are guaranteed 80 MHz of spectrum. Note that SAS was originally intended to operate on a highly granular geographic basis (census tract cell sized) with leases of only 1-3 year duration (to promote access by interested small entities). This original notion facilitates highly granular and rapid spectral re-use, in direct proxy to small cell operational dynamics. For example, the City of Philadelphia, with 369 sq km, has 19,000 Census tracts with an average of 1/3 sq km of area. However, recent considerations of the FCC seem to suggest a more “large business entity-

friendly” posture, with service footprints moving to county-sized plots and lease durations running to 7-10 years. In any event, the following figure exemplifies SAS’ priority considerations:

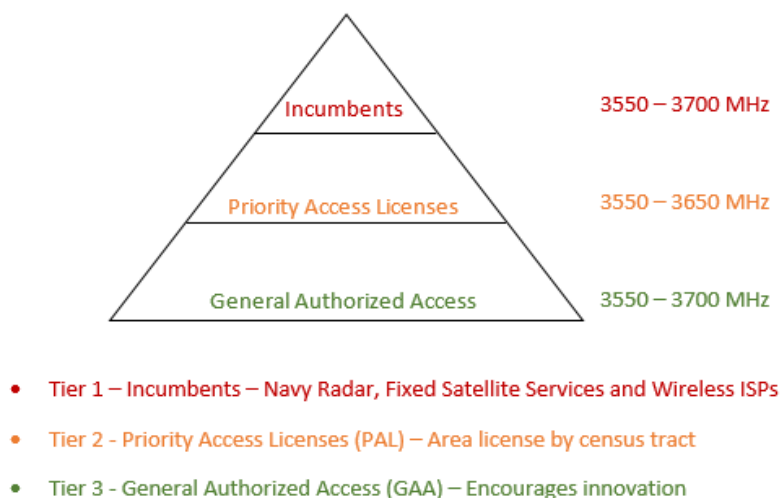


Figure 2 - CBRS Priority Tier Membership Distribution

Fundamentally, the SAS maintains a regionally referenced, curated database of potential users annotated by license type and is also informed by an Environmental Sensing Capability (ESC) device — essentially activity detectors for incumbents, such sensors deployed in proximity to the exclusion zone — and uses these information stores to dynamically arbitrate accesses on small-cell boundaries in the 3.5 GHz CBRS band. To underscore the scalable small-cell nature of CBRS, the FCC created the following radiated and conducted power envelopes for Citizens Broadband Radio Service Devices (CBSDs) which intend on leveraging the PAL and GAA tiers in the band:

CBSD Category	Maximum Conducted Power (dBm/10 MHz)	Maximum EIRP (dBm/10 MHz)	Maximum Conducted PSD (dBm/MHz)	CBSD Installations	Operations in 3550-3650 MHz	Operations in 3650-3700 MHz
Category A	24	30	14	- Indoor - Outdoor max 6m HAAT	Everywhere Outside DoD Protection Zone	Everywhere Outside FSS and DoD Protection Zone
Category B (Non-Rural)	24	40	14	- Outdoor only - Professional Installation	Outside DoD Protection Zone & requires ESC approval	Everywhere Outside FSS Protection Zone and DoD Protection Zone
Category B (Rural)	30	47	20	- Outdoor only - Professional Installation	Outside DoD Protection Zone & requires ESC approval	Everywhere Outside FSS Protection Zone and DoD Protection Zone

Figure 3 - CBSD Category A and B Power Signature Limits

The beauty of leveraging the band with unlicensed LTE means that scale economies for 3.5 GHz radios managed by 3GPP-based LTE narrowband protocols would make the silicon componentry available to minimize both cost and (potentially) battery use of the constrained end devices (CEDs) used in the IoT network. And of course, the LTE small-cell infrastructure in its entirety facilitates carriage of mobile phone services (as replacement for in-home landline dependency or perhaps as proviso of larger footprint, in-neighborhood proximate use scenarios.)

2.1. CBRS/LTE Home Inside/Out Signal AP Reach Analysis

An analysis of available RF link budget for a connection between a mid-home located 3.5 GHz CBRS/LTE AP (modeled as a dual-path device of 3 dBi antenna gain/path with assignable total EIRP of 1W, ½ W or ¼ W and a receiver NF of 6 dB) and a mobile transceiver (modeled as a dual-antenna device with +20 dBm total output EIRP and a receiver NF of 9 dB) was conducted. 20 MHz of channel bandwidth was presumed. A results matrix allowing for three different size home/lot combinations, across three different service grades (Downstream/Upstream Mbps as 25/5, 10/2 and 0.1/0.1 – the latter a voice service presumption but also relevant proxy for NB-IoT signaling) and three types of exterior material construction (wood siding, brick/Hardiplank or stone – all with e-glass windows) was established to test layout sensitivities across several concerns. The end intention of the exercise was to identify potential distribution schemes for exploit of the technology and what complexities might arise.

3.5 GHz propagation characteristics were measured in various open-air environments around Tampa to lump Fresnel and other diffractive effects with bulk loss tangent calculation. Recall that the ~ 10:1 wavelength advantage of the 3.5 GHz band versus mmWave results in a 20 dB lower hit to the link budget and much better nLOS and NLOS propagation. (The free space loss is described by $20 \log(d) + 20 \log(f) + 20 \log(4\pi/c)$ where d is the distance, f is the frequency and c is the speed of light.) A best-fit curve describing generic through-air path losses was extracted from the data:




Figure 4 - 3.5 GHz Propagation Study

Lumped-element material transition attenuations from available tabular data were used to peg 3.5 GHz losses as 3 dB for drywall/floor, 35 dB for e-glass, 3-5 dB for wood siding, 13 dB for brick/Hardiplank and 25 dB for 4-6" stone. Average spring/summer foliage losses for a mixed light forest were set in the range of 10-12 dB. Tabular service radii data is available in the section following the illustrations. The illustration below defines the architectural implications of inside/out home reach using CBRS/LTE:




Figure 5 - Schema of Inside/Out Propagation Studies for 3.5 GHz CBRS

In addition to inside/out possibilities, a backbone of 4W strand mount POPs may be used to extend coverage so that CBRS/LTE mobile devices may be used for the case of neighborhood roaming (this lies within the FCC guidelines of 10W maximum for urban areas and 50W for unpopulated rural tracts):

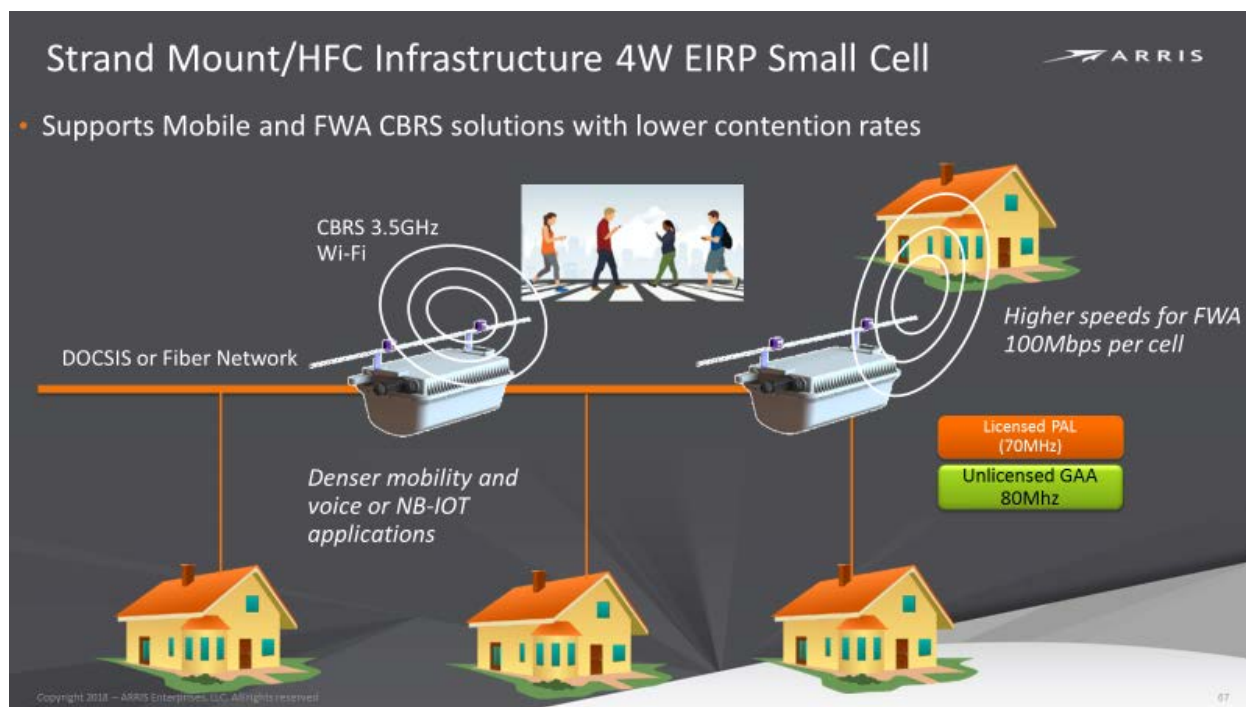


Figure 6 - Defining the Roaming (Home Exterior) Proposition for CBRS/LTE

2.1.1. Bungalow Data and Voice Services



Figure 7 - Sample Bungalows of Various External Construction

The Bungalow portion of the analysis presumed a 1,500 square foot dwelling on a 0.2-acre lot and an internal AP placement which would have met with two internal floor/ceiling transitions and an exterior wall in order to reach outside the home before encountering the foliage costs to propagation. The resultant service radii for the three grades of connection were calculated as follows:

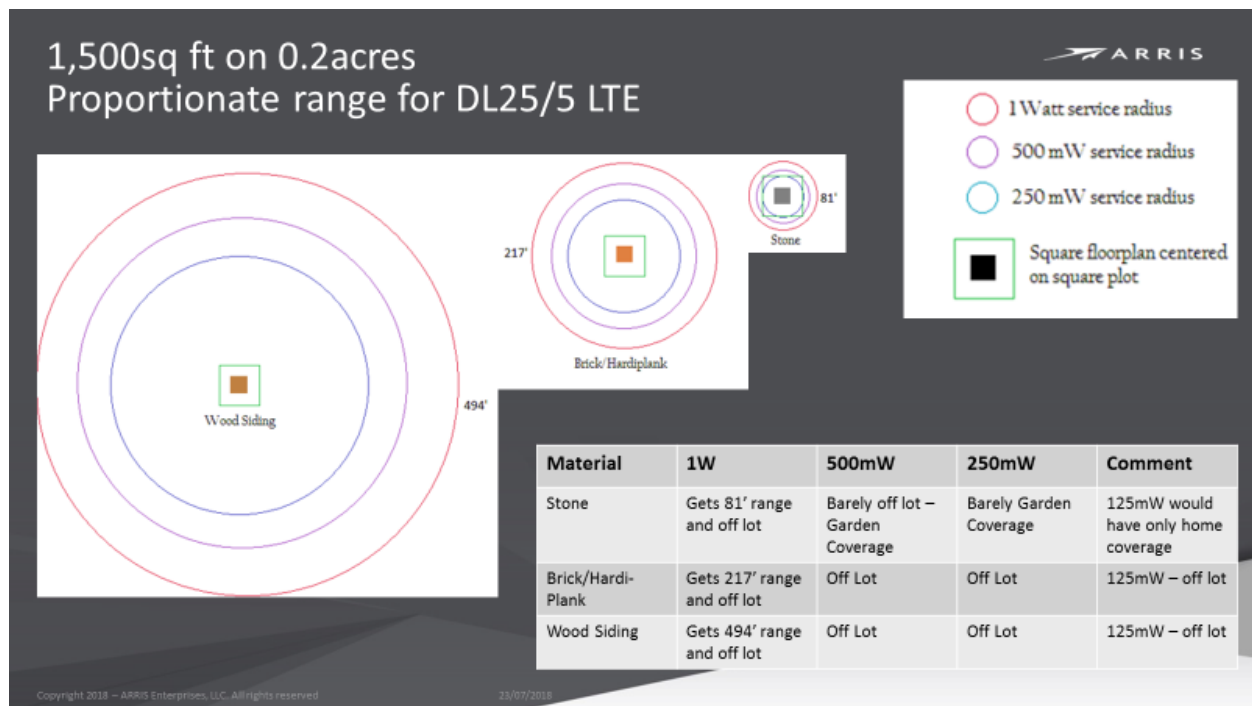


Figure 8 - Inside/Out Bungalow Service Reach @ 25/5 Mbps

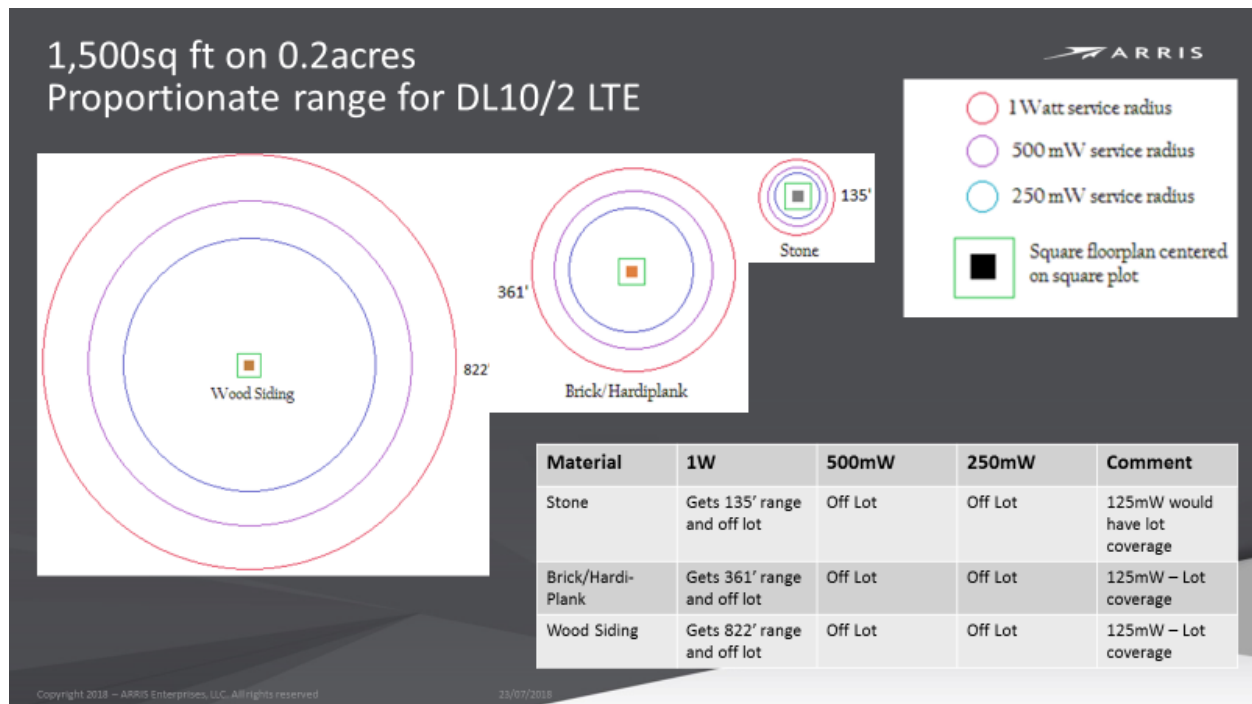


Figure 9 - Inside/Out Bungalow Service Reach @ 10/2 Mbps

1,500sq ft on 0.2acres
Proportionate range for 100kbps voice call

ARRIS

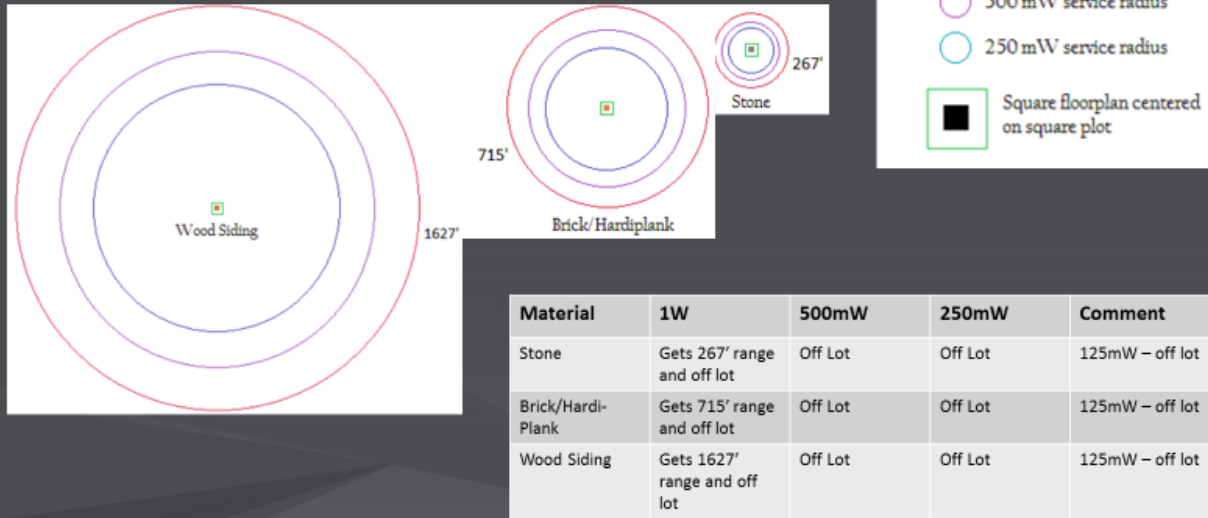


Figure 10 - Inside/Out Bungalow Service Reach for Voice Call

2.1.2. Average Home Data and Voice Services



Figure 11 - Sample Average Homes of Various External Construction

The average-sized home portion of the analysis presumed a 2,500 square foot dwelling on a 0.35-acre lot and an internal AP placement which would have met with three internal floor/ceiling transitions and an exterior wall in order to reach outside the home before encountering the foliage costs to propagation. The resultant service radii for the three grades of connection were calculated as follows:

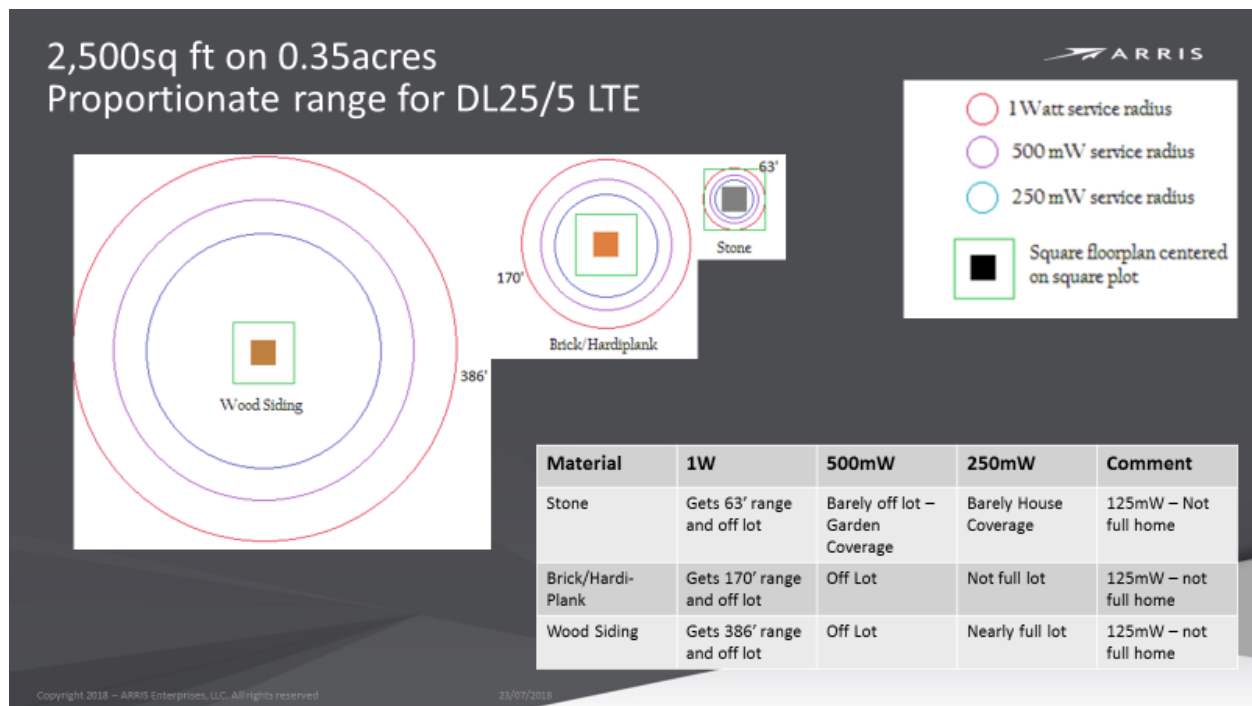


Figure 12 - Inside/Out Average Home Service Reach @ 25/5 Mbps

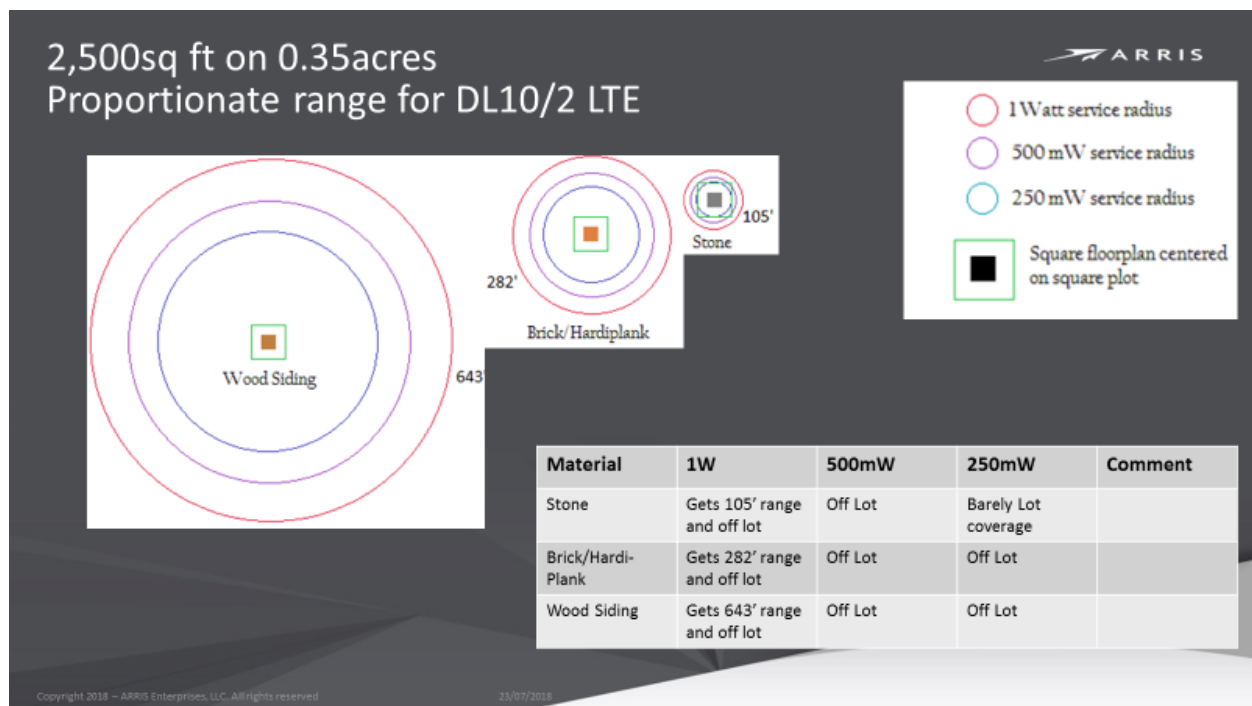


Figure 13 - Inside/Out Average Home Service Reach @ 10/2 Mbps

2,500sq ft on 0.35acres
Proportionate range for 100kbps voice call

ARRIS

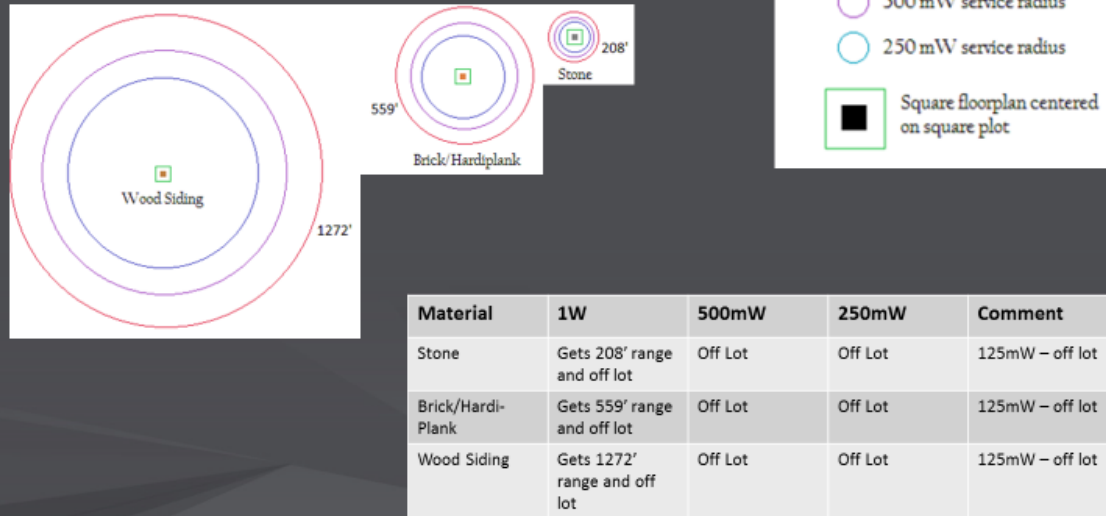


Figure 14 - Inside/Out Average Home Service Reach for Voice Call

2.1.3. Large Home (Mansion) Data and Voice Services



Figure 15 - Sample Mansions of Various External Construction

The mansion portion of the analysis presumed a 5,000 square foot dwelling on a 0.75 acre lot and an internal AP placement which would have met with three internal floor/ceiling transitions and an exterior wall in order to reach outside the home before encountering the foliage costs to propagation. The resultant service radii for the three grades of connection were calculated as follows:

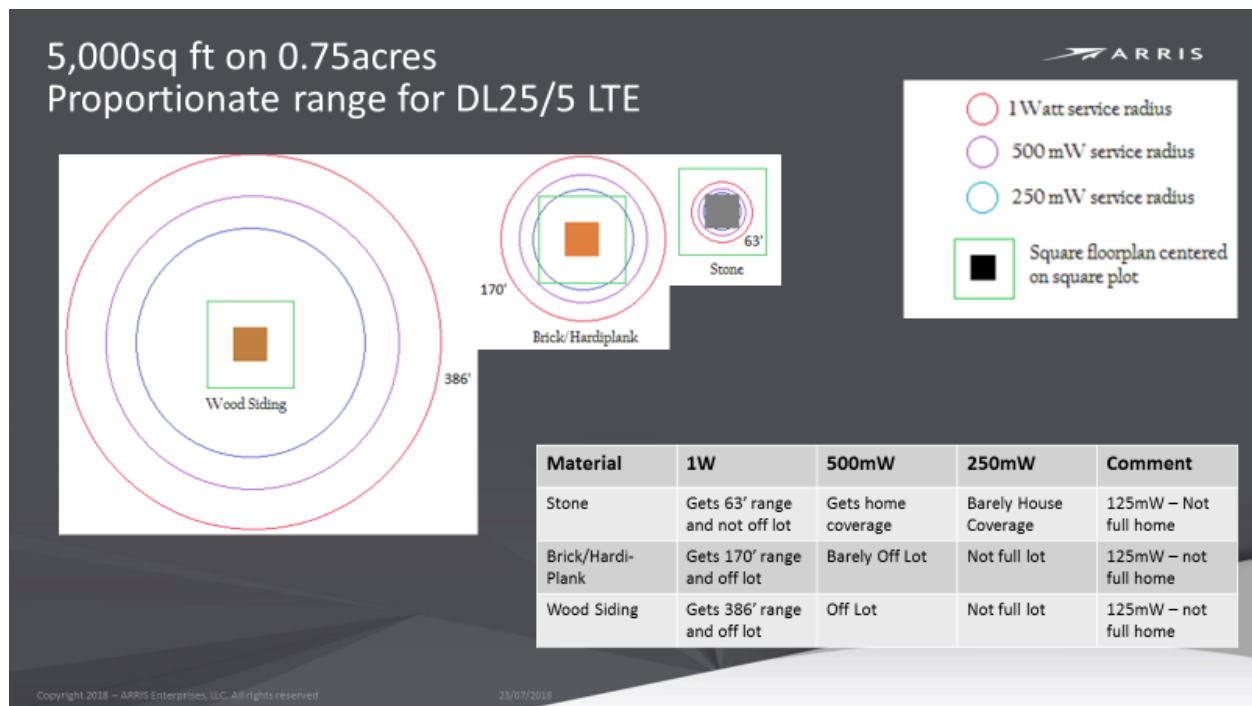


Figure 16 - Inside/Out Mansion Service Reach @ 25/5 Mbps

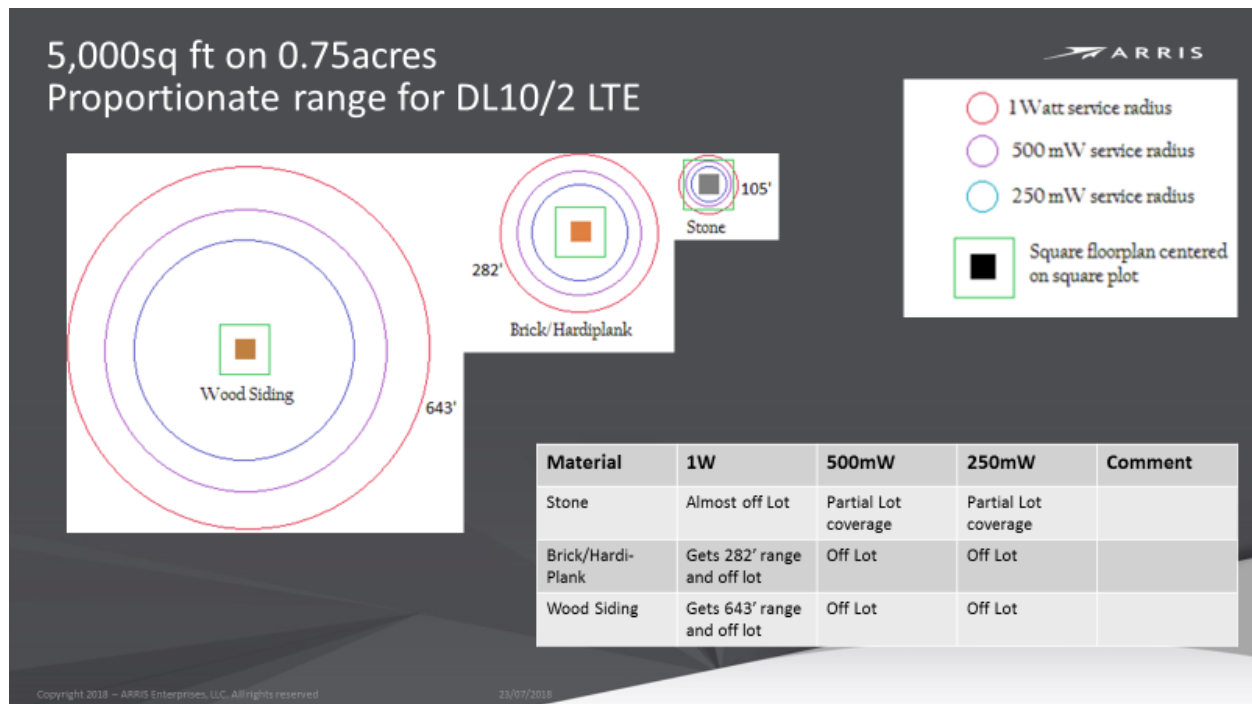


Figure 17 - Inside/Out Mansion Service Reach @ 10/2 Mbps

5,000sq ft on 0.75acres
Proportionate range for 100kbps voice call

ARRIS

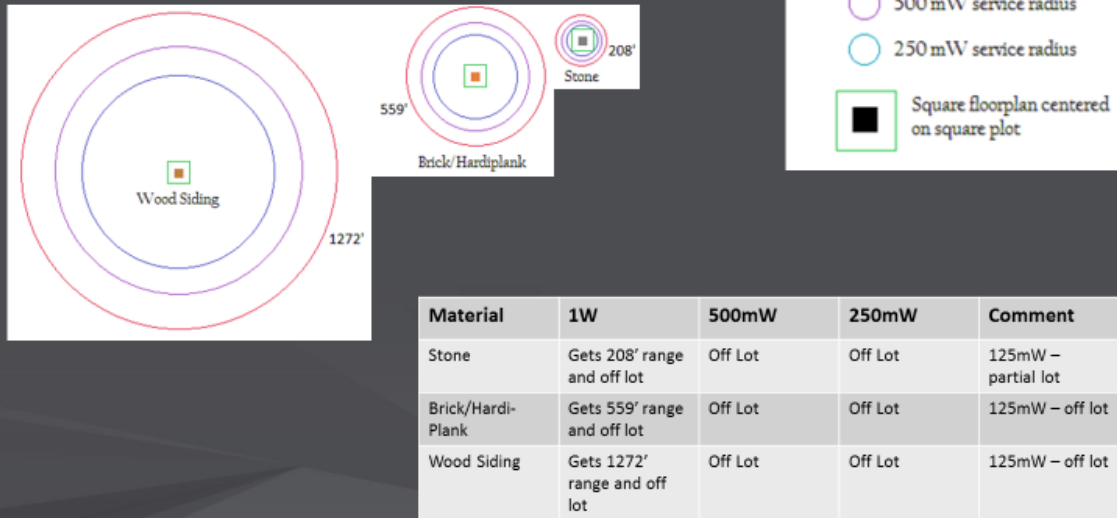


Figure 18 - Inside/Out Mansion Service Reach for Voice Call

2.2. Tabular Summary of Inside/Out 3.5 GHz CBRS AP Reach

AP Service Radius (feet) 25 Mbps (D) / 5 Mbps (U)

	Size (sq ft) \ Exterior	Stone	Brick	Wood
1W AP	1500	81	217	494
	2500	63	170	386
	5000	63	170	386

	Size (sq ft) \ Exterior	Stone	Brick	Wood
500mW AP	1500	63	170	386
	2500	49	132	302
	5000	49	132	302

	Size (sq ft) \ Exterior	Stone	Brick	Wood
250mW AP	1500	49	132	302
	2500	39	104	236
	5000	39	104	236

Figure 19 - 25/5 Mbps Service Radii Performance

AP Service Radius (feet) 10 Mbps (D) / 2 Mbps (U)

1W AP	Size (sq ft) \ Exterior	Stone	Brick	Wood
	1500	135	361	822
	2500	105	282	643
	5000	105	282	643

500mW AP	Size (sq ft) \ Exterior	Stone	Brick	Wood
	1500	105	282	643
	2500	82	221	502
	5000	82	221	502

250mW AP	Size (sq ft) \ Exterior	Stone	Brick	Wood
	1500	82	221	502
	2500	64	172	392
	5000	64	172	392

Figure 20 - 10/2 Mbps Service Radii Performance

AP Service Radius (feet) 100 kbps (D)/100 kbps (U)

	Size (sq ft)	Exterior		
		Stone	Brick	Wood
1W AP	1500	267	715	1627
	2500	208	559	1272
	5000	208	559	1272

	Size (sq ft)	Exterior		
		Stone	Brick	Wood
500mW AP	1500	208	559	1272
	2500	163	437	994
	5000	163	437	994

	Size (sq ft)	Exterior		
		Stone	Brick	Wood
250mW AP	1500	163	437	994
	2500	127	341	776
	5000	127	341	776

Figure 21 - Voice Call Service Radii Performance

2.3. 3.5 GHz CBRS Neighborhood Roaming Via HaaT (Home as a Tower)

Reviewing the data for inside/out coverage indicates that extremely dense exterior home construction (stone) puts paid to the notion that one might be able to roam outside very far with a band 48 (CBRS) mobile device – potentially not even being able to reach the limits of the property for a large (≥ 0.75 acre) lot before losing all but voice connection even with the most powerful 1W interior AP (not that such devices are all that desirable for an interior environment, with a ~ 400 cubic inch volume dissipating an estimated 26 W).

There is a qualifier which must be noted here regarding visible antenna placement within the home and the accurate observation that horizontal propagation from a low height within the dwelling (or even from the second floor) does indeed prove problematic. However, if the home is viewed as effectively a radome of sorts, then some amount of propagation relief can be had by elevating the antenna into the attic. This typically would involve a fair amount of cabling loss (say, 5 dB for 20 feet of coax at 3.5 GHz) and necessarily involve the use of an antenna element with sufficient gain to at least overcome the cabling losses. This is largely wasted effort in the case of wood-sided homes, since the roof aperture for the antenna pattern still involves end plates which cannot be mitigated and the relieved surfaces (as transition of roof sheathing, insulation and shingles) are not much less lossy than the wood siding itself. But for brick or stone homes, the pattern would certainly improve along the horizontal access where the dense

siding is replaced by roofing materials (unless, for example a slate, formed concrete or tile roof is involved). Aesthetic and access objections aside, the elevated interior antenna should show greatly improved footprint along at least one horizontal access compared to the buried case, presuming typical North American materials (since even composite, asphalt or stone materials are thin enough to reduce the comparative effects of several inches of masonry or stone.)

The potential of attic placement aside, the effect of ever-denser exterior construction in a real sample neighborhood on distribution of high power interior-only APs – and the resultant collapsing coverage footprint for high value data services are shown in this sequence of illustrations:



Scale: 200 ft/61 pels = 3.28 ft (1m) / pel

Total area shown: 136.3 acres
Total homes shown: 158

Wood-sided mansion (best case)
1W AP service radius is 386 ft for
25/5 Mbps

(Deployment every 6
homes shown)

(Effects of 1W APs in every 6th wood siding home)

St. Marlo Community Example

Figure 22 - Inside/Out Coverage for 25/5 Service Assuming Wood Sided Homes (1 x 6 Homes)



Scale: 200 ft/61 pels = 3.28 ft (1m) / pel

Brick mansion 1W service
radius of internal AP is 170 ft
(for 25/5 Mbps)

Deployment of AP every third
home shown

Total area shown: 136.3 acres
Total homes shown: 158

(Effects of 1W APs in every third brick home)

St. Marlo Community Example

Figure 23 - Inside/Out Coverage for 25/5 Service Assuming Brick Sided Homes (1 x 3 Homes)



Scale: 200 ft/61 pels = 3.28 ft (1m) / pel

With 4W strand and ~ 10-12 dB worth of foliage attenuation, expect ~ 700 ft service radius for 25/5 service.

Stone mansion (worst case)
1W service radius of internal AP is 63 ft for 25/5 service

Total area shown: 136.3 acres
Total homes shown: 158
Total 4W POPs: 6
Avg # homes/POP: 26

(Stone home 1W APs augmented by strand products)

St. Marlo Community Example

Figure 24 - Inside/Out Coverage for 25/5 Service Assuming Stone Homes (Yellow) and the Need to Augment with 4W Strand Mount APs (Red) (1 x 1 Home + 6 x 4W POPs)

As is evident in the progression above, 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 stone home placement become spaced by large lots – unless the coverage is augmented by exterior high-power APs.

This leads to a solution where we examine 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. Coverage maps of the same sample neighborhood are shown:



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 25 - 1W HaaT Coverage for 25/5 Service (1 x 10-16 Homes)

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):



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 26 - 4W HaaT Coverage for 25/5 Service (1 x 20-40 Homes)

2.4. Roaming User Considerations

Some data traffic notes are worthwhile here. As mentioned, the exterior AP's service reach also needs to scale with the potential number of roaming users within the service radius. Use of premium data does not necessarily suggest that downlink speeds are continuous (as in streaming). In fact, connected browsing users may exhibit duty cycles of only 25% or so, but we will assume worst-case streaming by all simultaneous users. Assuming further that the AP under study can be backhauled to the limit of its PAL channel carriage (70 MHz) and each of the simultaneous users is a different distance from the AP (but none outside the expected service footprint required for the 25/5 service – so in round terms at signal levels no worse than -80 dBm for the mobile device under consideration if we include the upstream data carriage considerations).

Next, we allow for distribution of the users equal distances from the AP (so the effective average signal level drives an MCS on their mobile devices to between 16- and 64-QAM – figure an average spectral density of around 4 bps/Hz). If the antenna is tri-sectored (and again, the users don't gather in a single sector) then you get a "x3" multiplier for the delivered spectrum. Under all of these (admittedly ganged) assumptions, your 4W AP would be delivering an ensemble to-mobile bitrate of 840 Mbps (280 Mbps to

each of 3 sectors) and you could be supporting 33 roaming users per AP (which ends up being roughly ~1 per household in the service area).

2.5. 1W HaaT Simulations

In order to corroborate the inferences produced by the lumped model calculations, direct simulation of existing Ruckus 1W devices (arrayed as those coverage calculations suggested above) was performed in order to validate the roaming footprint proposed. This type of simulation permits the otherwise averaged effects of multipath propagation, antenna pattern, and topological interactions to be directly calculated via aggregation of ray-trace power measurements. A side-by-side comparison of the two 1W HaaT findings is immediately below:



Figure 27 - 1 W HaaT Simulation Vs Calculation; ~ 2:1 Service Radius Improvement in Sim

The 2:1 service radius differential (~ 800 ft vs ~ 400 ft) amounts to roughly 7-8 dB worth of pessimistic additional link loss in the lumped calculations (shift in the path loss vs distance curve fit, or overly excessive foliage losses, for example). This might also be indicative of the differences between the environment measured in Tampa and the neighborhood characteristics outside Atlanta. The simulation also provides greater granularity in the topographic feature effects associated with bitrate support degradation with increasing distance from the AP, better showing how radially non-homogeneous the coverage can actually be:




Figure 28 - 1 W HaaT Simulation, Bitrate Variance in Coverage Map

If we presume that the lumped calculations serve as a reasonable indicator of the service level available at every intermediate point within their bounded radius, then a rough estimate of the mesh density required for 3.5 GHz 4W strands or pedestals for the cases of open range distribution and mildly forested scenarios can be made:

Service Reach of 4W strand mount to mobile @ 3.5 GHz

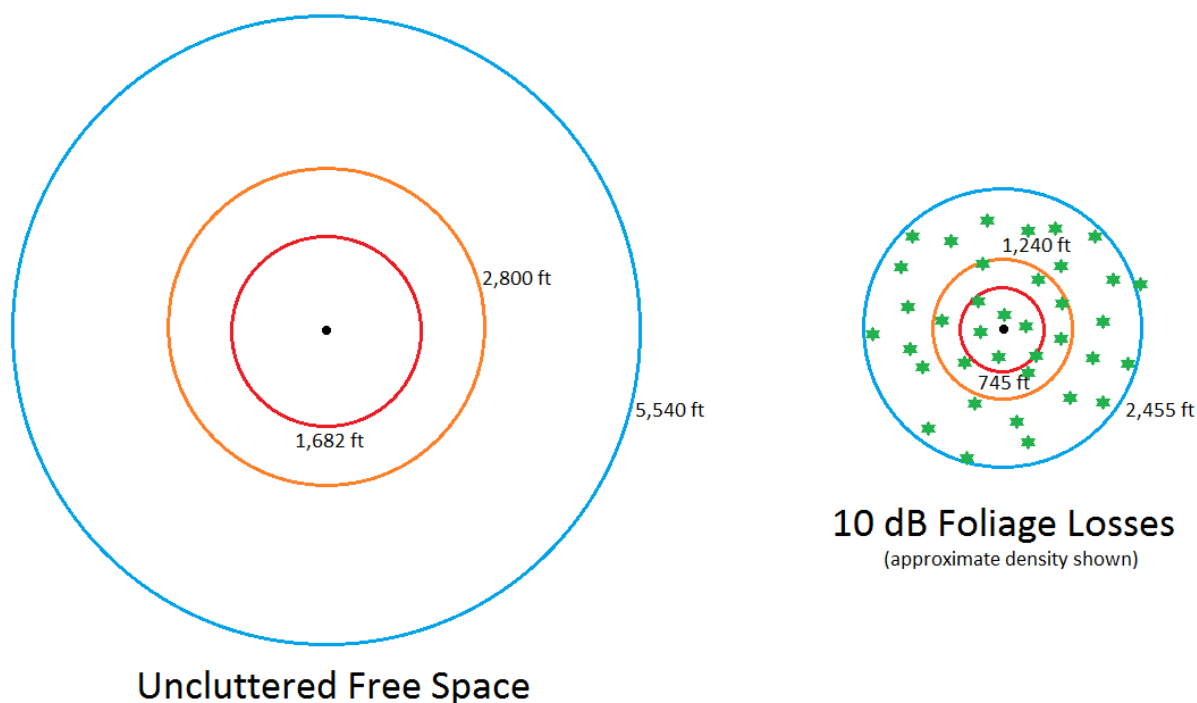


Figure 29 - Service Mounting Potential for Outdoor 4W CBRS Mesh

Recall that the simulation difference to the lumped estimation calculations yielded a 2x improvement; applied to the estimates above, it suggests that in open range conditions, 2 kilometer AP spacing might yield fairly good data coverage.

2.6. CBRS/LTE and IoT

A quick address of the implications of using 3.5 GHz CBRS as a potential out-of-band solution for IoT is in order. As might be expected, CEDs operating at the 2.4 GHz ISM band via the NFC MAC of choice would certainly benefit from either less competition in-band or the emergence of another band option outright. At issue is the cost of MAC/PHY in the new band and in particular, the availability of scale economies from widespread leverage of that new band (better still, from that standpoint, if the band use is unlicensed). CBRS may offer some relief in this aspect, since its tiered SAS support anticipates both licensed users (who might then bear the cost of scale economies in question) and GAA participants (who then benefit as an “interested second market” from MAC/PHY chip solutions which have to be developed for the license-based crowd).

At the moment, there is a cost penalty associated with CEDs moving from a legacy 2.4 GHz NFC radios and onto NB-IoT support offered by LTE at 3.5 GHz. (Hence, the interest in coupling other services – voice and premium data – into the move). However, catalytic cost benefit from licensed spectrum

leverage requires watching to see when/if a cost inflection point is reached. Certainly, there is market pressure building which seeks to resolve the wireless service contention at 2.4 and it must be solved in some fashion. This is a situation whose dynamics beg monitoring.

2.7. Timing Considerations Across LTE and Cable Domains

The issue of transit between an LTE-based domain (roaming) to a home interior which is beholden to DOCSIS timing considerations is under active consideration by CableLabs. Without exposing intellectual property interests in this area, it can be said that the two domains' synchronization in general is covered by a common reference to GPS timing. However, the LTE domain (and in particular, 5G) is set to trigger on a finer granularity of latency than DOCSIS (roughly one order of magnitude, if not almost two). This difference is bridgeable, as a general rule, if the LTE network apprises DOCSIS of the impending handoff so that the latter can schedule the required wireline packet availability (essentially slaving, via alert messaging, the DOCSIS network's chunkier operation to the LTE's near-1 msec latencies.)

3. LoRa

LoRa has been more of an unknown commodity than 3.5 GHz CBRS due to its recent adoption timeline and lack of MVNO interest. LoRa amounts to a purpose-built out-of-band IoT service network scheme supporting very low-power endpoints and a native distributed star topology designed for robust and (somewhat) timely relay of IoT small-packet, spread-spectrum narrowband communications. On the cloud network edge, LoRa operates in the USA at the very well characterized ISM 900 MHz band which features an improved through-air loss tangent and better materials penetration than 3.5 GHz CBRS or any of the Wi-Fi bands – as would attend inside-out communication.

LoRa also promotes an opportunistic leverage of any conveniently available wireline IP network backhaul due to its ability to curate and cache repeat IoT packets at its edge aggregation points. This, plus its deep RF link budget, seem at first blush to facilitate an organic location protocol for its base stations. As more IoT clients get seeded, only rough triangulations should be needed to calculate where best to locate new bases (and the proximity of suitable wireline backhaul nodes can be baked into the estimations).

Furthermore, LoRa projects a naturally hardened aspect given the “repeated packet” culling and roundtrip latency calculus which can be done at its base stations. The effective star topology by default should then define a best-path solution based upon first successful edge discrimination even though there may be redundant receptions at multiple base stations. More to the point, perhaps, is that LoRa allows network service providers to abstract IoT management considerations away from operation of the main IP network (IoT verticals are constructed as a secure tunnel of application server(s) x addressable CEDs – the ‘x’ proxying the IP network backhaul as mere crosspoint switch). Such a schematic is captured below:

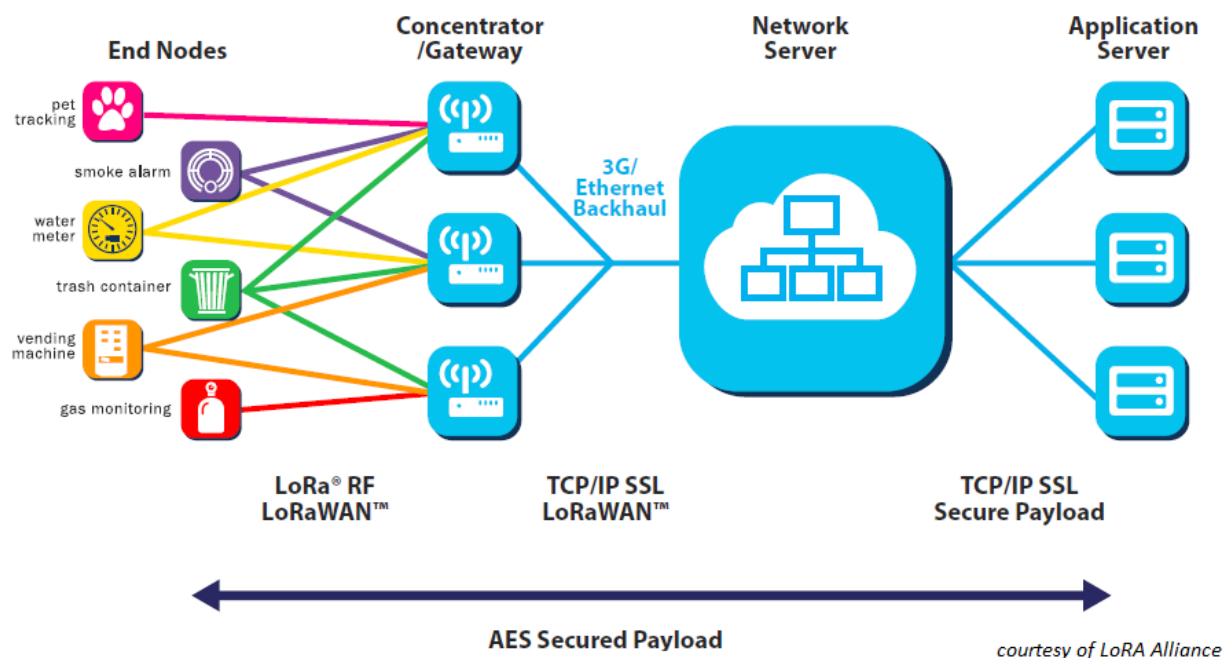


Figure 30 - Schematic of LoRa Edge Network

LoRa's approach to leverage of the ISM 900 MHz band involves the use of randomly channel-hopped, low-bitrate CSS (chirped spread spectrum). The approach is essentially chipping a data stream and then uses that to modulate a chirp waveform on an uplink-prioritized band bifurcation which places 500 kHz downlinks in the upper ~ 5 MHz of the band and two classes of uplink (64 x 125 kHz or 8 x 500 kHz channels) in the lower ~ 13 MHz of the band. The arrangement renders itself as follows:




Figure 31 - LoRa US ISM 900 MHz Band Occupation

As mentioned above, from a link noise management perspective, LoRa uses a spread spectrum chirp and chipping of the underlying data to perform constant-envelope modulation of the selected channel carrier. Part of its noise adaptation mechanism is to apply additional spreading of the signal by essentially chipping a lower rate stream with a higher sampling rate and buying discriminator margin in (very) roughly similar fashion to an OFDM/QAM stream backing off its constellation density (MCS reduction –

lowering its spectral density). The cost to bitrate versus noise margin for some of the spreading factors is listed in the following table:

SF (Spreading Factor)	Chips/Symbol	SNR limit	Time on Air for 10 byte packet (ms)	Bit Rate (bps)
7	128	-7.5	56	5470
8	256	-10	103	3125
9	512	-12.5	205	1758
10	1024	-15	371	977

Figure 32 - LoRa US Spreading Factor Implications to Noise Margin and Bitrate

Another illustration of the relationship between higher rate chipping, threshold of acceptable performance, and the impact of these parameters on sustained bitrate and range is captured here:




Figure 33 - LoRa Tradeoff of Sustained Bitrate for Better Range

3.1. Base Stations and CEDs

The LoRa cloud-edge base stations (which can use various backhaul technologies to attach to the cloud) serve as one terminus of the 900 MHz link. At the IoT actuator or sensor end lie three types of client devices, partitioned per their respective battery draws and communication latency needs. Type A devices are strictly the very lowest power sensors which randomly transmit (based upon event or internal watchdog timing) and wait two slotted periods for a base station ACK. (This specified random access behavior generates some quite-unwanted 2nd order effects, about which more later.) Type B devices are also battery powered devices but higher energy consumers as they are actuators which must process a timing beacon to constrain control loop latency (and establish wake/sleep periods). And Type C devices are high power actuators expected to operate off AC mains and thus maintain a constant wake state. The implications of the parsing are expressed below:

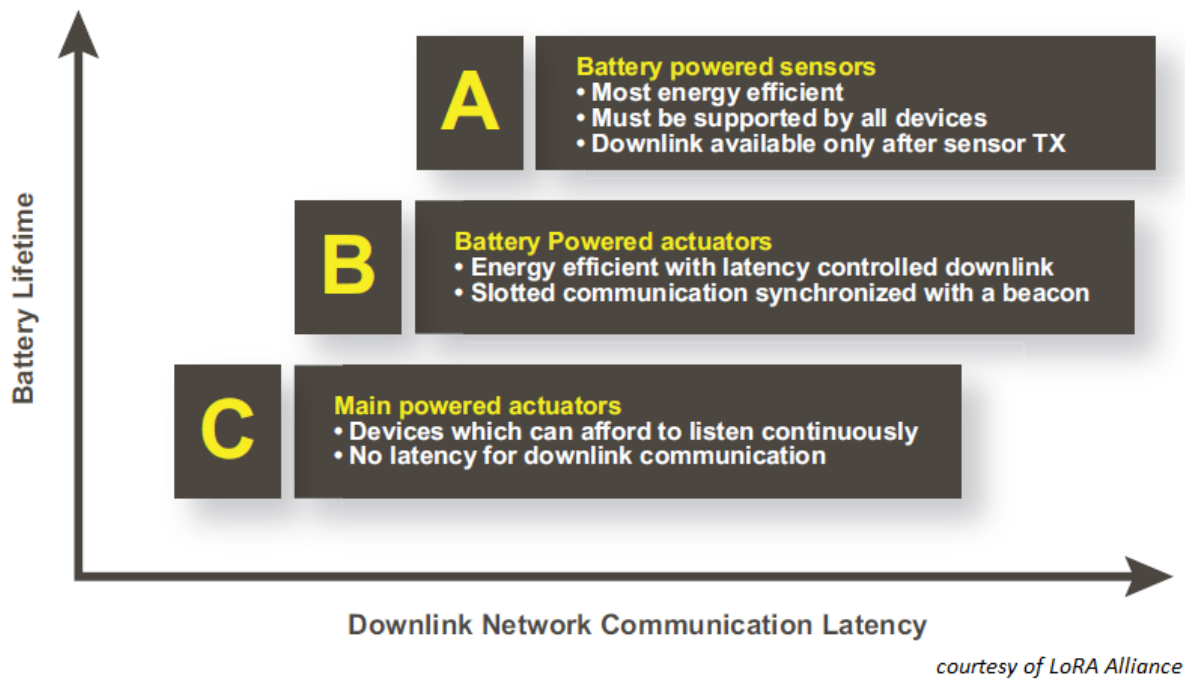


Figure 34 - LoRa Client Mix: Sensors and Actuators

3.2. Link Specifics

The associated link parameters are tabulated below. Considerations of integrated RF power under different spreading considerations suggest an actual maximum power of perhaps +21 dBm (and such is in keeping with maximizing battery life). Note especially the deep link budget even with such modest transmit power:

Frequency band	902-928MHz
Channels	64 + 8 +8
Channel BW Up	125/500kHz
Channel BW Dn	500kHz
TX Power Up	+20dBm typ (+30dBm allowed)
TX Power Dn	+27dBm
SF Up	7-10
Data rate	980bps-21.9kpbs
Link Budget Up	154dB
Link Budget Dn	157dB

courtesy of LoRA Alliance

Figure 35 - LoRa US ISM 900 MHz Link Parameters

3.3. Link Budget and Service Throw (Range)

LoRa offers two compelling recommendations for its consideration: 1) it operates in the recently (1985) created and sparsely used (as in: only partial band exploits) 903-928 MHz ISM space and 2) it employs a robust CSS modulation scheme with coding gain and random (though mask-controlled) channel assignments to extract link budgets approaching 160 dB in some cases – promoting huge potential operating range for the RF link from base station to addressable CED (and more importantly, back).

The following is tabular performance data of a LoRa endpoint device using commercially available silicon:

Mod	Data rate	Frequency (MHz)	Sensitivity (dBm)
LORA	SF7BW125	902.3	-125.1
LORA	SF7BW125	908.7	-125.8
LORA	SF7BW125	914.9	-125.9
LORA	SF10BW125	902.3	-133.7
LORA	SF10BW125	908.7	-134.5
LORA	SF10BW125	914.9	-134.7
LORA	SF8BW500	903	-122.7
LORA	SF8BW500	907.8	-123.4
LORA	SF8BW500	914.2	-123.6

courtesy Semtech

Figure 36 - Measured Sensitivity of Commercial LoRa Product

The numbers above correspond to operation at a PER threshold of 10 %. Using a web-available LoRa calculator for receiver sensitivity (<http://www.rfwireless-world.com/calculators/LoRa-Sensitivity-Calculator.html>) and seeding it with the appropriate spreading factor and BW numbers yields a calculated NF of 4 dB (averaged across all 9 data points). This is an excellent implementation. More to the point, the realizable bounds of +20 dBm transmit and -134 dBm receiver sensitivity play out as easily surpassed 20+ km LOS reach. (City reach is topography-specific: as in all NLOS wireless propagation calculations, the link budget is reduced by through-material transitions – getting out of the CED’s housed environment – and lost scattering in addition to the classic frequency-dependent free-space losses.) Propagation models to handle multiple diffractive paths are beyond the scope of this paper yet are worth separate investigation – the Egli model with its VHF/UHF television heritage seems a reasonable choice in areas where tree/building interferers are not common. The long story short is that 20 km LOS is a reasonable range for LoRa and this might be reduced by a factor as high as 10 in extremely dense urban environments.

As an example set of calculations and observations, there is an Egli model calculator available on the internet (<https://www.commscope.com/calculators/qegli.aspx/>). Using this resource and seeding values for 4 meter heights for both the base station and CED (so, 2nd residential floor for both) yields a loss estimate of 142 dB at 915 MHz and 3 km spacing. (This provides a generous 12 dB budget of lumped losses to transit buildings, for example).

Semtech builds silicon for LoRa radio implementations and offers the following performance observations (green balloons are successfully exchanged transmissions; had there been dropped packets these would have been shown in red):




Figure 37 - Sample LoRa Suburban Mobile Inside/Out Connectivity @ 1300' Radius



Figure 38 - Sample LoRa Rural Mobile Inside/Out Lossless Connectivity @ 49 Km (!) Radius

3.4. Geolocation of Clients as Calculated Benefit

The LoRa network topography is accurately termed an extended star. However, the fact that multiple base stations receive transmitted CED packets offers an opportunity to use packet arrival statistics (timestamp and RSSI) to build a triangulated representation of their location since base stations are GPS timed and located themselves. As long as at least 3 base stations receive a particular packet, rough estimates of client locale can be built from the near-intersections of weighted radii – the weights associated with signal strength (RSSI) or what is called TDoA (timed difference of arrival). As shown in the following figure, both techniques are inherent in LoRa and so do not represent costly appropriation of additional capability but merely exploit of a simple calculus on existing message data. The scale of the drawing is a bit misleading; examination of the error terms makes it plain that the accuracy of the TDoA exploit is roughly 10x that of the triangulated RSSI. As might be expected, the addition of additional base stations – particularly in a geographical pattern which yields as close to equal path lengths as possible – produces the most accurate results.




Figure 39 - LoRa TDoA (Time-Difference of Arrival) and RSSI Geolocation Accuracy

The tracking capability has only recently become the subject of interest from the LoRa Alliance; a formal whitepaper regarding implementation details was released in 1Q 2018. Part of the engineering studies performed to validate error sources (and thus, suitability) of the GPS-based timing in what could be problematic multipath environments yielded the following estimations for circular error probability (CEP) in the location estimates produced. Noteworthy (and expected) aspects confirmed that long-throw rural estimations produced tighter CEP results than progressively more spread/scattered results as path diffractions densify in more urban settings:




Figure 40 - LoRa Timestamp Error Contribution to Geolocation Error

3.5. The Issues of Scale

In the client device description in this paper, a bookmark regarding class A devices employ of what amounts to ALOHA signaling was lodged. This is perhaps the most cringeworthy shortfall in the LoRaWAN protocol, as of course the damage to throughput (even under the condition of randomized channel selection) bears the unmistakable imprint of an ALOHA asymptote. Though the convergence to poor throughput is mitigated by the random channel hopping scheme and the orthogonality of the chirping modulation (essentially, your throughput approaches the sum of multiple simultaneous ALOHA schemes, one for each channel and spreading factor). In the IEEE Communications Magazine of January 2017 the research paper “Understanding the Limits of LoRaWAN” addressed the standard’s shortfall in regards to scalable applications and some of it is referenced here.

Note that it is not only an ALOHA congestion issue but one that is subject to FCC rules on band occupancy dwell and duty cycle. In brief, the FCC puts occupancy restrictions on the CSS spectrum employed by LoRa. For the 125 kHz channels, this specification is no more than 400 msec of per-message transmission every 20 seconds. For the 500 kHz channels, the restriction is less restrictive: no more than 400 msec every 10 seconds. This impacts the time on-air, packet sizes and spreading factor as follows:




Figure 41 - LoRa Message Dwell Time Implications to Payload and SF

Note that the dwell time specification precludes use of spreading factors 11 and 12 in the US. Furthermore, the most noise-immune SF left to use (10) must be restricted to the transport of no more than 24 MAC payload bytes.

The resultant ensemble, asymptotic throughput behavior looks very familiar to those familiar with ALOHA congestion – albeit with the inclusion of a hard limit due to duty cycle off time restriction:




Figure 42 - LoRa Congestive Behavior Due To Closed Loop ACK of ALOHA Upstream Messages

4. LoRa Compared Vs LTE Narrow Band IoT Options

LTE has taken several swipes at establishing a scheduling mechanism to handle IoT signaling needs, appearing originally to accept that radio costs and power draws would likely disqualify it from ever scaling down to CEDs themselves. A shared LTE host (mobile phone) for 2.4 GHz ISM NFC-based CEDs has been one key driver for establishing the necessary small-packet handling priorities in the larger network. The other, much more critical aspect, has been the need for extremely low latency industrial IoT (IIoT) control environments to be in place to handle the real-time needs of the manufacturing sector(s). As it turns out, the lack of determinism in control loop latency and very modest signaling bitrates for LoRa disqualify it from competing in that role, so there appears to be a natural gap between high value/high accuracy/low-latency commercial IoT applications with LTE support and cost-sensitive, light industrial, asset-tracking or consumer-end (and, as regards latency, more casual) IoT support which LoRa can underpin. On the plus side for LoRa are its much less expensive implementation BOM (for both ends of the RF link) and extremely usable low-current modes (principally for class A CEDs, though some B's might qualify). A tabular breakdown of the differences follows:

Feature	LoRaWAN	LTE Cat-1 2016 (Rel12)	LTE Cat-M 2018 (Rel13)	NB-LTE 2019(Rel13+)
Modulation	SS Chirp	OFDMA	OFDMA	OFDMA
Rx bandwidth	500 KHz	20 MHz	20 - 1.4 MHz	200 KHz
Data Rate	0.98 - 21.9 kbps	10 Mbit/sec	200kbps – 1Mbps	~20K bit/sec
Max. # Msgs/day	Unlimited*	Unlimited	Unlimited	Unlimited
Max Output Power	20 - 30 dBm	23 - 46 dBm	23/30 dBm	20 dBm
Link Budget	154 dB	130 dB+	146 dB	150 dB
Battery lifetime - 2000mAh	105 months		18 months	
Power Efficiency	Very High	Low	Medium	Med high
Interference immunity	Very high	Medium	Medium	Low

*** < 400 msec / msg dwell. Ultimately capped by FCC duty cycle limits. (upstream limit)**

courtesy of LoRa Alliance

Figure 43 - US LoRa Comparison to LTE as LPWAN

Conclusions

The home's casual and largely organic adoption of Wi-Fi as its favored brand of wireless communications has seen this haphazard marketing play challenged first by self-handicapping via unmanaged standards supersession, then peer overcrowding and pre-emption (the penalties for success), and now finally by direct competition from unlicensed co-participants in the 2.4 GHz ISM band associated with the emergence of various IoT vertical businesses. While Wi-Fi's assimilation of the 5 GHz band and the introduction of airtime management have begun the process of distributing and scheduling RF energies in and out of the old band, such relief-valving involves a protracted remediation schedule and does not completely resolve some of the service competition issues associated with sharing 2.4 GHz among so many perspective clients. Two new band opportunities have become available for consideration, both as options in the service contention solution space and as outright disrupters as regards the ability they give providers to mine new opportunities in the home: 3.5 GHz CBRS/LTE and 900 MHz LoRa.

Both wireless technologies provide options to move at least IoT data out of the crowded Wi-Fi space and into alternate bands for backhaul and in doing so, either harden the IoT services involved against casual or perhaps even targeted interruption. Both feature new radio packaging for CEDs (though LoRa's is less expensive). Both involve investment in the overlaid wireless portion of the backhaul – though LoRa's longer wavelength and throw indicates that the distribution of concentrators (base stations) can be sparser than the seeding of 3.5 GHz support. Scale economies seem to favor LoRa as well – but only to the limit where upstream ALOHA-based congestion throttles the star aggregation scheme. For its part, however, 3.5 offers a much broader support bandwidth which would enfranchise a wider palette of IoT devices (most specifically cameras – whose ad hoc home utilization seems to easily outstrip other smart home connected appliances of more modest signaling requirement.)

As the pitch of complaints against oversubscription of the 2.4 GHz ISM space mount, it will do well for operators to examine other band solutions for the hosting of emerging smart home and aging-in-place businesses which will demand (perhaps life-critical) service availability at the three nine's level and beyond; 3.5 GHz CBRS/LTE and LoRa offer potential solutions to this problem.

Abbreviations

AC	Alternating current
ACK	Acknowledgement message
AP	Access point
BLE	Bluetooth Low Energy
BOM	Bill of material
bps	Bits per second
BW	Bandwidth
BYO	Bring-your-own
CBRS	Citizens Band Radio Service
CBSD	Citizens Band Radio Service Device
CED	Constrained end device
CEP	Circular error probability
CPE	Consumer premises equipment
dB	Decibel
dBm	Decibel referenced to 1 milliwatt
DOCSIS	Data-over-cable service interface specification
DoD	Department of Defense
EIRP	Effective isotropic radiated power
ESC	Environmental sensing capability
FCC	Federal Communications Commission
FDM	Frequency Division Multiplex
FEC	Forward error correction
GAA	General authorized access
GHz	Giga-hertz
GPS	Global positioning system
HaaT	Home as a Tower
HFC	Hybrid fiber-coax
HD	High definition
Hz	Hertz
IoT	Internet of Things
ISBE	International Society of Broadband Experts
ISM	Industrial, Scientific and Medical
ISP	Internet service provider
kHz	Kilo-hertz
Km	Kilometer
LTE	Long Term Evolution
MAC	Media Access Control
Mbps	Mega-bits per second
MCS	Modulation

MHz	Mega-hertz
Msec	Milliseconds
MVNO	Mobile virtual network operator
NB	Narrowband
NF	Noise figure
NFC	Near-field Communication
OFDM	Orthogonal frequency division multiplexing
PAL	Priority access license
PER	Packet error rate
PHY	PHYsical Layer
QAM	Quadrature amplitude modulation
RF	Radio frequency
SAS	Spectrum Allocation System
SCTE	Society of Cable Telecommunications Engineers
SF	Spreading factor
SNR	Signal-to-noise ratio
TDD	Time division duplex
TDM	Time division Multiplex
UHF	Ultra high frequency
VHF	Very high frequency
W	Watt

Bibliography & References

AN 1200.04 Application Note : *FCC Regulations for ISM Band Devices: 902-928 MHz*, copyright 2006 Semtech

AN 1200.22 Application Note : *LoRa Modulation Basics*, copyright 2015 Semtech

Can a Fixed Wireless Last 100m Connection Really Compete with a Wired Connection and Will 5G Really Enable this Opportunity?, J.R. Flesch et al; copyright 2017 SCTE-ISBE and NCTA

Geolocation Whitepaper, LoRa Alliance, January 2018

LoRaWAN 1.1 Regional Parameters, copyright 2018 LoRa Alliance

LoRaWAN 1.1 Specification, copyright 2017 LoRa Alliance

Understanding the Limits of LoRaWAN, Ferran Adelantado et al, IEEE Communications Magazine, January 2017

Toward Automated Intelligent Resource Optimization for vCMTS Using Machine Learning

A Technical Paper prepared for SCTE•ISBE by

Kieran Mulqueen

Research Project Manager
Intel Labs, Intel
Collinstown Industrial Park, Leixlip, Co. Kildare
353 (1) 606 3789
kieran.m.mulqueen@intel.com

Michael O'Hanlon

Principal Engineer
Network Products Group, Intel
Dromore House, East Park, Shannon, Co. Clare
353 (61) 777 808
michael.a.ohanlon@intel.com

Marcin Spoczynski, Intel Labs, Intel

Brendan Ryan, Network Products Group, Intel

Thijs Metsch, Intel Labs, Intel

Leonard Feehan, Intel Labs, Intel

Ruth Quinn, Intel Labs, Intel

Table of Contents

Title	Page Number
Toward Automated Intelligent Resource Optimization for vCMTS Using Machine Learning	1
Table of Contents	2
Introduction.....	3
Enabling Insight-Driven Management and Orchestration of NFVI Nodes	3
Utilizing Machine Learning in a vCMTS scenario	5
Initial Results Demonstrate Key Areas for Optimization	8
Conclusion.....	9
Abbreviations	11
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1 - The analytics-driven MANO is supported by long-running observations in a (continuous running) background flow, and a runtime-driven foreground flow in which we deal with fine-grained VNF requests and rebalancing decisions²	4
Figure 2 - Overview of vCMTS scenario configuration in a real-world context.....	6
Figure 3 - Testbed setup detailing the network flows between the traffic generation host and several instances of the vCMTS data plane	7
Figure 4 - Sample metrics as collected by the telemetry system Energy usage of the sockets, as well as memory, are shown in relation to metrics about the usage of the CPU (Instructions Per Cycle and Retiring).....	8
Figure 5 - Showcases variance in energy measurements for two key metrics when two vCMTS VNF instances either share a CPU core (shared) or are pinned to separate cores (exclusive)	9

Introduction

Virtualization of Cable Modem Termination Systems (CMTS) provides a large range of benefits for operators within the cable industry. Ensuring that this virtualized CTMS architecture can meet required capacity and performance objectives, both current and forthcoming, provides essential assurance to operators in determining whether to adopt this approach. Typically, human-derived policies are being used for the management of the virtualized CMTS (vCMTS), however these are generally static and can result in over- or under-utilized resources and, eventually, in missing service objectives.

The introduction of Network Functions Virtualization (NFV) paves the way toward autonomous management of Virtualized Network Functions (VNF) like vCMTS. Autonomous NFV management using machine learning (ML) shows huge potential to optimize such a system, benefiting both service providers and customers. Machine learning allows for the generation of models that ensure reliability and dependability of the overall system while minimizing running costs and improving service assurance.

To generate analytic models capable of enabling autonomous management capabilities, the NFV Management and Orchestration (MANO) components need to be presented with key behavioral insights of the VNFs and infrastructure resources.

Here we describe an approach that allows derivation of such insights, which facilitates its use in various models that can be used in a Network Functions Virtualization Infrastructure (NFVI). We also show initial proof points demonstrating areas where machine learning can be used for immediate effect.

Enabling Insight-Driven Management and Orchestration of NFVI Nodes

Software Defined Infrastructures (SDI) as the foundation of NFVIs enable the decoupling of VNFs from the physical infrastructure they run on, which allows for innovations in service delivery: mainly by enabling VNF agility and service assurance.¹ Management and orchestration activities can play a role at various scales. This includes the management of clusters of nodes, as well as the management of individual nodes.

The physical infrastructure is handled as a resource pool, which allows for ease of VNF management. Mapping the VNF instances to the physical infrastructure is the task of the MANO components. By leveraging machine learning-based analytics to, for example, predict resource demands, it is possible to better facilitate various Service Level Agreements (SLAs) of the VNFs while in parallel achieving better resource utilization and reduction of TCO.

To achieve this level of autonomous management of VNFs and infrastructure resources in an NFVI, several components are required:

1. VNF and infrastructure resource aware MANO components, such as the NFV Orchestrator (NFVO); VNF Manager (VNFM) and Virtual Infrastructure Manager (VIM);
2. Background flow analytics enabling observation of VNFs and infrastructure resources over longer time frames;

3. An Information Core which captures insights on application and resource behavior;
4. Foreground flow analytics utilizing the Information Core to deal with VNF requests and rebalancing decisions

Using these concepts, autonomous MANO of VNF instances and the NFVI can be achieved as described in Figure 1.

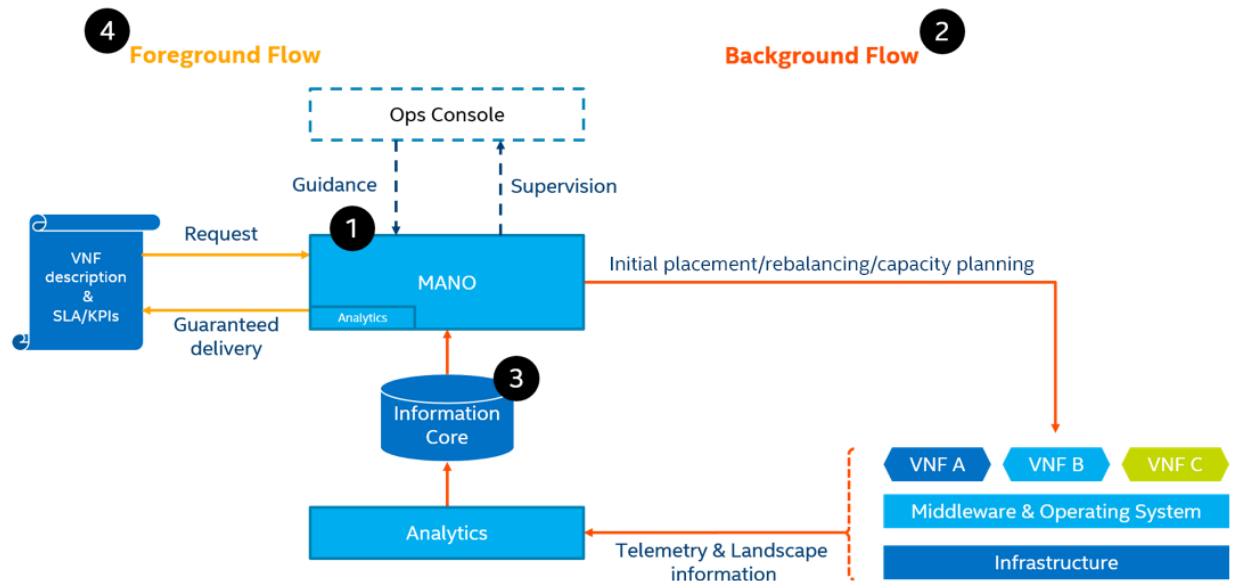


Figure 1 - The analytics-driven MANO is supported by long-running observations in a (continuous running) background flow, and a runtime-driven foreground flow in which we deal with fine-grained VNF requests and rebalancing decisions²

Within the MANO-enabling components, machine learning derived insights can play several roles for various activities:

- a. Determining appropriate **initial placements** of the VNF components allows for tighter packaging. This can be carried out using classification of the VNFs based on available data, allowing for the prediction of future behavior when placed on given resources. This can also predict interactive and interfering behavior resulting from placement with VNF instances sharing resources.
- b. Periodic **rebalancing** of the VNF to resource mapping can be triggered only when the benefit derived from migrating, scaling, or tuning the VNF configurations would outweigh the costs of the rebalancing activity.
- c. **Capacity planning** allows the forecasting of infrastructure resource capacities.

This system will enable the autonomous management of NFVI: to ensure that an autonomous system can still be supervised and guided it is proposed that it connect to the appropriate operations console.

Machine learning-based analytics that are embedded in the background flow allow for creation of insights in the form of models that capture behavioral patterns of the VNFs and the infrastructure.³ The models are eventually stored in the Information Core. As this background flow is enabled to run continuously, the models present in the environment can be regularly updated and adapt to changes.

Since various machine learning algorithms exist, the resulting models also have various forms. This includes, but is not limited to, decision trees, regression models, or neural network (NN)-based models. In addition to this, machine learning can facilitate the automated testing of alternate model types, which given the increased amount of data collected over time, might be more suitable than the original model architectures provided to the system. For example, various deployed model outputs can be traded off against each other and, if threshold levels of accuracy are reached, replace those currently in action. Also, continuous exploration of the original model's hyperparameters can be achieved through the mechanism described earlier.

It must be noted here that the background flow can either be enabled on single nodes on the NFVI or be offloaded (if needed, based on available compute capacity). This would mean that models trained elsewhere simply need to be embedded on the nodes for scoring and inference.

When the VNFs are deployed, the VNF manager, as well as the NFV Orchestrator, can use several other machine learning based analytics techniques within the foreground flow. This mainly includes notions of enabling anomaly detection. If the infrastructure itself is predicted to move to an undesirable state an intervention can be made prior to its deterioration. Given the past behavior of classified VNFs and the infrastructure resources at play, the future behavior of VNFs can be predicted and compared to the actual behavior to determine any possible mitigation actions. These mitigations include pausing, killing, or relocation of the VNFs or tuning configurations, such as altering the clock frequency of the CPU or activating available offloading capability. Further, an excessively anomalous result during the foreground analytics points toward suboptimal models being embedded in the MANO stack, and requires further mitigation of the same.

Overall, the integration of machine learned models is key to enabling a responsive, adaptable system. Through the foreground and background flow, enabled analytics models could be created to generate more precise models (e.g., for use in capacity planning activities), rather than ones based entirely on historical data.

Utilizing Machine Learning in a vCMTS scenario

Introducing the insights derived from machine learning can provide solutions to the numerous activities described earlier, such as optimized initial placement, rebalancing, and capacity planning. While the vCMTS VNFs can be relatively static in terms of infrastructure requests, differences in user activity, channel configurations, subscriber density, and modem encryption types can impact the required infrastructure resources. Further, it is likely that vCMTS VNFs will be executed alongside other VNFs, making appropriate placement a key problem to address. To deal with such a mix of VNFs and their requirements, the available physical infrastructure can be considered a resource pool, providing opportunity for optimization on several fronts, while maintaining adherence to specific performance Key Performance Indicators (KPIs).

Figure 2 gives a high-level overview on how a set of vCMTS VNF instances are deployed on physical infrastructure using [Kubernetes](#) as a VIM.

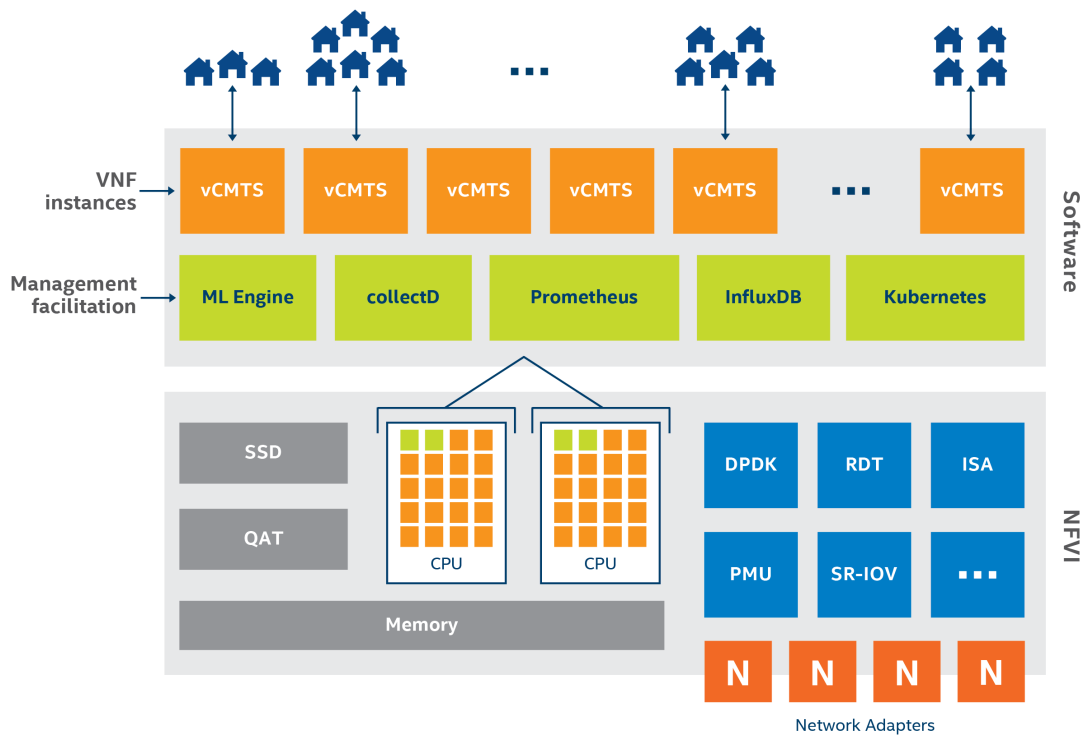


Figure 2 - Overview of vCMTS scenario configuration in a real-world context

For preliminary investigation, we have set up the following environment to demonstrate how an ideal initial placement of the VNF instances on single host NFVI can be achieved. Better initial placement can have a huge benefit for the infrastructure resources provider in terms of efficiency and performance and, therefore, profit.

Our baseline VNF is an Intel-developed vCMTS data plane pipeline. This was previously used to investigate performance on an Intel® Xeon® Scalable processor-based platform⁴ and contains a DOCSIS MAC* data-plane running on IA heavily utilizing the Data Plane Development Kit (DPDK).⁵ Each VNF instance is defined as a service group handling a collection of end users and is pinned to a core of the CPU.

Our testbed is based on Intel's latest architecture in a two-socket configuration, with two Intel® Xeon® Gold 6148 processors. Each processor contains 20 cores with a default frequency of 2.4 GHz. Furthermore, two Intel® QuickAssist Technology (Intel® QAT) adapters are installed, providing hardware acceleration for offloading of crypto- and compression-based tasks. Four Intel® Ethernet Network Adapters are installed and are connected to a separate traffic generation host allowing for high bandwidth testing. Each Network Adapter provides two ports, and hence up to eight physical network connections.

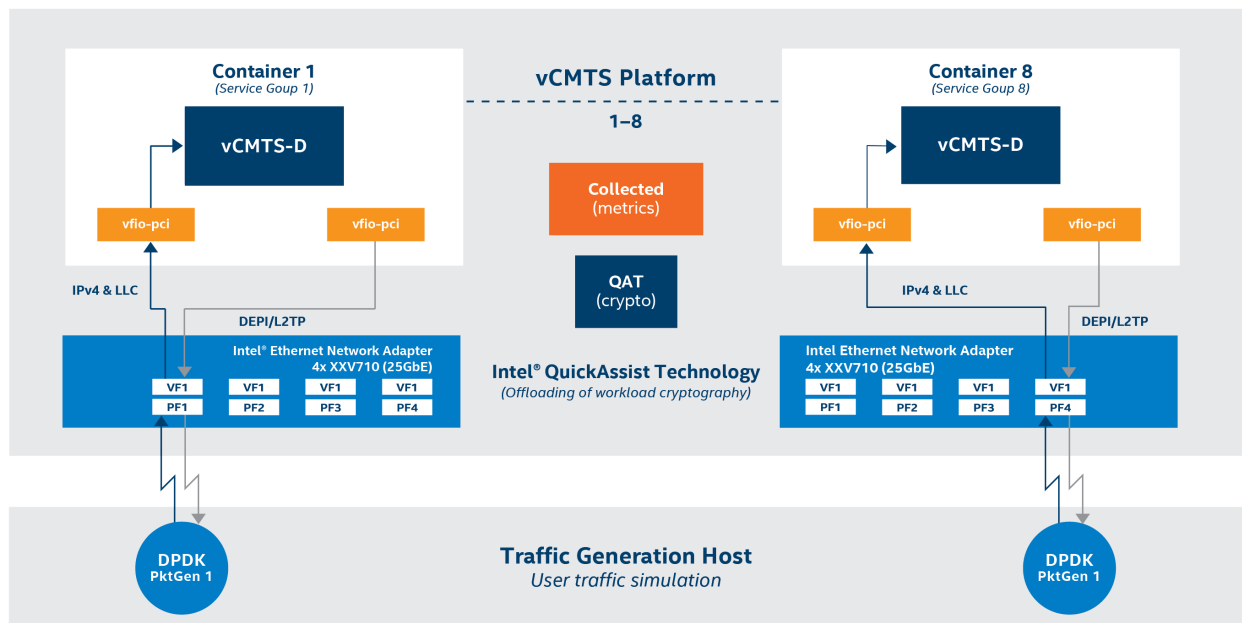


Figure 3 - Testbed setup detailing the network flows between the traffic generation host and several instances of the vCMTS data plane

Figure 3 shows the network flows from the traffic generation hosts to several vCMTS data plane instances. This setup leverages platform differentiating features, such as Intel Ethernet Network Adapters technology and Intel QAT.

[CollectD](#)* is used for gathering telemetry data from the infrastructure, the middleware, and VNF instances. [Intel® Resource Director Technology \(Intel® RDT\)](#) and [Intel® Performance Counter Monitor \(Intel® PCM\)](#) allow for detailed instrumentation and management of the platform. These technologies allow gathering of detailed hardware-level telemetry data, such as statistics about the usage of the individual cores and caches within the processors. The open source metric monitoring tool [Prometheus](#)* is used to consume the output of CollectD, aggregate the data, and perform derivation of our custom set of additional metrics.

Telemetry data is a key input for understanding the behavior of the VNF instances, as well as the infrastructure resources.⁶ Metrics have various levels of granularity (for example, per socket, per core, or per thread). As an example, high levels of utilization of the CPU, the associated caches, and the number of Instructions per Cycle (IPC) can indicate when computation itself has become a limiting factor.^{7,8} Furthermore, metrics on the usage of memory and I/O provide insights into the VNF's impact on other subsystems of the platform. Next to this, a landscape makes the relationship between entities in the system explicit. For example, the metadata describing which VNF instances were associated with which cores in the system can be used by the data pipeline in the analytics component.

Preselection of key useful metrics allows for more lightweight machine learned models to run efficiently, facilitating the foreground analytics flow. Detailed metrics—or statistical information on the same—are stored to the Information Core for later processing and future model generation. Derived metrics can also be generated, which provides a set of more complex data that captures less obvious component interactions. For example, the interplay between and impact on performance of the L1 and L2 caches can be rolled up into a single, more useful, metric.

Figure 4 shows a subset of metrics as gathered by the telemetry system, which can potentially be used to train machine learned models. The initial focus for this paper is on metrics such as the energy usage of the CPU (package energy) and memory (RAM energy). The IPC metric indicates how many instructions per cycle the CPU is performing and is a good indicator of overall system performance. The closer the IPC metric is to the theoretical maximum, the more efficient the system is performing (similar to miles per gallon). A drastic reduction of the IPC metric would indicate that the CPU is most likely waiting on memory I/O and cannot process faster.




Figure 4 - Sample metrics as collected by the telemetry system

Energy usage of the sockets, as well as memory, are shown in relation to metrics about the usage of the CPU (Instructions Per Cycle and Retiring)

Even for a single compute host NFVI about 200 different metrics can be collected from various subsystems. A small time window is shown in the diagram with a few data points; in a production system, large volumes of data points can be expected. It is important to be able to select those metrics which are key for the scenario at hand. This can be partly automated by a feature selection process that determines those metrics which are significant enough or show a high enough probability to be useful to predict a certain outcome.

Initial Results Demonstrate Key Areas for Optimization

To investigate the potential benefits of using machine learning to optimize initial placement of VNFs on a single host NFVI, experiments were run that compared the behavior of two vCMTS instances. In one case, the instances were pinned to two separate cores and can exclusively make use of them, while in the other case, two instances share a single core. The latter case is a simple example of dense packing and ultimately of a better return on investment as the processor cores are more optimally utilized. Note that

although we seek to achieve dense packing of VNF instances, it must not come at the cost of a reduction in the performance of the vCMTS instances and hence an impact on the end user. Being able to apply machine learning optimization by determining the most beneficial number of VNF instances sharing a resource is hence the driving goal.

As a baseline measurement of the two scenarios, we investigated performance by configuring the system to have 500 users subscribed to each VNF instance, exclusively using AES encryption and one OFDM channel. The total throughput for each VNF instance was 5 Gbps.

Behavioral differences were seen when comparing two different usage scenarios. As Figure 5 shows, the average package energy measured by the system was ~161 Joules/s when two VNF instances were pinned to individual cores (exclusive), while when pinned to a single core (shared) the overall power usage was reduced by 32 percent. The IPC metric on the other hand shows less of a reduction, which demonstrates that the system is not yet saturated. This means that although the vCMTS VNF instances are more tightly packaged, the system is not under-provisioned. This is also confirmed through the experiments (regardless of whether the VNF instances are sharing CPU cores or not) which show a similar throughput.




Figure 5 - Showcases variance in energy measurements for two key metrics when two vCMTS VNF instances either share a CPU core (shared) or are pinned to separate cores (exclusive)

Conclusion

Tighter packaging of VNF instances on shared resources of an NFVI is beneficial to make the best use of the available capacity. However, it is crucial to understand the behavior of VNF instances in such scenarios to ensure that their service level performance is not impacted. Such insights on what the optimal configuration is can be derived using machine learning insights. These insights can then be embedded in the MANO decision-making process while allowing for an autonomous adaptive system. This is especially necessary if the overall system needs to adapt to context changes, such as changes in user activity, channel configurations, and subscriber density.

Overall, the benefits of using machine learning approaches for derivation of insights that can be used in a MANO stack are the following:

1. An adaptive system that can automatically adapt to the addition of new VNF types or customer behaviors, reducing the need to manually change set placement policies.
2. Higher quality of service, due to multiple layers of anomaly detection and fault checks in service behavior data.
3. Reduction of OpEx through tighter packaging of VNFs, ensuring power savings when VNFs share infrastructure resources.
4. Coordination of multiple VNF types, enabling a vCMTS to not only execute on the same server infrastructure as other VNFs, such as augmented reality (AR), IoT, etc., but to leverage the differences between these VNFs to improve collective execution and save on OpEx across multiple fields.

vCMTS combined with insight-driven MANO allows for improved utilization of the resources, while offering consistent levels of Quality of Service (QoS). As a result, cable service providers can achieve new levels of optimization of the NFVI in order to deliver better service to their customers, while reducing TCO.

Abbreviations

CMTS	Cable Modem Termination System
CPU	Central Processing Unit
IPC	Instructions per Cycle
I/O	Input/Output
KPI	Key Performance Indicator
MANO	Management and Orchestration
ML	Machine Learning
NFV	Network Functions Virtualization Infrastructure
NFVI	Network Functions Virtualization Infrastructure
NFVO	NFV Orchestrator
SDI	Software-Defined Infrastructure
SLA	Service Level Agreements
TCO	Total Cost of Ownership
vCMTS	Virtualized Cable Modem Termination System
VIM	Virtualized Infrastructure Manager
VNF	Virtualized Network Function
VNFM	VNF Manager

Bibliography & References

1. S. Krishnapura, et al., *How Software-Defined Infrastructure Is Evolving at Intel*, Intel Corporation, 2015, <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/how-software-defined-infrastructure-is-evolving-at-intel-paper.pdf>
2. T. Metsch et al., *Apex Lake: A Framework for Enabling Smart Orchestration*, Middleware Industry '15, Proceedings of the Industrial Track of the 16th International Middleware Conference, <https://dl.acm.org/citation.cfm?id=2830016>
3. R. Khanna R. et al., *Autonomic Characterization of Workloads Using Workload Fingerprinting*, <https://ieeexplore.ieee.org/document/7015482/>.
4. Maximizing the Performance of DOCSIS 3.0/3.1 Processing on Intel® Xeon® Processors <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/vcmts-docsis-architecture-study.pdf>
5. H. Wippel, *DPDK-based implementation of application-tailored networks on end user nodes*, 2014 International Conference and Workshop on the Network of the Future (NOF), Paris, 2014, pp. 1-5., <https://ieeexplore.ieee.org/document/7119762/>
6. A. Yasin, *A Top-Down Method for Performance Analysis and Counters Architecture*, 2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), <https://ieeexplore.ieee.org/document/6844459/>
7. A. Herdrich et al., *Cache QoS: From Concept to Reality in the Intel® Xeon® Processor E5-2600 v3 Product Family*, <https://ieeexplore.ieee.org/document/7446102/>
8. M. Schwarzkopf, *Operating System Support for Warehouse-Scale Computing*, PhD thesis, <https://people.csail.mit.edu/malte/pub/dissertations/phd-final.pdf>

Estimated results reported above may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information about benchmarks and performance test results, go to www.intel.com/benchmarks.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© Intel Corporation

Trends on the Role of Internet of Things Technology in Senior Care

Market Trends, Technology Drivers, Initial Learnings

A Technical Paper prepared for SCTE•ISBE by

Bruce McLeod
Executive Director, Engineering
Cox Communications

Table of Contents

Title	Page Number
Table of Contents	2
Abstract	3
Content	3
1. Demographic Change and Rapid Demand growth for Senior services in the United States.....	3
2. IoT Technologies with potential to provide or enhance caregiving services to the elderly	10
3. IoT Architecture	12
4. Current and Near Term IoT Opportunities	12
Conclusion.....	13

List of Figures

Title	Page Number
Figure 1 – Projected Number of Children and Older Adults	4
Figure 2 – Multigenerational Households	5
Figure 3 – Multigenerational Households and Youth	6
Figure 4 – Workspans (US Bureau of Labor Statistics)	7
Figure 5 - US Census Bureau Statistical Brief, 1995, Revised Oct 31, 2011	8
Figure 6 - Average Costs in Georgia	9
Figure 7 – Unpredictability of Long-Term Care	10
Figure 8 – The 4 Stage IoT Solutions Architecture	12

List of Tables

Title	Page Number
Table 1 – Senior Living Mode	10

Abstract

Demographic shifts in the United States project a rapidly aging population. The number of citizens aged 65 and above will exceed the number of children by 2030. Longer life spans, smaller families, relocation of working age children away from aging parents have all contributed to a burgeoning opportunity for targeted technology to fill growing gaps in care. When elderly become unable to live fully independently and in-home care options are limited, what options exist? Assisted Living is a growth business in response to this phenomenon. Remote diagnostic health management and telepresence are largely absent. Residents generally lack technology savvy needed to avail themselves of such quality of life enhancing technology. This discussion explains specific proof of concept and test use cases to demonstrate benefits of IoT, Robotics/Telepresence, and AI in this emerging market. It pinpoints where and how Service Providers can deliver service packages to a growing industry.

Content

1. Demographic Change and Rapid Demand growth for Senior services in the United States


As of Aug 10, 2018, the total population of the United States as reported by the census Bureau was 328,325,857. Population change occurs at the rate of one birth every 8 seconds, one death every 12 seconds, and net gain of one citizen every 12 seconds (immigration + births) – deaths. Projections from the Census Bureau portray a significant aging of the overall population, with the Adults 65+ category growing from 15.2% in 2016 to 23.5% in 2060. Meanwhile the percentage of children <18 is projected to decline from 22.8% to 19.8%. The number of Adults 65+ is projected to exceed the number of children in the country by 2033.



An Aging Nation

Projected Number of Children
and Older Adults

For the First Time in U.S. History Older Adults Are
Projected to Outnumber Children by 2035



United States[™]
Census
Bureau

U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
[census.gov](https://www.census.gov)


Source: National Population
Projections, 2017
www.census.gov/programs-surveys/popproj.html

Figure 1 – Projected Number of Children and Older Adultsⁱ

Another significant aspect of the demographic shift is the ratio of working age adults per dependent Youth or Adult 65+ able to fund subsidies for caregiver services. In 2009, the number of people living in multi-generational households was 17%, or 51.5 million, where aging parents received some help with ADL's (Activities of Daily Living), accounting for 21% of Elderly needing care. In 2018, that statistic has increased to 20%, or 64 million people.ⁱⁱ While much of this trend can be attributed to higher overall cost for housing, it also reflects the larger economic trend of labor cost reduction and automation that is reshaping the economy. In the Senior Care market, opportunities to improve service and reduce cost are just now coming under scrutiny.

One-in-five Americans live in a multigenerational household

% of population in multigenerational households



Note: Multigenerational households include at least two adult generations or grandparents and grandchildren younger than 25.
Source: Pew Research Center analysis of 1950-2000 decennial censuses and 2006-2016 American Community Survey (IPUMS).


PEW RESEARCH CENTER

Figure 2 – Multigenerational Households

Of note is the shift in this statistic is that the ratio of parents living with children has greatly declined since 1940 while children and adult children living in their parents' home has increased as shown in the chart below. (Fig 3)

Young adults most likely age group to live in a multigenerational household

% of population in multigenerational households



Note: Multigenerational households include at least two adult generations or grandparents and grandchildren younger than 25.

Source: Pew Research Center analysis of 1940-2000 decennial censuses and 2006-2016 American Community Survey (IPUMS).

PEW RESEARCH CENTER

Figure 3 – Multigenerational Households and Youth

This oldest segment of the population, as well as most of the Baby Boom generation, are expected to pursue a lifestyle of aging in place. This is one example of where IoT technologies can provide convenience, cost reduction, reduced stress on caregivers, while also enhancing the quality of life and well-being of Seniors.

Demographic shifts (figure 4) over the next two decades will exacerbate deficiencies in availability and affordability of Senior care services for the Baby Boom generation. In 2030, they will begin to turn age 85, entering the point in their life where assistance with one or more ADL's will likely be needed.




Figure 4 – Workspans (US Bureau of Labor Statistics)

Seniors unable to live with caregiving relatives must rely on services when confronted with the need for assistance with Activities of Daily Living. Services may then be obtained from municipalities, For-profit and Non-profit service providers, local and remotely located relatives, and individual volunteers. Government assistance is reserved for individuals that qualify for Medicaid. The Institute on Aging, a Non-Profit Senior Care organization in San Francisco, characterizes this segment of the Senior population with the following facts:

Living to 85+

- In 1900, only 100,000 Americans lived to be 85+.
- By 2010, that number had grown to 5.5 million. This is the fastest growing age group of elders.
- By 2050, the 85+ age group will reach 19 million—24 percent of older adults and five percent of the total population.
- Some researchers say the 85+ group will grow even faster than this, because death rates at older ages will decline more rapidly than the U.S. Census Bureau predicts.

Living Alone

- Of the older adults who were living outside nursing homes or hospitals in 2010, nearly one third (11.3 million) lived alone.
- Older women are twice as likely as older men to live alone (37 percent and 19 percent, respectively). In 2010, 72 percent of older men lived with a spouse, only 42 percent of older women did.
- Living arrangements differ by race and ethnicity. Older non-Hispanic White women and Black women are more likely than women of other races to live alone (39 percent each, compared with about 21 percent of older Asian women and 23 percent of older Hispanic women).
- The likelihood of living alone increases with age. Among women age 75+, almost half (47 percent) lived alone in 2010..ⁱⁱⁱ




Figure 5 - US Census Bureau Statistical Brief, 1995, Revised Oct 31, 2011

As the US population shifts to more elderly citizens, the role of Internet of Things technology will only increase as we grapple with the cost and burden of care. Many WWII and Silent Generation Seniors in the upper age brackets of 85 and beyond, express strong desire to retain independence, age in place for as long as possible, and do not wish to burden their children for long term care. Yet, long term care costs are extreme, ranging from \$3000-\$4000/per month for Assisted Living residence with only lodging, meals, activities, excursions, and housecleaning included. Any additional service such as laundry, medication management, or personal care to assist with ADL's comes at a premium that can quickly add hundreds or thousands of dollars to the monthly cost. At its most extreme, Memory care for residents with Alzheimer's or other cognitive issues can cost up to \$10,000 per month, depending on region. Medicare

does not cover any form of assisted (Custodial, or ADL) living. Seniors not prepared with Long Term Care Insurance must deplete their entire nest egg before becoming qualified to obtain Medicaid subsidies for long term care. A common scenario is for a widow to sell her home and use the proceeds to pay Assisted Living facility costs until depleted. The Long Term Care Insurance market has recently collapsed to only 15 providers due to the burden of current claims on underwriters. In 2018, LTCI is difficult to obtain and expensive with premiums that have increased dramatically since 2012. A snapshot of 2016 Care costs from Mutual of Omaha: (Figures 5 and 6)

Home Health Care	Per Hour
Home Health Aide	\$20.54
Licensed Practical Nurse (LPN)	\$57.08
Registered Nurse	\$68.14
Assisted Living Facility	Per Month
One Bedroom Unit	\$3,656.54
Nursing Home	Per Day
Semi-Private Room	\$200.46
Private Room	\$216.63


Figure 6 - Average Costs in Georgia

Annual cost for home health aide is based on services received 44 hours per week, 52 weeks per year. Annual cost for LPN and RN can be determined based on the actual number of hours services are required. Source: Mutual of Omaha Insurance Company's Cost-of-Care Study conducted by LTCG, 2015; released April 2016
Source available upon request.

The Unpredictability of Long-Term Care

HOW LONG ...


The duration of paid care among 65-year-olds who will need it someday varies widely, but for many it is under one year.



PERCENTAGES EXCEED 100% BECAUSE OF ROUNDING. SOURCE: DEPARTMENT OF HEALTH AND HUMAN SERVICES

HOW MUCH ...


The median annual cost of nursing home care depends on your state.



PRICES ARE FOR A SEMIPRIVATE ROOM. THE MEDIAN ANNUAL COST IN ALASKA IS \$292,000. SOURCE: GENWORTH

YOUR COSTS

One in four people now age 65 will face over \$50,000 in life-time out-of-pocket long-term care expenditures.



SOURCE: DEPARTMENT OF HEALTH AND HUMAN SERVICES

Figure 7 – Unpredictability of Long-Term Care

2. IoT Technologies with potential to provide or enhance caregiving services to the elderly

Senior living modes and opportunities for available and rapidly emerging IoT technology to provide or enhance services needed by Senior groups, focused on the 85+ age group:

Table 1 – Senior Living Mode

	Senior Living Mode						
IADL, ADL Need	Alone	Senior Village	Minor In-Home	Significant In-Home	Constant In-Home	Assisted Living +	Skilled Nursing
Monitoring	STW	STW	STW	STEM	STEM	STWE	STWE
Alert- Personal	STW	STW	ST	STW	STW	STW	M
Alert- Physical	STW	ST	ST	ST	M	STW	M
Communication	ST	ST	ST	ST	N/A	ST	S, T
Medications	ERW	ERW	ERW	ERW	ERM	ERW	STWE
Transportation	T	T	T	T	N/A	N/A	X
Visitation	TR	TR	TR	TR	TR	TR	X

Cooking	R	T R	T R	N/A	N/A	N/A	X
Housework	E R	E R	R	N/A	N/A	N/A	N/A
Personal Finance	T	T R	T R	N/A	X	T R	X
Home Maintenance	E	N/A	E	N/A	N/A	N/A	N/A
Hygiene	E W	E W	E W	E W	X	X	X
Eating	E W	E W	E W	E W	X	X	X
Dressing	R	R	R	E W	X	X	X
Mobility	R	R	R	R	X	X	X
Continence	W	W	W	E W	X	X	X
Cognitive	N/A	R T	R T	N/A	N/A	R T	N/A

S = Security System, W = Wearable Sensor, E = Environmental Sensor, R = Robot, T = Telepresence/com

M= Medical, X = Incompatible with Living Mode, N/A= Not Applicable

Definitions:

Security System/Smart home – Includes standard sensors (cameras, motion detectors, smart home devices) for home physical security where access is shared with Care persons or loved ones. Two-way voice capable.

Wearable Sensor – Includes ID and location devices, call buttons, through-skin sensors/dispensers for measuring body functions, accelerometers.

Environmental/Health Sensors – Includes all fixed detectors not considered ordinary Security/Smart home devices. Includes air quality detectors, lavatory fixtures with biometric and sample processing, as well as the standard functions of temperature, humidity, and air pressure.

Robot – A programmable machine that performs a task. Examples include medication dispensers, meal delivery, cleaning, companion or control device.

Telepresence – Two-way video and mobile screen platforms to interact with Seniors and allow remote visitation. Examples include Beam and Double.

Medical – Specific medical devices that are currently out of scope for integration with commercially available IoT components but includes Tele-Health service systems such as Trapollo.

3. IoT Architecture

IoT systems vary widely in composition, capability, and purpose. A general model proposed by J. R. Fuller of Online source Tech Beacon describes stages of IoT as shown in Figure 7:

The 4 Stage IoT Solutions Architecture

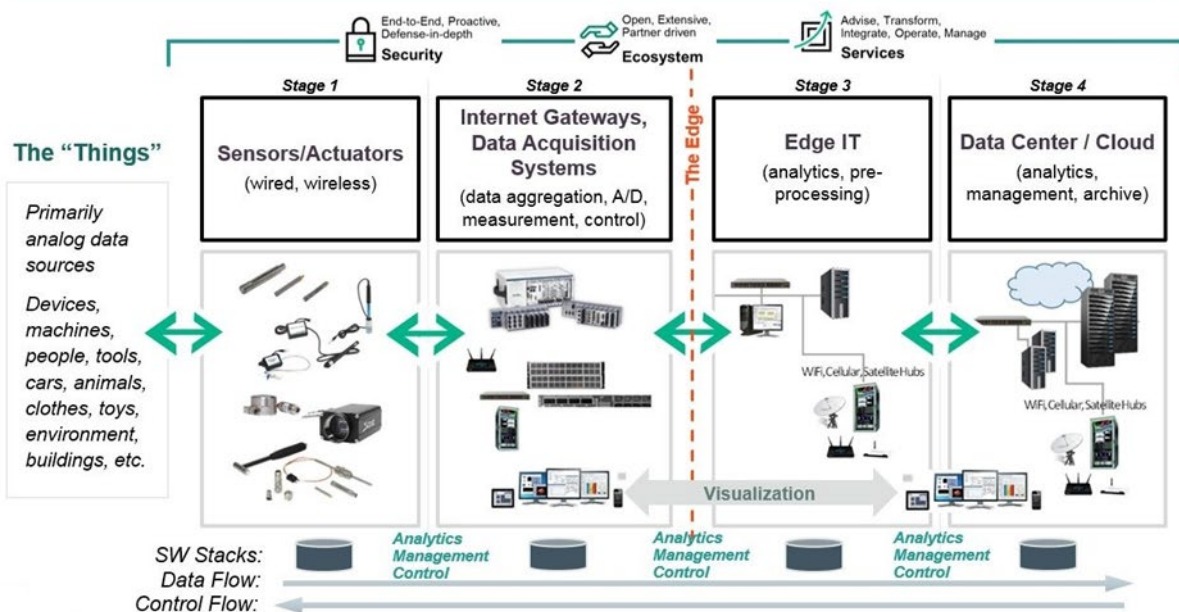


Figure 8 – The 4 Stage IoT Solutions Architecture

The scope of this paper does not include detailed comparisons of all technical approaches but a general assumption is that all sensors and data acquisition devices can be accessed, controlled, and managed for their purpose.

4. Current and Near Term IoT Opportunities

Historically, Aging in Place accommodations have consisted primarily of physical changes to the residence, making it as barrier free as possible or reasonable. Aging in Place is a popular choice among Seniors and creating new demand. Stair-lifts or elevators may be needed in a multi-story home. Ramps at door sills or other obstacles, lavatory adjustments, and transportation arrangements are those basic elements.

Traditional security systems and increasingly available wireless integrated services and user-installed that include modern features such as remote management of many smart devices. For Seniors, door locks, doorbell cameras, and additional detection can form a basic IoT starter system. For caregivers, such a system helps in determining if the Senior is moving about as expected. For a family care giver, these two items provide the means to safely leave the home, yet stay in touch with their Senior and this can be a great relief. It can lessen the risk of a third party caregiver misbehaving with the Senior or their property. A best practice of using IoT technology with Seniors is that you cannot expect someone in their mid-80's or above to learn how to use a new technology. For most Seniors of advanced age, smart device management may be completely confusing.

An alternative interface for Seniors is voice control, such as an Amazon Echo, that can tie all services into a central hub and simplify the user interface. Service Providers are deploying packages, such as Cox Communications Homelife, to simplify the user interaction with the system. Cox is currently researching interfaces that Seniors are more receptive to. One such example is the Social Robot Platform, Jibo. Cox trialed this approach for “Homelife Care” in Oklahoma City this year with strong user ratings and feedback.

IoT technology as a whole is still maturing while gaining momentum through rapid consumer and industrial automation adoption. Operating systems vary widely and consumer systems require some administration, so if you plan to provide your Senior loved one with a system, current choices are either well-integrated but expensive systems or DIY open standards components. Provider based services usually come with a monitoring service monthly fee.

In Assisted Living Facilities, at present, the use of IoT technology in facilities is generally not present. Other than an in-person visitation, it can be difficult to know if your loved one is receiving adequate care and value for their expense. At Cox, we are intending to test the use of mobile telepresence, using the Double Robotics device to enable remote visitation of residents. We will also investigate the utility of camera-based monitoring. While the Senior Care industry is still undergoing massive growth to meet demand, it remains daunting to see how the industry can fulfill the demand for service without the ability to lower cost and improve service. The majority of Seniors needing care will not arrive with several hundred thousand dollars to finance their stay or several children to provide home care. Americans seeking to care for their aging parents will expect, and will gravitate to, operators and services using IoT technology as an integral part of the service offering. Senior Care professionals that do not embrace IoT technology as table stakes for the future can expect to suffer the market consequence of ignoring this technology mega-trend.

Conclusion

Smart Home and Industrial IoT technologies are currently capable of providing beneficial features for the growing need for Senior Care. It’s not yet clear which solutions will ultimately emerge as the most successful but a clear opportunity for MSO growth is IoT Technology for Senior Care.

ⁱ US. Census Bureau

ⁱⁱ D’Vera Cohen, Jeffrey Passel, Pew Research Fact Tank, updated April 5, 2018

ⁱⁱⁱ Institute on Aging, <https://www.ioaging.org/aging-in-america#womanliving>

Using AI to Improve the Customer Experience

A Virtual Assistant Chatbot

A Technical Paper prepared for SCTE•ISBE by

Bernard Burg,
Fan Liu,
Abel Villca Roque,
Sunil Srinivasa,
Ryan March,
Comcast

1050 Enterprise way, Sunnyvale CA 94089
+1 408 900 85 75
bernard_burg@comcast.com

Tianwen Chen,
Comcast
1110 Vermont Av NW, Washington DC 20005

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
1. Chatbot Architecture	3
2. Introduction to Machine Learning	4
2.1. Supervised Learning	4
2.2. Unsupervised Learning	5
2.3. Reinforcement Learning	5
3. Domain AI.....	5
3.1. XRE React: What to do after an XRE Error?	6
3.2. XRE Predict: Who needs our help now?	7
3.2.1. Call Predictions	7
3.2.2. Silent Sufferers.....	8
3.3. XRE Prevent: Can we solve the problems before they surface?	9
4. Decision Engine (DE).....	10
4.1. Overview	10
4.2. Multi-Armed Bandit (MAB) Algorithm Overview	11
4.3. Linear Contextual MAB: Introduction	12
4.4. Linear Contextual MAB on the XA App: Experimental Setup	12
4.5. Linear Contextual MAB on the XA App: Policy Evaluation Results	13
Conclusion.....	14
Abbreviations	15
Bibliography & References.....	15

List of Figures

Title	Page Number
Figure 1: Overview of ChatBot.....	3
Figure 2: System Refresh, Restart and DoNothing models.....	6
Figure 3: System Refresh and Restart Deployment Models.....	6
Figure 4: Precision, Recall	7
Figure 5: Call Predictions Model Flow and Results	8
Figure 6: Initial Clusters	9
Figure 7: Meta Clusters with callers and silent sufferers	9
Figure 8: Root Cause Analysis and Asynchronous Sampling	10
Figure 9: Decision Engine Data Flow.....	11
Figure 10: Linear MAB Machine Learning Flow.....	12
Figure 11: Linear MAB Policy Evaluation: Net Rewards.....	14
Figure 12: Linear MAB Policy Evaluation: Positive (left) and Negative (right) Rewards	14

Introduction

AI and Machine Learning (ML) are becoming pervasive, allowing new applications for chatbots. This paves the way for operational transformations in the field of the cable telecommunication industry, where chatbots will soon be able to field some customer contacts to solve their issues in real-time -- thereby reducing waiting queues while enhancing service quality, informed by consistent answers and considering all available network information in real time.

This article describes the steps service providers can use to build such a chatbot. It and drills into the AI/ML elements in charge of performing installation diagnostics, predicting the behavior of the systems and users, and pinpointing the root causes critical to fixing issues before they even become visible to customers. These highly specialized AI/ML algorithms feed their propositions into a decision engine (DE), so as to understand a much larger context, and apply game theory to make optimal choices for the customer.

1. Chatbot Architecture

A chatbot is a computer program that provides an interface that allows customers to interact with machine learning systems via auditory or textual methods. Through Chatbots, customers can report service issues and get personalized answers or help, using the best knowledge of our machine learning systems. Further, Chatbots can obviate the need for customers to navigate websites or call for support. An overview of the Chatbot is provided in Figure 1.

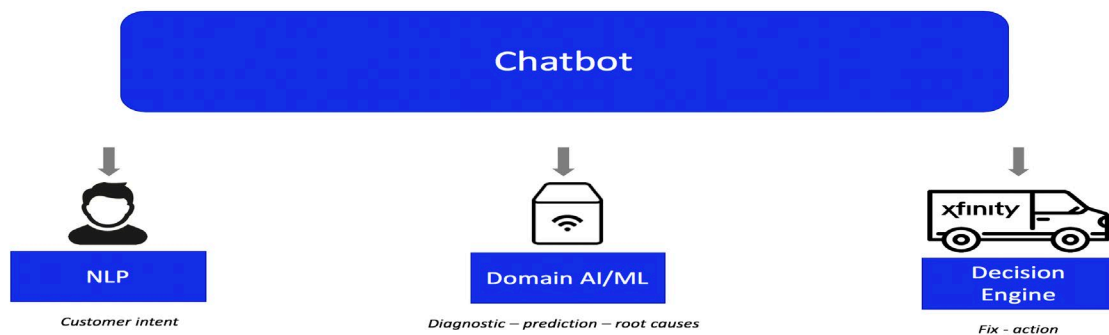


Figure 1: Overview of ChatBot

A chatbot dialog unfolds as follows: A customer interacts, either by typing a query into a mobile device, or by directly talking into a microphone which transmits the voice signal to a speech-to-text algorithm to produce a script of the query. This query is analyzed by a Natural Language Processing (NLP) module that extracts the intent of the customer. In some cases, these intents are implicit in the query and need to be translated by NLP into an explicit statement, often formulated in technical jargon. For example, a query like “it takes very long to download images” might be translated into a more generic intent like “internet speed slow,” known in the domain knowledge.

This intent is forwarded to the Machine Learning/AI domain for investigation. These ML/AI algorithms have access to a wealth of information, such as high speed internet data, in-home

WiFi data, and many kinds of system and video errors including RDK and XRE. Ideally, all of this information is stored in a data lake, with nationwide coverage and a history duration spanning over several months to years. These data lakes allow ML algorithms to learn the customer behavior or performance of our current networks, and to extract salient events, called “features,” which in turn enable algorithms to assess the situation of a customer through a kind of fingerprinting.

One method used by ML to fix issues is actually to identify “error fingerprints” from the past and learn how they were fixed. All things being equal, applying the same fix should solve the same issue with high probability. Additionally, ML algorithms can predict future calls or issues that customers are experiencing based on similar “fingerprints” seen in the past. Domain AI/ML finally presents its diagnostics, predictions and recommendations of actions to the DE (decision engine.)

The DE gets recommendations from domain AI/ML as input features. These recommendations might be incoherent, in competition with or even in contradiction with each other, given that they’re created by highly specialized ML algorithms -- each of which has a deep but narrow understanding of the world. The DE aims to understand a larger context to select the right decision.

In its simplest implementation, a DE can be implemented by a rules engine that can choose the best solution for a problem. However, these rules engines are static solutions, and would constantly need re-evaluations over time to make sure they still are optimal. A superior implementation of the decision engine uses a recommendation system to take into consideration larger context domains, along with domain knowledge, to continuously choose the optimal action and fix the issue reported by the customer. Our particular implementation uses what’s called a “multi-arm bandit algorithm.”

2. Introduction to Machine Learning

In this section, we define the machine learning algorithms used in this paper. Machine learning algorithms try to learn the underlying patterns in data. Depending on the ways that ML algorithms learn patterns, they are categorized into three families: supervised learning, unsupervised learning and reinforcement learning.

2.1. Supervised Learning

Supervised learning is a guided learning approach, meaning that it learns the pattern of the data that could optimally predict the provided labels. In our system, we use classifiers to discern similarities and differences in data, and to assign them to categories based on similarity. Examples of such classification algorithms include identifying objects in a video frame, identifying the underlying sentiment in a customer service message, or associating a log message from a set-top with a specific error class. The technologies powering classifiers range from the simple -- decision trees and random forest algorithms-- to the very complex, such as deep neural networks. The choice amongst available technologies is typically related to the level of complexity and the number of features in the underlying data. Image classification, for instance,

has been shown to benefit greatly by neural networks and its derivative technologies, such as convolutional neural nets.

2.2. Unsupervised Learning

In contrast to supervised learning, unsupervised learning aims to determine underlying patterns without any sort of guidance or provided labels. In our system, we use clustering algorithms to group similar data into clusters, and the clusters do not have any preset labels. They are typically used to understand the behavior of data or to look for any significant deviations in data. For example, clustering algorithms could be used to look at anonymized smart home data to build user profiles, such as early risers, late risers etcetera, based on common behaviors. Clustering algorithms range from the simple, such as k-means clustering, to the complex, like agglomerative hierarchical clustering.

2.3. Reinforcement Learning

Reinforcement learning is a type of ML inspired by behavioral psychology, and concerned with how software agents take actions in an environment so as to maximize some notion of cumulative reward. Reinforcement learning is studied in many disciplines, including control theory, game theory, simulation-based optimization, and more.

Reinforcement Learning is also directly applicable to the operations improvement in the cable telecommunication industry, as our networks collectively qualify as complex systems to be optimized, despite the impossibility of building a mathematical model of such a system. Early implementations of reinforcement learning have been studied in optimal control theory to tackle this very issue.

Reinforcement learning can also be used in gaming theory to select the best action to perform while optimizing gains. For cable operators, this translates into selecting the best actions to perform on an account, while minimizing the disruptions of service. Such an approach is used in our Decision Engine.

3. Domain AI

Three domain-specific AI modules are under development:

- High Speed Data
- WiFi
- Video domain models including XRE, which stands for Cross-platform Runtime Environment, and is a platform-independent protocol for distributed applications.

This paper describes the example of XRE. Three problems are assessed in this study. First, how to react after an XRE error or a sequence of XRE appears on a device. Second, how can consequences of such XRE errors or sequences be predicted, in terms of how important are they to our customers, and how much they affect our service? Third, how can these errors be prevented from happening in the future, and can problems be solved before they even surface?

3.1. XRE React: What to do after an XRE Error?

This first study investigates three actions that can be performed after observation of an XRE error or a sequence of XRE errors for a single device during a time window of an hour. An example of an XRE sequence for user a is $a_1 = \{\text{XRE-03059}, \text{XRE-03059}, \text{XRE-10007}\}$. Whereas at the same time, user b might receive the sequence $b_1 = \{\text{XRE-10007}\}$, and user c might not receive any error during the time window, hence $c_1 = \emptyset$. All of these sequences represent fingerprints of the users' statuses, and they can be used to perform predictions.

Three independent ML models are used to predict the effectiveness of three actions that may be performed to fix the XRE errors (see Figure 2):

- **System Refresh:** account refresh + cable card refresh + firmware reset + reboot
- **System Restart:** simple box reboot
- **Do Nothing:** estimates the XRE error attrition over time, and recommends not to act

The goal of each of these actions is to eradicate the XRE errors. Success is claimed if there are no errors in a duration of one hour after the action has been performed.

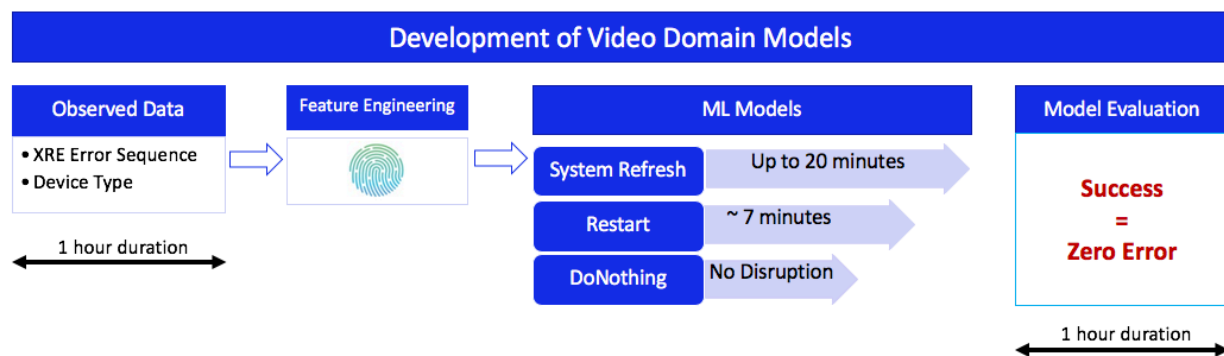


Figure 2: System Refresh, Restart and DoNothing models

For each of these actions, a supervised ML algorithm learns a classifier to organize samples into two classes: Success or failure. During the prediction phase, new examples that have never been seen are presented to the classifier, which will predict the probability of success of the action (see Figure 3).



Figure 3: System Refresh and Restart Deployment Models

Current ML results allow the predicted performance of the system refresh model with 70% of precision and 87% of recall, meaning that 70% of the recommended system refreshes

successfully eradicate the XRE errors, and 87% of the effective system refreshes are captured by the ML model. See a visual representation of precision and recall in Figure 4: Precision, Recall. Similarly, the ML results predict the performance of a system restart model with a precision of 84% and a recall of 99%.

The outputs of these models are stored in a repository for assessing their validities when the actions, system refresh or restart, are performed. These actions can either be performed by the customers, or by reinforcement learning models executing the ML recommendation directly. The reinforcement learning models can check the validities of the actions immediately for self-calibration. In essence, reinforcement learning is comparable to a feedback loop in control systems.




Figure 4: Precision, Recall

3.2. XRE Predict: Who needs our help now?

Comcast measures customer satisfaction by using the Net Promoter Score (NPS). NPS scores send a very clear message regarding the number of interactions required to fix an issue. Customers are fairly happy when systems work flawlessly or get repaired before having to call. First call resolutions are generally well accepted, but mark a dip in NPS. Given the impact of repeat interactions on NPS scores, it makes much sense to study the severity of the XRE fingerprints in aiming to predict and later preempt the customer calls.

3.2.1. Call Predictions

The call prediction scenario for ML is similar to the previous example, as it also uses supervised learning. In this case, ML learns to predict the calls based on XRE errors that occurred over 24 hours. The modeling of these errors is also more precise, as it records each of the errors into hourly buckets. In Figure 5, the studied device received {11 XRE-00021, 5 XRE-03056, 33 XRE-03059} at 15:00, {7 XRE-00021, 5 XRE-03056, 22 XRE-03059} at 16:00, etc. The goal of the ML is to differentiate the callers from the non-callers based on these XRE models, and, given a new sequence the ML algorithm never encountered before, predict if this customer will call.

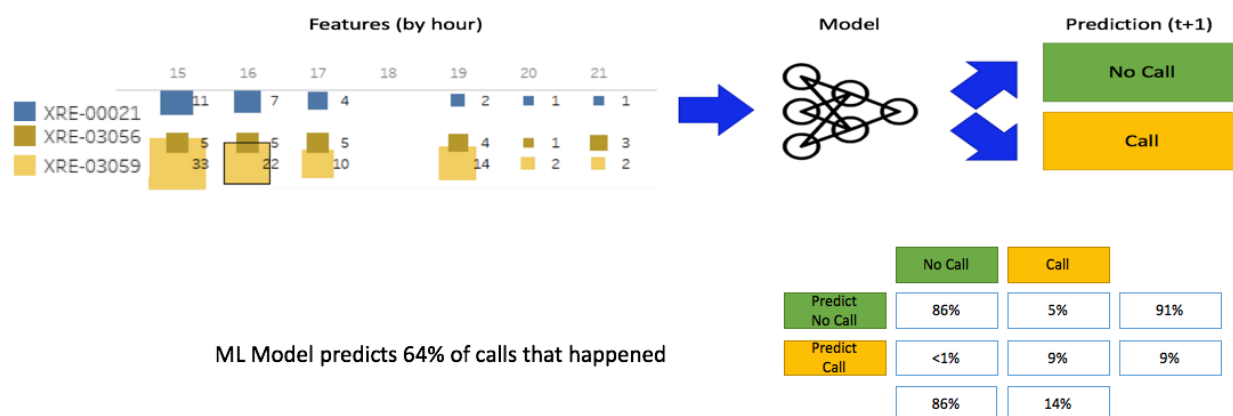


Figure 5: Call Predictions Model Flow and Results

The ML model predicted 64% of the calls that will happen, which is quite a valuable result because the learning sets are unbalanced -- meaning there are very few callers as compared to non-callers. But the most exploitable result of this study is the high accuracy in predicting a call. If the ML predicts a call, it is wrong in less than 1% of the cases. However, the algorithm only catches nine percent of the total number of calls. This result is interesting in the sense that we can predict some calls with a very high accuracy, and thus are poised to preempt these calls by fixing the underlying issues.

3.2.2. Silent Sufferers

Previous methods used for detecting quality of service issues are typically based on user feedback, gathered through chats and calls. Implicitly this means that no news means all is right, right? Wrong.

There exists a small category of customers who never use our communication channels, even while facing issues. Hence, they are called silent sufferers. These customers are undetectable with previous methods, they are unaccounted for in NPS, and user histories indicate that silent sufferers often unsubscribe without ever calling.

No supervised learning method can address silent sufferers. We present instead the use of an unsupervised ML method. The design principle is to introduce a distance measurement between users. The ML measures a fingerprint of a single XRE error over a 24 hour period of time. The unsupervised method used is the k-means algorithm, clustering users into classes of similarity according to this fingerprint measurement. Experiments show that 90% of the users are clustered into the main class (see Figure 6). The devices of these users show a nominal behavior where everything works fine, up to specification. Other classes are called outliers, which group similar users. Some of these classes prove to have a nominal behavior whereas others might face unusual behaviors, for example, degraded service. These classes are called the sufferer classes.




Figure 6: Initial Clusters

We defined several types of logic to combine the classes of sufferers across different XRE errors into meta-classes, so as to reinforce the predictive power of the weak predictor classes. Such results are described in Figure 7. Since the k-means algorithm performs clustering regardless of calls, the resulting meta-classes of sufferers contain, in general, both callers and silent sufferers. For each of these classes independently, an algorithm studies the causes mentioned by the callers to gather a better understanding of this class. The knowledge acquired through the callers is shared to the silent sufferers and allows action on their issues without having them to ever call.




Figure 7: Meta Clusters with callers and silent sufferers

3.3. XRE Prevent: Can we solve the problems before they surface?

The best way to prevent XRE errors is to understand their root causes and to tackle them before they can propagate and morph into customer-impacting XRE errors. It can be done with supervised learning but requires advanced feature engineering, because such models rely on a number of data sources of different natures/semantics. In particular see Figure 8:

- Hardware data: XRE telemetry, RDK telemetry, xFi data transfer telemetry
- Networking data: Spectra data in the cable, WOPR, PHT data
- Application data: XRE messages
- Behavioral data: Anonymized click and tune data generated by users.

In a first approach, these data were bucketed hourly, in the same manner as in 3.2.1. Because of the large number and the heterogeneity of the data sources, this approach is unlikely to work because of the different sampling rates in each of these data sources. Some of the data are sampled hourly while others are sampled at the micro-second level (e.g. the clicks and tunes). By using an hour as the common denominator, the millisecond-rate information of the clicks and tunes is washed away, and the clicks and tunes lose most of their precious information.

To overcome such difficulties, we developed our own asynchronous sampling methods into our models, inspired by [3]. Asynchronous sampling allows the capture of all the information without using buckets.

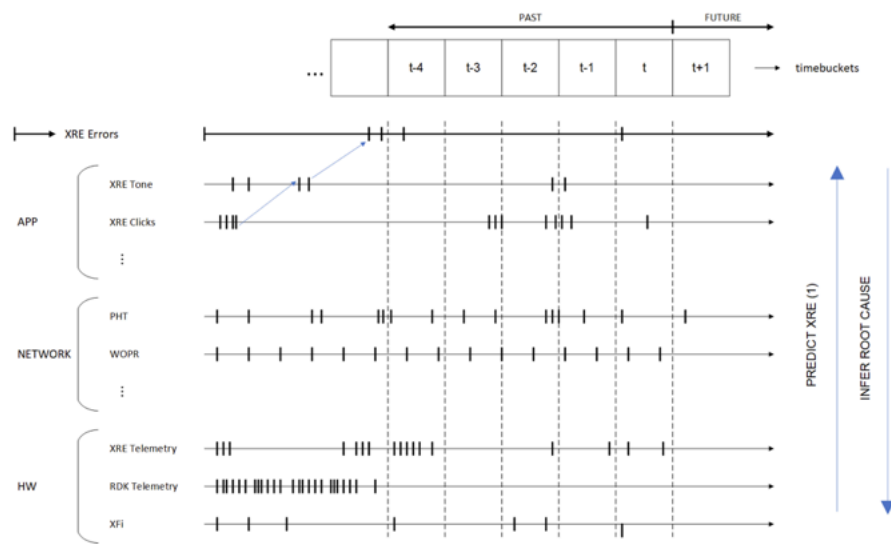


Figure 8: Root Cause Analysis and Asynchronous Sampling

4. Decision Engine (DE)

Each of the 'depth-first' domain AI algorithms provides recommendations based on their narrow but deep understanding of the building blocks of a cable operator's infrastructure. The decision engine operates in a 'breadth-first' fashion: It is in charge of collecting these recommendations, their contexts, the user's intent and making sense of it all, to execute the best actions that will improve the customer experience.

4.1. Overview

Figure 9 gives an overview of the decision engine workflow. The customer intent is acquired through the NLP module; the customer context is acquired through a virtual assistant guiding users in the resolution of their tasks; the customer's experience is provided by the domain AI algorithms presented earlier -- they analyze many sources of raw data, perform a diagnostic of the situation and propose actions to fix issues. Together, the customer's intent, context and experience form the overall context or state of the environment. Given this state, the job of the decision engine is to take the right action at the minimal cost. In the process, it might have to make trade-offs between the actions proposed by the domain AI, as some of these actions, though successful, might come at a cost in terms of user impact (see details in Figure 2). For

example, the restart of a set top box takes about 7 minutes, whereas the more powerful refresh operation fixes potentially more issues but can take (worst case) up to 20 minutes. In some cases, the domain AI even calculates the natural error attrition of the system and suggests not to perform any action, so as to not adversely impact the user.

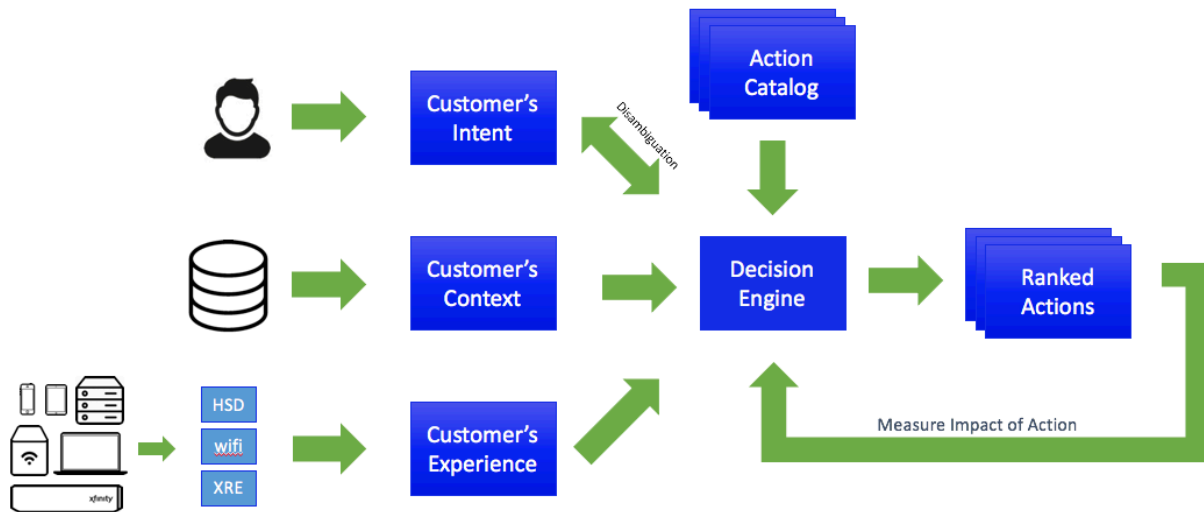



Figure 9: Decision Engine Data Flow

For a given context, the decision engine consults its action catalog, assesses the benefits and costs of the actions, and finds the best tradeoff to fulfill the customer's intents. In reinforcement learning terminology, the decision engine takes as input a state (or context), and predicts the best action(s) to take so as to minimize cost (or maximize reward).

4.2. Multi-Armed Bandit (MAB) Algorithm Overview

The operation of the decision engine is driven by the multi-armed bandit (MAB) algorithm [1], which is widely used for single-step decision-making problems. The name 'multi-armed bandit' references a gambler (generally, a bandit) at a casino with several arms, trying to play the right slot machines so as to optimize their winnings.

MAB is often depicted in cartoons as an octopus in front of several slot machines (each assumed to provide a specific payout) and trying to use its several arms to determine the slot machine that provides the highest payout.



source: Microsoft Research

An intrinsic characteristic of MAB algorithms is the '**exploration/exploitation**' tradeoff. In the context of the slot machine payout, this is what the terms mean:

- Exploitation: play the machine believed to have the highest payout
- Exploration: play untested machines to learn if there are higher-paying ones

Thus, the best long-term strategy may involve short-term sacrifices. The octopus will need to explore pulling the levers of several machines before it can start to exploit the best one.

4.3. Linear Contextual MAB: Introduction

Here, we used a 'linear' version of the MAB algorithm, which is depicted in Figure 10. Given a context, the linear MAB uses several linear regression models (one for each action) to predict the reward for the (state, action) pair. Specifically, for each training sample comprising state, action and the corresponding reward, the linear contextual MAB trains a linear regression model that maps state to reward. Note that with the data sample comprising action index 'k', only the action model 'k' is trained. At test (inference) time, the bandit algorithm predicts the reward for each action model and thus can rank the actions based on the rewards yielded. Picking the best action all the time creates the exploitation scenario. Exploration is achieved by either using an upper confidence bound on the regression fit or using an ϵ -greedy policy, which is a way of selecting random actions with uniform distribution from a set of available actions.

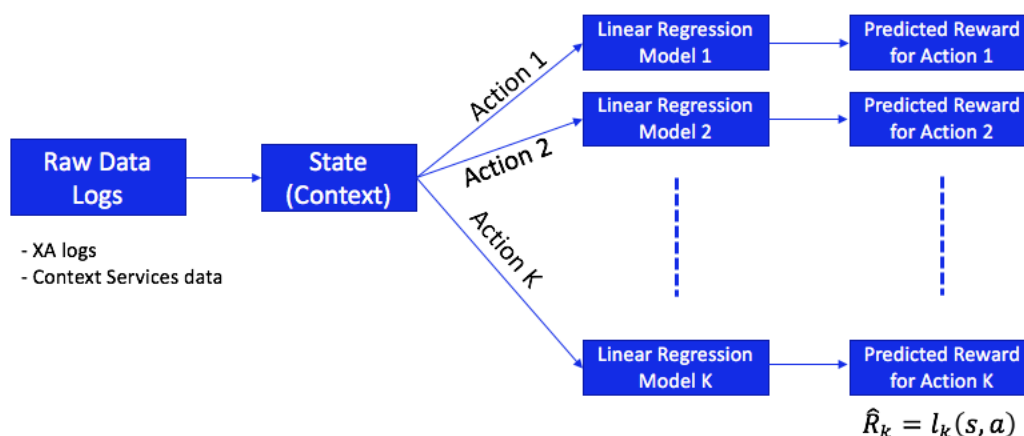


Figure 10: Linear MAB Machine Learning Flow

In general, the linear regression models may be replaced by any ML model, such as random forest, fully connected networks or generic neural networks, to create non-linear versions of the bandit architecture.

4.4. Linear Contextual MAB on the XA App: Experimental Setup

We now present the results on the linear MAB policy evaluation. This experiment was conducted using data gleaned from the Xfinity Assistant (XA) app. The problem was to correctly present the actions (or buttons) to a customer, on the XA app, that s/he would most likely click on while navigating through the app. The more relevant the presented actions are, the more engaged the customer will be, and the more likely his/her question or concern is successfully addressed.

The state, action, and reward details for this problem is described below.

State: We used the raw data from both the XA app logs and context services data. The context services data contains anonymized and individualized information such as service mix,

appointment events and present and future outage information. From this raw data, we extracted features to use as state (context) for our linear regression models.

Reward: In the XA logs, we used the following scheme to assign rewards to buttons. In the set of presented buttons,

- If the user clicked any of the buttons, that specific button gets a reward of +1.0, and other presented buttons get rewarded +0.5.
- If the user did not click on any button (which may happen if the user uttered something else or abandoned/closed the app), all buttons get rewarded -1.0.

Action: There are over 300 types of buttons in the button catalog. We only considered the 100 most frequently occurring buttons as our actions, and grouped all remaining actions into another action bucket. This means, in the context of Figure 11, we would be training 101 action models, i.e., $K = 101$.

4.5. Linear Contextual MAB on the XA App: Policy Evaluation Results

Given the data represented by the (state, action, tuple) pairs, we can provide results on training and policy evaluation [2]. In order to study how the bandit algorithm learns over time, we used the following methodology: In each iteration, we trained the bandit on 25,000 data samples and subsequently evaluated it on 2,500 test samples. We repeated this for roughly 1,000 iterations (which equates to one epoch for the amount of data we have). For evaluating the policy, we used the following unbiased offline evaluation methodology: When the action chosen by the algorithm matches the action chosen by the user in the data log, we add the corresponding reward. Otherwise, the data sample is ignored.

Figure 11 plots the net rewards obtained by the linear MAB over time. The increasing overall reward means that the algorithm is learning to explore and exploit the various patterns in the data over time. When we examined more closely into the positive and negative rewards that make up the net rewards (see Figure 12), we noticed that the algorithm was attempting to increase the positive rewards and avoid the negative ones, which is what we desired.




Figure 11: Linear MAB Policy Evaluation: Net Rewards




Figure 12: Linear MAB Policy Evaluation: Positive (left) and Negative (right) Rewards

Conclusion

The AI/ML building blocks described in this paper are currently being tested in a production environment, and some of the early results presented here have been obtained on live data. These results are encouraging and show that the productization of the whole Chatbot is just a matter of time.

There is actually a good fit between existing AI/ML methods and the cable telecommunication industry. We generate an overwhelming wealth of data, the interpretation of which far exceeds human capabilities. Our current AI/ML methods are barely scratching the surface of the possible, and focus on low hanging fruit to equal the performance of existing customer care. The example of the silent sufferers described in this paper is just an illustration of a case where AI/ML can reach beyond existing human resources to detect issues that are invisible to the human eye and simple query/filter systems or supervised learning. Unsupervised and reinforcement learning is well positioned to open new and radically beneficial frontiers in the operational transformation of the cable telecommunication industry.

Abbreviations

AI	artificial intelligence
DE	decision engine
HSD	high speed data
MAB	multi-armed bandit
ML	machine learning
NLP	natural language processing
NPS	net promoter score
RDK	reference design kit, http://rdkcentral.com/
XA App	Xfinity assistant app
XRE	cross-platform Runtime Environment

Bibliography & References

1. L. Li, W. Chu, J. Langford, and R. E. Schapire, "A Contextual-Bandit Approach to Personalized News Article Recommendation," *In the 19th International Conference on World Wide Web (WWW)*, 2010.
2. L. Li, Wei Chu, John Langford, and Xuanhui Wang, "Unbiased Offline Evaluation of Contextual-Bandit-based News Article Recommendation Algorithms," *In the 4th ACM International Conference on Web Search and Data Mining (WSDM)*, 2011.
3. Binkowski, Mikolaj, Gautier Marti, and Philippe Donnat. "Autoregressive Convolutional Neural Networks for Asynchronous Time Series." 2018.

Using Historical Traffic Data to Schedule Service Interruptions for Minimum Customer Impact

A Technical Paper prepared for SCTE•ISBE by

Jason Rupe

Principal Architect

CableLabs

858 Coal Creek Circle, Louisville, CO 80027

720-313-2434

j.rupe@cablelabs.com

Colin Justis

Associate Engineer

CableLabs

858 Coal Creek Circle, Louisville, CO 80027

303-661-3470

c.justis@cablelabs.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Background	3
1. Problem Statement.....	3
2. Related Work and Models	4
3. Data Review and Analysis	4
4. Formulated Approach to the Problem	8
Models Description.....	9
5. General Model Requirements	9
6. Competing Methods	9
7. Chosen Solution	9
8. Model Validation.....	11
Implementation Approach	12
9. System Description	12
10. Process Approach	13
Field Trial Verification.....	13
11. Field Trial Plan and Design	13
12. Performance Measurements.....	14
Findings.....	14
13. Importance of the Model	14
14. Trial surprises and findings	15
Conclusion.....	15
Acknowledgements	15
Abbreviations	16
Bibliography & References.....	16

List of Figures

Title	Page Number
Figure 1 – Usage for a single MAC address, over the time of the day	6
Figure 2 – Usage for a single MAC address, over the time of the week	6
Figure 3 – Four different MAC addresses showing very different usage patterns.....	8
Figure 4 – Daily usage for a single interface group of MAC addresses	8
Figure 5 – Rainbow plot of the conceptual model.....	10
Figure 6 – General concept of forecasting a linear trend, as it applies to the usage model's general exhibited overall trend of usage growth over time	11
Figure 7 – Application Front End	13

Introduction

Maintenance is necessary, but service disruption isn't. Some cable system repairs will impact service in ways that customers notice, but can be necessary and urgent. Fixing service while a customer is not using services is far better. But to do that without bothering the customer requires a usage forecast model.

With historical usage data of service classes, we created a simple model that predicts the amount of data being consumed by end devices in a cable plant. Using this model as an indicator of usage by customers, we can determine the best time for a repair interval of a defined duration, allowing a technician to time necessary but disruptive operations to minimize the disruption to customers. We analyzed the data, tested the model, then built a simple application based on the model which can specify the best time(s) to conduct a service disruptive repair for a defined duration, for a given set of end devices to be impacted.

The application is about to be tested in a trial. Customer call-in rate will be used to measure the effectiveness of the projected schedule, as compared to a baseline.

Background

1. Problem Statement

While an outage must be addressed immediately, impaired service is not always immediately addressed for various reasons, and an impaired network providing sufficient service is a Proactive Network Maintenance (PNM) opportunity. Addressing impaired service or PNM work, where service providers will schedule the maintenance work, is the concern of this paper.

Some maintenance required on HFC networks will impact service. But it is not reasonable to coordinate and schedule all maintenance activity with all affected customers directly. Further, customers don't want to be bothered in that way, and would prefer they not be impacted by maintenance at all. Therefore, minimizing the impact of maintenance on customers is a valuable undertaking. But it is also a difficult one; you can't just ask them if they are using the service, then go do the maintenance if they are not. If it is a large number of customers, you can't coordinate the maintenance reasonably either. So, a service usage measurement or forecast method is necessary.

But a forecast is actually better than a real time measurement method. If a truck needs to roll for a measurement to take place, then there is already a cost involved. If you can measure traffic in a center without rolling a truck, then you don't know if the usage will change by the time you decide to send the technician. So, to solve this problem, a forecast is necessary.

An open question is whether it is necessary to predict whether a customer is actively using a service or not, or whether it is better to estimate a level of usage or utilization of services, at a given future time. Certainly, we expected that availability of information would influence our interpretation of the problem, as well as how to address it. Our predictive model, implied by our framing of the question, would need to provide an accurate, actionable result. Therefore, a measure of effectiveness for our solution must address the heart of the issue, which is how do we best avoid reducing our customers' ability to use services over the HFC plant when they want to use them.

If we can use Internet Protocol Detail Record (IPDR) data to identify usage patterns in the data, by edge device in the network, we have an indication of how much disruption would be experienced at a given time if service was disrupted. This in turn could be used to schedule maintenance to avoid impacting

customers' use of the service. But for this to work, customer usage has to be reasonably predictable; we need a useful model and implementable method.

2. Related Work and Models

Forecasting models and methods are a long-standing area of applied mathematics (operations research) work. The classic book "Operations Research in Production Planning, Scheduling, and Inventory Control," by Johnson and Montgomery contains a chapter on basic forecasting methods. The well-known methods explained in this seminal textbook include regression methods, moving average methods, exponential smoothing methods, adaptive control, Bayesian methods, and Box-Jenkins models. While there are many more methods to consider, the above methods are a sufficient set to start with for our consideration.

Craig Marlow and Nick Pinckernell did some initial work to identify the opportunity. Their approach was to build a Bayesian model of whether a customer was using the service at a given time or not. While a yes or no result on usage at a given time is useful toward scheduling maintenance, we thought it more useful to focus on how much a service is being utilized, setting up for future possible enhancements where we prioritize important service classes.

This approach also allows a balance between timely maintenance and service disruption. Waiting to repair might risk a service impacting event. Further, knowing that a customer is using a service would discourage the maintenance event from being scheduled, even when that usage is very minor, perhaps at a minimal level over a long period of time, and perhaps over what is still the best option for maintenance. There could be some customers who never stop using the service in at least some small way, which would prevent any maintenance at all.

Much work has been done for decades with forecasting, and numerous forecasting methods are worthy of consideration for this particular problem. Approaches such as moving average are useful, especially with consideration to time of day, day of week, seasonal, and overall trend effects. Because we expect customer usage to be affected by the day of week and time of day, perhaps even day of month or day of year, we considered methods that would allow for such correlation effects in the forecast model.

3. Data Review and Analysis

We use Internet Protocol Detail Record (IPDR) data to indicate the amount of traffic on the network, by MAC address and by defined service type. We used filtered data which contained 15 minute traffic data by service type and by MAC address. For this early analysis and proof of concept, we decided to aggregate the service types into one estimate of traffic over the 15 minutes, for each MAC address in the data. From these aggregated data, we began examining the aggregated traffic trends among single MAC addresses, and various random groups.

It may be useful in future versions to exclude some types of traffic, or to weight the traffic types according to criticality. We leave those options to future work.

We began by looking at averages of the data by time of day, and day of week, for both single MAC addresses, and groupings such as interface groups. This first step is important for validating and invalidating our assumptions about the data, its quality, and the general behavior of the traffic statistics.

Figure 1 below shows some time of day data for a single MAC address, with bold lines indicating the average (red), 30 minute moving average (orange), one hour moving average (green), and two hour

moving average (blue). Clear time of day trends are observable, but with a high degree of variability during some times of the day more than others. These results told us that time of day would matter clearly.

Figure 2 below shows week-long trends for the same MAC address, for every 15 minute interval in the week. This figure shows similar patterns each day, but clear differences going into and out of the weekends. As in the previous figure, bold lines indicate the average (red), 30 minute moving average (orange), one hour moving average (green), and two hour moving average (blue). Once again, we see the high variability during some times of the day, but clearly there are times of the day that are generally lower than others in usage, and the day of the week matters somewhat too.

After examining several individual MAC addresses over several weeks, we found that specific MAC addresses had very different patterns from others, so clearly not all devices are being used in the same way, under the same usage patterns. We found many examples that demonstrate why it is not sufficient to just predict general usage patterns; the specific MAC address matters. See Figure 3 for some different MAC addresses showing different usage patterns.

Further, predicting individual MAC address traffic would be important for one or small groups of MAC addresses, but larger groups would likely exhibit a general trend. By looking at groupings of MAC addresses, individual differences became less important, and groups of addresses tended to look the same. In other words, general usage patterns were good predictors for when to do maintenance impacting large groups of customers; specific forecasts would be less important. See for example the interface group of MAC addresses shown in Figure 4. At this large of an aggregation, the usage tended to follow very closely this pattern no matter the MAC addresses in the grouping.

While each MAC address did have important differences, there is still much we can say generally which is of use. Ideal service times differ significantly between MACs, but tend to coincide with early morning hours (before 7:00am), generally. When constrained to typical work hours (say, 8:00am to 5:00pm), service times must be analyzed per MAC in order to minimize disruption, as generally the differences between MAC addresses becomes important during those times of the day.

By averaging over daily and weekly usage, an expectancy can be obtained for the time of day and time of week. The most direct way to project on internet usage is to perform a moving average over a rolling period, such as 4 weeks.

Note that we intend to use this information to predict whether services are in use over a particular end device. Generally, we assume a customer location to be aligned to one MAC address, though that is not a critical assumption to the project.




Figure 1 – Usage for a single MAC address, over the time of the day







Figure 2 – Usage for a single MAC address, over the time of the week






Figure 3 – Four different MAC addresses showing very different usage patterns




Figure 4 – Daily usage for a single interface group of MAC addresses

4. Formulated Approach to the Problem

Based on our analysis of the data, we learned several ideas which framed our model.

- Usage data exhibit clear patterns which might be exploitable to minimize disruption of service during a repair.
- Clear, significant differences in usage patterns by MAC address leads us to predict individual MAC usage patterns independently.
- Grouping MAC usage models in small groups was important to predict the best times to impact service for the group. But as groups got larger, such as an interface group, a general model was likely sufficient for many times of the day.

- Visually, it was clear to see there were time of day, and day of week effects that were important for almost all MAC addresses studied. We further suspected there was an overall increasing trend of usage too. While we did not have enough data to find an effect for time of the year, we have strong suspicions that there are effects due to holidays, summer vacations, etc. Thus, we recommend:
 - Obtaining a year's worth of data to find annual patterns to add as effects to any chosen model, and
 - Understanding in some way (predicting) the risk of a forecast, especially over days where there are no data to contribute an annual effect to the model.

From these observations, we decided a simple model predicting the amount of usage on each MAC address would be a useful first model. Further, as we could see clear effects, we sought a linear model of these effects as a simple, sufficient way to predict usage for short periods of time into the future, say a day to a few days. While not ideal, it was sufficient for our proof of concept and trial.

Models Description

5. General Model Requirements

We recognized that interrupting lower amounts of usage might still be better than interrupting higher amounts, even if there is usage, so we focused on predicting the amount of usage over whether services were being used or not. Further, we considered classifying usage into discrete levels of usage as a compromise between the on-off and full fidelity of usage level, but later decided to stay with a direct approach as a starting point, not having enough knowledge to set finite usage levels.

6. Competing Methods

We considered a Bayesian approach which would predict whether service was being used or not. This did not meet our criteria for predicting usage level. We considered a Bayesian approach to determine usage level, but again decided to stay simple for our first model, if we found a simple approach that appeared reasonable, which we did.

We also considered several linear prediction models, some using various forms of moving average, as forecast models. Without a large amount of data to do a serious comparison of model methods, we did not have a reason to go with a complicated model for the proof of concept.

7. Chosen Solution

When it comes to predicting internet usage for a single customer, or group of customers, there tends to be consistency in the usage for time of day and time of week. However, the usage can still vary slightly from day to day, or week to week. For short-term projections, it is ample to predict future internet usage using linear projection from a line of best fit for the projected data.

Essentially, the model we used for predicting future usage is as follows. Given a few weeks of data is all we had to work with, we recognized there was no way to model for special days of the year. Therefore, we used the available data to get a weekly effect, day of week effect, and time of day effect. Then, we combined these linearly to form projections into a short future of a week.




Figure 5 – Rainbow plot of the conceptual model

An example for day-to-day, or week-to-week, projections is shown in Figure 5. This figure is a 3-D rainbow plot where the time of day (or week) is treated as independent of the number of days or weeks passed; color indicates a particular day or week, and paleness represents the depth or time of day or week. The time of day or week can be measured in hour, half-hour, or 15 minute intervals. Internet usage is measured as bytes passed (usually megabits (MB) or gigabits (GB)) for each of those time intervals. Finally, the lines of best fit are calculated independently for each time of day or week; the slope and y-intercept can be different for different time intervals of the day or week.

The simplest way to calculate a line of best fit is through method of least squares. In that case, the line of best fit becomes

$$y^* = \text{corr}(x, y) \frac{\sigma_y}{\sigma_x} (x - \langle x \rangle) + \langle y \rangle, \quad (1)$$

where x is the time of day or week, y is the internet usage rate, and y^* is the usage for the line of best fit. By convention, standard deviation is denoted by σ , and $\text{corr}(x, y)$ is the correlation between x and y . The Root Mean Squared (RMS) error becomes

$$\text{Error}_{RMS} = \sqrt{1 - \text{corr}(x, y)^2} \sigma_y \quad (2)$$

Because the standard deviation is the RMS error for an average mean, (2) shows that the RMS error for a least squares line of best fit projection is always no greater than a simple flat projection of a moving

average. In other words, there is nothing to lose by performing a linear projection, regardless of whether internet usage varies significantly day-to-day or week-to-week.

Once we have a line of best fit, it can be projected into future days or weeks for some short time period. A hypothetical example for a given time of day or week is shown below in Figure 6.




Figure 6 – General concept of forecasting a linear trend, as it applies to the usage model's general exhibited overall trend of usage growth over time

The slope can still vary over long time periods, so it is important not to project too far. At most, the projected usage should not be as long as the measured usage. The measured internet usage would be on a moving time window. For instance, a queue of one month's worth of data could be collected, and then the next day or week would be projected. After the next day or week, a new day's or a week's worth of data would be added to the model. Potentially, as older data becomes less helpful to the prediction, it would be removed from the model. A weighting of older data, such as in an exponentially weighted moving average model, would be best, as in a Box-Jenkins approach.

From looking at simple averages, the best service times can be found, along with a measure of the reliability of a chosen service interval for every grouping of customers possible.

8. Model Validation

We tested the model using different data sets at different CMTSs. Further, we took a trained model and used it to predict results for times which we had data. As the chosen model showed promise for being a sufficiently accurate model as we could measure that, we decided to continue to conduct a trial where we could measure the real impact of the approach against actual customer calls of interrupted service.

Implementation Approach

9. System Description

To support the trial of the model, the model was encoded and implemented with a system and process.

The model described above was implemented in a two-stage process for the sake of the field verification trial; an implementation would be very similar. The two stages involve a processing stage, and then an application stage.

The processing stage receives updated IPDR data periodically, and applies the modeling described above to form a new set of predictions for the forecasted horizon. The output from this stage forms a table of expected usage by MAC address for each of the 15 minute intervals, which is the resolution of the data, over the horizon being predicted. This output table is the input to the next stage.

The application stage is a web-based or locally-cached application which performs table lookups. The application front end created by CableLabs, shown in Figure 7 below, is the user interface which collects the MAC addresses to potentially be disrupted, the window within which the repair needs to be scheduled, and the duration of the service interruption expected. The application performs a table lookup of the MAC addresses, adds the predicted usage for each 15 minute interval over the duration of the scheduling window, then does a moving average of the disruption window size over the duration window, reporting on the lowest usage candidates. The application reports the top few options for minimal impact based on the overall usage statistics calculated.



Project Presence

No more guessing.

Find the best service time.

Window Start*

Saturday, July 1, 12:00 AM

Window End*

Sunday, July 2, 2:00 AM

Duration (minutes)*

120

MAC Addresses*

00:0b:b6:1b:19:08

00:0b:b6:10:47:c4

00:0c:e5:2e:4f:58

Best Time

Saturday, July 1

1 PM for 120 minutes

Good Alternatives

Saturday, July 1	1:15 PM
Saturday, July 1	1:30 PM
Saturday, July 1	1:45 PM
Saturday, July 1	2 PM

Worst Time

Saturday, July 1	3:15 AM
------------------	---------

FIND THE BEST TIME

Figure 7 – Application Front End

10. Process Approach

Each evening, new data are gathered for the node, and fed into the model. The model is updated with the new data, creating a new forecast of 15 minute usage predictions for each MAC address on the node. The resulting file is uploaded to the application server. The application is then updated to connect to the new file, and tested. Once confirmed as functional, the application is able to process off of the new model results table. Because the application exists separate from the data file that is updated periodically, there is no duration over which the application is down. Instead, the application simply updates to the new information when made available. The application then processes updated intervals each time it is used.

At the start of the work assignment day, service personnel bring up the application through a browser. They then use the application to assign the work times for the work done in the area covered by the application.

Field Trial Verification

11. Field Trial Plan and Design

For the field trial, we selected two CMTSs in the Logan, Utah area. We pulled data from this location before and after summer break for the local schools, as a comparison to determine whether the model

results change across this known usage change. These data were used to form the reference model for the field trial.

Each evening, we extracted new IPDR data in 15 minute increments for the past 24 hours, and incorporated the new data into the old. This updated data set was used to train a new model for the next work day. The back-end model ingests the updated data, and the output is a table of 15 minute predicted usage for the future week for each MAC address.

As of this writing we have yet to begin the trial, so the rest of the planning and design have yet to be tested. We anticipate the following general activities to follow. All trial participants will need to be briefed as to the changes in the operations steps when assigning work and conducting maintenance. Those assigning the work will need to add the step of using the front-end tool to determine the best times to conduct the maintenance, based on their best information about what needs to be done, how long it will take, and who will be impacted. If changes in the field are necessary, a line of communication needs to be established so that technicians and those assigning the work can agree as to any adjustments. This action will help maintain the integrity of the trial. Further, field technicians who conduct the maintenance will need to record the times when service was disrupted and restored to be sure it did or did not overlap the times indicated by the model. A good model that can't be followed is not very useful, so we must track its usability as well as its accuracy.

12. Performance Measurements

The key measure of performance for this trial is the number of customer call complaints per impacted customer per unit of service interruption time. We track this by collecting from records the number of customers who call in to indicate a service interruption during the maintenance time for the experiment group, and comparing this to the control group. To normalize for each maintenance event, we take the number of customers who call in to complain out of the group of interrupted customers, and divide that by the number of interrupted customers times the interruption duration. Each maintenance event has one measure of performance result; if the maintenance action required more than one outage, we simply add the performance measure for each outage for that maintenance event. We collect this measure of performance for the experimental group, and for a control group as well. Then we calculate basic statistics from the measure of performance including mean, standard deviation, and confidence bounds. Finally, we calculate statistical confidence bounds and conduct statistical tests to determine whether we can say with confidence whether the results indicate effectiveness of the solution.

Given the measure of performance offered, we expect the measure of performance to be distributed Poisson, so simple statistical tests on Poisson parameters should apply.

Findings

13. Importance of the Model

While the trial will reveal some information with which to determine the importance of the model, a single trial will not reveal very much. And because we had access to very limited data, we do not consider our model to be necessarily the best, but sufficient.

Instead of simply adopting the model given in this paper, we suggest, in a full implementation, that multiple models be formed, tested, and used in competition, with a long-term adopted model to result from the experience of field use. The model reported here was the result of a few competing models considered and compared on multiple merits, including accuracy of prediction. Had we more data to work

with in the development of these models, perhaps a different model would result. More experience, and more data, are in order.

A good model can be created from the data only, but there are implementation differences that can make one model better than another, and one approach to implementation better than another as well. Consideration of specific applications is important (network, operations, environment, OSS, etc.).

14. Trial surprises and findings

The trial is planned to begin in September. We hope to have preliminary results to report in our presentation at the Expo.

Conclusion

By using available usage information from the network, a simple model can be created to predict the traffic through an end point on the network. Service class usage in 15 minute increments can be used to form simple predictive models of usage, projecting a few days into the future. This prediction of usage can be used to schedule maintenance so that the impact on traffic is minimized. The expectation is that by impacting the least amount of network traffic we reduce the impact on customers. This lower impact should be measurable through a lower customer call in rate, so that fewer customers will call to complain about service outages when the model is used to plan the maintenance activities. This idea is to be tested in a field trial soon.

The prototype built for the trial demonstrates that a process, using a simple model based on IPDR data, and utilizing a web-based front-end interface, can provide work assignment windows for planned maintenance which would interrupt service.

The analysis we conducted showed definite patterns in usage at network end points (MAC addresses), and in groupings of end points. While there were considerable repeating patterns in the data across MAC addresses, not all MAC address usage patterns were the same. But in most cases, there were large differences in the peak usage compared to the minimum usage, and long periods of time over which usage remained mostly low. The analysis suggested a model to predict usage by MAC address was achievable and could be useful for scheduling maintenance.

By predicting usage on single MAC addresses, then clustering the models for the group of MAC addresses to be impacted, a merged model was created for each maintenance case. By searching the time over which the maintenance was desired to be scheduled, the best times to schedule a maintenance activity for any given duration could be found easily. By creating a front end to the model, and a process by which we could update the model as frequently as daily, we were able to build a prototypical solution which could be trialed to prove the concept further.

Acknowledgements

A special thanks to Larry Wolcott who drove this work and made sure it met success, and for all his input and help with the paper and trial. More thanks go to the software team at CableLabs who supported the front-end for the trial. Also, a big thanks goes to Jay Zhu and John Phillips at CableLabs who stepped in to support the work when a primary author was on leave.

Abbreviations

GB	gigabits
HFC	hybrid fiber coax
IPDR	Internet protocol detail records
MB	megabits
PNM	proactive network maintenance
RMS	root mean squared

Bibliography & References

Lynwood A. Johnson, Douglas C. Montgomery, *Operations Research in Production Planning, Scheduling, and Inventory Control*, John Wiley & Sons, New York, Copyright 1974, ISBN 0-471-44618-1.

Douglas C. Montgomery, *Introduction to Statistical Quality Control*, John Wiley & Sons, New York, Copyright 1985, 1991, ISBN 0-471-51988-X.

What Gets Measured Gets Done / What Gets Analyzed Gets Transformed

Analytics for a Wider/Deeper Network View

A Technical Paper prepared for SCTE•ISBE by

Venk Mutalik

Comcast

Philadelphia, PA

Venk_mutalik@comcast.com

Dan Rice

Comcast

Philadelphia, PA

Daniel_rice4@comcast.com

Karthik Subramanya

Comcast

Philadelphia, PA

Karthik_Subramanya@comcast.com

Jon-en Wang

Comcast

Philadelphia, PA

Jon-en_wang@comcast.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	4
Overview	4
1. Common Mode Disturbance	6
1.1. Field Measurement	6
1.2. Lab Simulation.....	7
1.3. Mitigation Design.....	8
1.3.1. Lab Test of Mitigation Modulation Profiles.....	10
1.4. Field Test and Deployment Results	12
2. Data Analytics	14
2.1. Leading and Lagging Indicators	17
Conclusion.....	19
Abbreviations	20
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 - Opportunity development process.....	5
Figure 2 - Example switching power supply circuit ²	6
Figure 3 - Coupling decoupling network vs. CM with loose connector, on an HFC network with max-hold spectrum analysis.....	7
Figure 4 - CMD noise time-domain analysis	8
Figure 5 - Example DOCSIS modulation profile design with transient noise.....	9
Figure 6 - Balancing robustness vs. efficiency to maintain total bonded capacity.....	9
Figure 7 - Modulation Profile Efficiency per burst size model.....	10
Figure 8 - 1 CMD interference source, throughput improvement in upstream and downstream.....	11
Figure 9 - 2 CMD interference sources, throughput improvement in upstream and downstream.....	11
Figure 10 - Packet Loss, CER and CCER for UDP packets for default and new modulation profile	12
Figure 11 - CER and CCER before and after modulation profile configuration change	13
Figure 12 - Single Interface example where configuration was modified in the middle of a CMD noise event over several days	14
Figure 13 - Green node with 3 upstream channels.....	15
Figure 14 - Green node 3 channel SNR and TX correlation.....	15
Figure 15 - Yellow Node with 3 US channels CM TX Power and SNR	16
Figure 16 - Red node with 4 US channels CM TX Power and SNR	17
Figure 17 - Network performance cross-correlation matrix.....	18
Figure 18 - Cross-correlation matrix dashboard	19

List of Tables

<u>Title</u>	<u>Page Number</u>
Table 1 - Opportunity development work streams	5

Introduction

With the acceleration of technology in homes comes a corresponding increase in the number of switching power supplies potentially impacting the upstream plant. More and more in-home electronics devices -- Internet-connected appliances, battery chargers, LED lights, video set-top boxes (STBs), broadband gateways and cable modems -- come with switching power supplies inside of them, which contribute to an age-old issue known as Common Mode Disturbance, or CMD.

This is happening coincident with the industrial shift away from traditional centralized architectures to distributed architectures. Distributed Access Architectures (DAA) are on the rise because of a growing need to fulfill newer needs, such as low-latency and high-speed applications. Yet the traditionally persistent issue that is Common Mode Disturbance (CMD) continues to impact networks in negative ways. While tools have improved dramatically in the last few years in addressing such pesky problems as CMD, it continues nonetheless to impact even the more modern fiber deeper and distributed networks.

Specifically, the rise of CMD noise, in part triggered by the explosion of Internet-connected CPE in our customers' homes, catalyzed within Comcast an impairment identification and mitigation framework described in this paper. The "identification" portion of the framework is informed by machine-level telemetry data, to better measure the impairment; and the mitigation portion of the framework is enabled by advanced data analysis. (Hence the title, "what gets measured gets done; what gets analyzed gets transformed.")

Our intent is to provide new insights into age-old problems, as well as a framework for analyzing old and new problems alike. New, machine-informed ways of looking at the traditional time and frequency domains of RF information can help to create a "Taxonomy of RF Impairments" -- a first in the industry (to our knowledge), which has developed as a playbook to tackle impairments. This work will ultimately lead the industry toward an effective use of cable assets and aid in the creation of a more elastic, low latency network.

Overview

Comcast has initiated a suite of projects to address some of the systemic operational challenges related to the evolving Hybrid-Fiber-Coaxial (HFC) and optical networks by applying some new advanced technologies, and by innovating on our existing platforms. Several examples are highlighted in this paper, including mitigating switching power supply noise and new data analytics. The first, RF ingress mitigation, is described in detail below. The opportunity development process is described in Figure 1, with four discovery and solution development stages that feed the solution development and deployment funnel, as shown in Table 1.

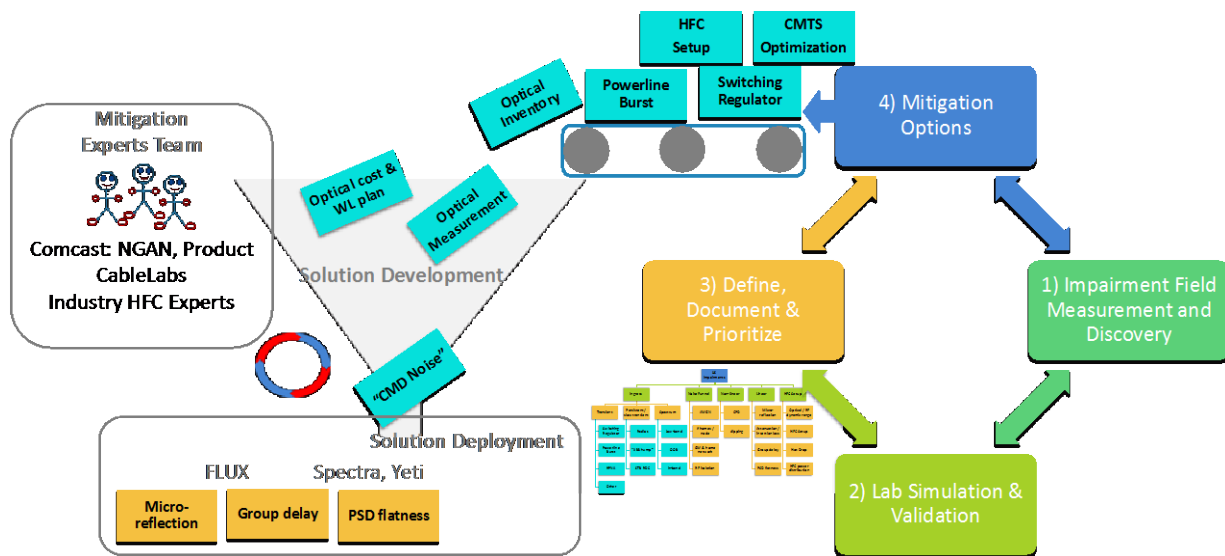


Figure 1 - Opportunity development process

Table 1 - Opportunity development work streams

Stage	Description
Impairment Field Measurement and Discovery	Opportunities to improve operational efficiency are discovered doing field measurements using lab grade Test and Measurement (T&M) equipment and leveraging Comcast's PNM and other OSS tools. Currently a field measurement campaign is ongoing to identify the root cause of the most impactful challenges to the access network that are causing operational expense and impacting customer experience.
Lab Simulation & Validation	Once these challenges are discovered in the field, they are evaluated in-depth in the lab to characterize the specifics and test different scenarios, to understand the impact and opportunity for improvement.
Define, Document & Prioritize	After the challenges are characterized in detail in the lab, they are documented and prioritized for the design of mitigation approaches. These details are added to a growing taxonomy of HFC impairments, with comprehensive descriptions characterizing the impairment. As the taxonomy grows, we plan to continually update the industry with additional descriptions for CMD and other noise.
Mitigation Option & Solution Development	For each of the prioritized challenges, we developed a set of options to mitigate or reduce the impact. One of the mitigation options detailed in this paper relates to CMD Noise. Other options under development for CMD noise issues range from customer communication channels to new self-install-kit (SIK) connectors to new low-cost HW that can block the noise from getting into the network. Once the mitigation options are prioritized, solution design is completed.
Deployment	Deployment of the solutions include field trials to evaluate both network performance metrics, such as Modulation Error Ration (MER), uncorrectable codeword error ratio (CER) and packet loss, and operational business metrics such as call-in rates (CIR), tickets, and truck rolls. Based on the efficacy of the solution, it can be deployed across the network.

In recent years Comcast has developed some very effective operational tools described in other SCTE publications.¹ This paper describes another such development and approach, focused on a growing HFC impairment, with an opportunity to improve the customer experience and operations. In essence, a new take on operationally-hardening a mature technology. Future focus areas, as shown in Figure 1 include additional advanced data analytics and optical measurement solutions.

1. Common Mode Disturbance

1.1. Field Measurement

With the acceleration of technology in homes there is an increasing number of switching power supplies. Switching power supplies are increasing with the growth of home electronics, Internet-connected appliances, battery chargers, LED lighting, video set-top-boxes (STBs), broadband gateways and cable modems (CMs), among many other uses. These power supplies convert AC line power to different DC voltage levels required for the consumer electronics circuits, through methods such as switching the current into a capacitor whose voltage is monitored and controls the frequency of the switch. One example power supply circuit is shown in Figure 2.²

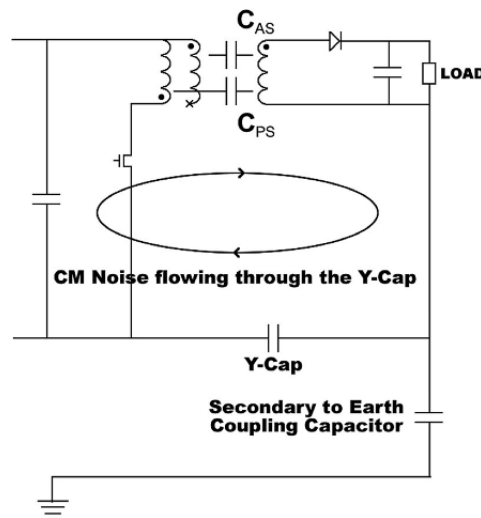


Figure 2 - Example switching power supply circuit²

The initial in-rush current when charging the capacitor may result in a noise current spike onto the ground of the device. This noise is referred to as Common Mode Disturbance (CMD) by power supply engineers. Because the coaxial cable's outer conductor is grounded, it can become a path for the noise

¹ L. Wolcott, J. Heslip, B. Thomas, R. Gonsalves; A Comprehensive Case Study of Proactive Network Maintenance, SCTE TEC EXPO 2016

² Y. P. Chan, B. M. H. Pong, N. K. Poon and J. C. P. Liu, "Common-mode noise cancellation in switching-mode power supplies using an equipotential transformer modeling technique", IEEE Tran. Electromagn. Compat., vol. 54, no. 3, pp. 594-602, 2012

current. When there is an imbalance in the coaxial transmission path, the common mode noise current converts to differential mode current.³

Examples of imbalance in the coaxial transmission path include loose connectors with poor ground continuity, shield break on a cable, bending⁴ (kinks) in the coaxial cable, and impedance mismatches. When an imbalance occurs, mode conversion occurs. In other words, the common mode noise from that home will couple into the cables that funnel data from all the other homes in that serving area, to the Cable Modem Termination System (CMTS), thus impacting the performance for all devices on the upstream signal path. Characterizing this type of noise, to ensure equipment attached to a cable network does not negatively contribute to the HFC noise levels, has been standardized as part of the SCTE 249 IPS standard based on a coupling- decoupling network.⁵

One example of CMD noise coupling into the HFC through a loose coaxial connector on a CM is shown in Figure 3, along with the same noise measured per the SCTE 249 test method. Another similar picture of the noise, as seen by a spectrum analyzer with min and max hold located near the CMTS, is also shown in Figure 3, with the noise coupling into the network at an impactful level underneath the 23.7 MHz carrier. Note that both pictures show the peak of the noise power between 20 and 25 MHz -- with impact into a lower DOCSIS 3.1 upstream carrier centered at 17 MHz typically placed below the 23 MHz Carrier.




Figure 3 - Coupling decoupling network vs. CM with loose connector, on an HFC network with max-hold spectrum analysis.

1.2. Lab Simulation

Upon laboratory investigation, this noise signature appears to be highly impactful to the DOCSIS signals, from the perspective of a spectrum analyzer on max hold or the SCTE standard CDN. In fact, it can create service-impacting DOCSIS codeword errors, resulting in packet loss. A deeper lab perspective was

³ A. Axelrod, K. Povolotski, and S. Nir, "Experimental study of DM to CM conversion in elements of data communication links," in Proc. Int. Eur. Electromagn. Compat. Symp., Sorrento, Italy, Sep. 2002, pp. 435–440.

⁴ Xinglong Wu, Flavia Grassi, Sergio A. Pignari, Paolo Manfredi, Dries Vande Ginste, "Circuit interpretation and perturbative analysis of differential-to-common mode conversion due to bend discontinuities", Electrical Design of Advanced Packaging and Systems Symposium (EDAPS) 2017 IEEE, pp. 1-3, 2017.

⁵ SCTE IPS TP 228 expected to become SCTE 249 standard before SCTE TEC Expo

obtained, beyond field measurement, by characterizing the signal with a Vector Signal Analyzer -- a high-speed sampling scope for signal analysis. While this noise appears to be intractable and a potentially significant customer experience detractor in the frequency-domain, mitigation opportunities become even more pronounced in the time-domain. CMD noise, viewed in the time-domain, is shown in Figure 4.




Figure 4 - CMD noise time-domain analysis

The CMD noise is actually a very periodic burst noise. The burst rate depends on the frequency of operation of the switching power supply, typically 50 to 75 kHz or a period of 14 to 20 usec. Other switching regulators have been seen up to a 200 kHz switching rate. When the in-rush current of the regulator spikes, the duration of the noise burst is only about 1 usec, or less than 10% of the time. Measured across a variety of make and models of Consumer Premise Equipment (CPE) and power supplies, the time-domain shows a very tight bound of period and duration across the equipment.

1.3. Mitigation Design

The Reed-Solomon (RS) method of forward error correction (FEC) used in the DOCSIS 3.0 upstream signal path reaches a rate of diminishing return with the amount of overhead vs. error correction performance with respect to Additive White Gaussian Noise (AWGN) or time-invariant noise. That said, Reed-Solomon encoding is a great coding scheme for dealing with transient noise sources, especially with the support of an interleaver. From a spectrum analyzer perspective, the noise appears to be time-invariant and very difficult to mitigate without sending a truck to a customer's house. From a time-domain perspective, it falls into the area where D3.0's error correction can have its maximum benefit.

Without going through detailed modeling, Figure 5 visualizes the concept. Based on a default D3.0 burst profile configuration that had been used on the network, a new, D3.0 burst profile was designed to eliminate the packet loss caused by the CMD noise. The CMD noise, at a 14 to 20 usec period, will impact every single codeword. In fact, one of the key metrics used in operations to identify when CMD noise is coupled into the network is a very high rate of correctable codeword error ratio (CCER) because it is causing errors to every codeword if it is coupled in at a high level.

If each time the burst of noise hits a codeword, and it is at a high enough power level to cause an error, it will error four RS symbols for the given modulation and symbol rate (64 QAM, 6.4 MHz) in the example. At the current 20 usec period, this burst can hit the short data grant three times and the long data grant five times with the default configuration. In the default configuration, only three bursts can be fixed in the

short data grant, and four bursts can be fixed in the long data grant. As a result, the performance is variable for short packets and very bad for long packets or concatenated data bursts. By restructuring the short and long data grant, the short data grant codeword can fix more than three bursts, and the long data grant can fix four bursts. The long data grant is re-defined so that only four noise bursts can impact the codeword, based on the time-domain characteristics of the CMD noise. Similarly, the Unsolicited Grant Service (UGS) data grant can be re-defined to eliminate the packet loss caused by CMD to the voice packets. If there are multiple sources of CMD noise, a similar analysis and modulation profile can be developed to manage the uncorrelated noise sources. Additionally, applying a D3.0 block interleaver can add additional margin when there are multiple CMD sources, which are uncorrelated in timing with each other and the data traffic and codewords.

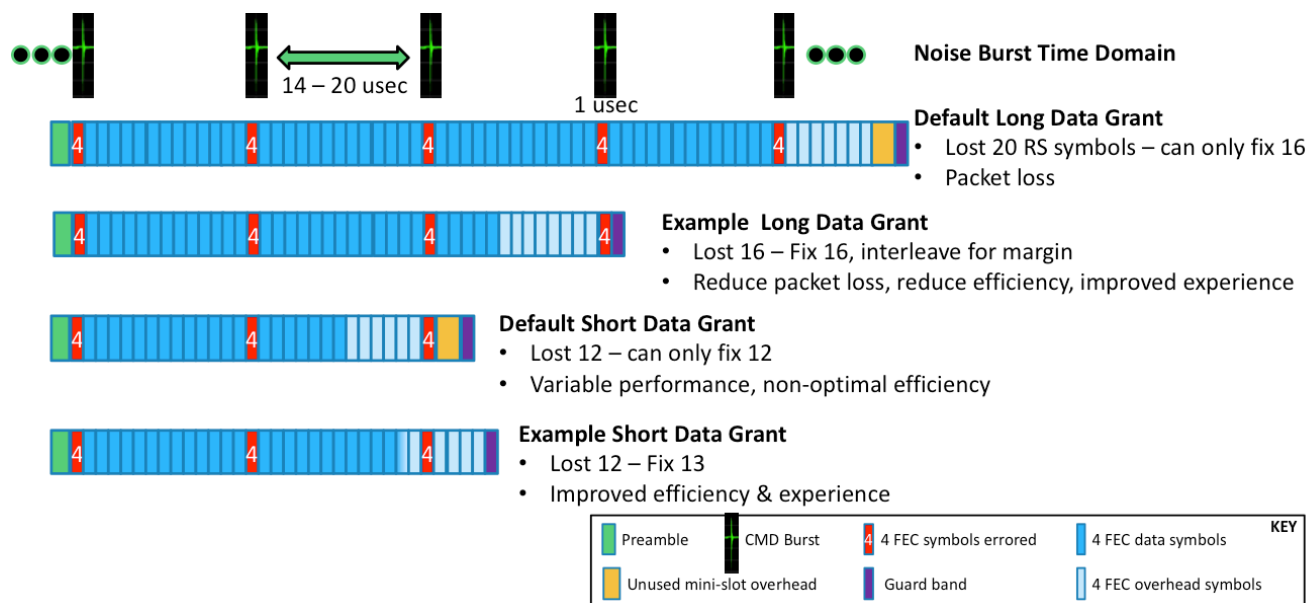


Figure 5 - Example DOCSIS modulation profile design with transient noise

When adjusting modulation profiles and D3.0 US channel parameters, it's important for all the configuration “knobs” to be set compatibly and collectively. By also adjusting the symbol rate, max burst size for the short data grant, preamble length, and guard time, the overall efficiency across packet sizes can be improved. When these techniques are applied to the lower two channels, most impacted by CMD noise, the result is increased robustness. Applying more efficient lower overhead FEC to the upper two channels can improve capacity and efficiency. The overall capacity of all four bonded channels can be maintained while mitigating the impact of the noise on the user experience as shown in Figure 6.




Figure 6 - Balancing robustness vs. efficiency to maintain total bonded capacity

The overall efficiency for each channel can be modeled based on the packet size or concatenated transmission burst, as shown in Figure 7.

- Because of default modulation profile inefficiencies for small packets, the example improves data efficiency by ~5% across all channels.
- Efficiency is reduced by ~4% for larger packets across all channels
- Efficiency is reduced by ~8% for Channels 1 and 2 in lower spectrum
- Speeds are improved by fewer TCP slow starts and re-transmissions, as shown in next section

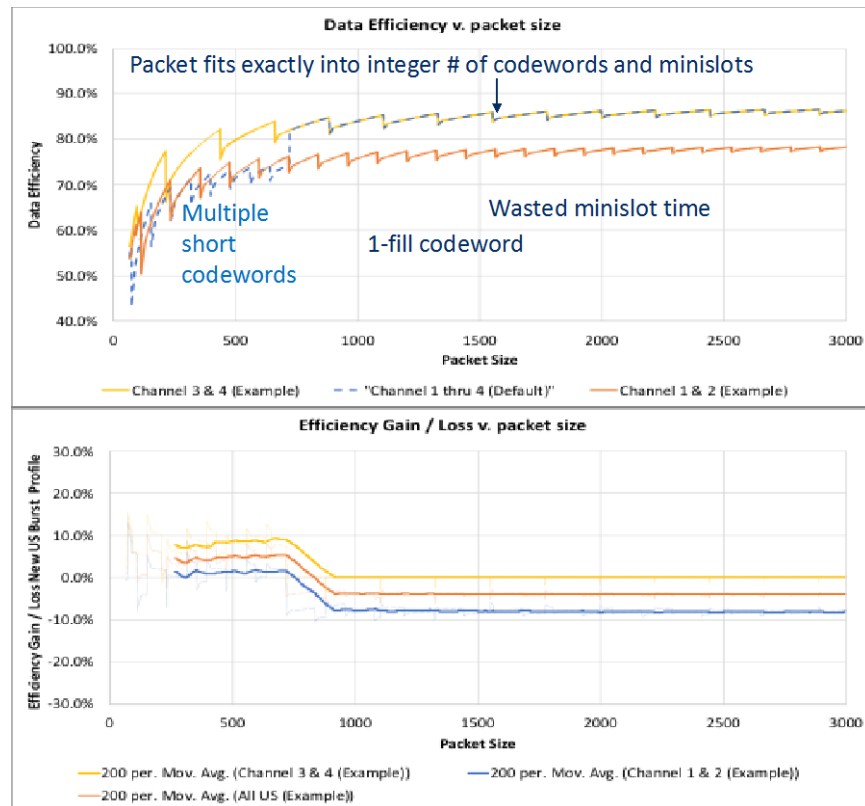


Figure 7 - Modulation Profile Efficiency per burst size model

1.3.1. Lab Test of Mitigation Modulation Profiles

These example configurations were tested in the lab to determine if the D3.0 FEC design could improve the customer experience. The new modulation profiles were tested with both TCP and UDP traffic while measuring the loss from the traffic generator. The level of CMD noise coupled into the upstream was increased until a significant impairment was caused, and then the new profile was applied to test the improvement in customer experience. Figure 8 and Figure 9 show the improvement in customer experience for TCP traffic. As the CMD is added, it impacts the bottom two channels, and the throughput of the network dropped in half as the upstream packet loss was increased for the two lower channels. As the new profiles are configured, the upstream and downstream throughput improves back to levels equivalent to no impairment, even though the noise is still in the channel. Because the TCP performance in the downstream is also impacted by the packet loss in the upstream, the downstream performance is also impacted by the upstream CMD noise, and improved by the new profile. The packet loss is reduced to a level that is no longer impactful to the customer experience for both one and two uncorrelated CMD interference sources.

Carrier level to CMD Noise Peak = 21 dB




Figure 8 - 1 CMD interference source, throughput improvement in upstream and downstream

Carrier level to CMD Noise Peak = 21 dB





Figure 9 - 2 CMD interference sources, throughput improvement in upstream and downstream


Figure 10 shows the CER, CCER and UDP Packet Loss. Because the CMD noise is hitting every codeword, there is a high level of CCER, but many codewords are uncorrectable and causing packet loss. After applying the new modulation profile, the CER drops to an acceptably low level, reducing the packet loss. The CCER increases because all the errors are now being corrected, which reduces the packet loss and improves the customer experience.

IPERF TCP/IP Throughput, 6 dB attenuation of CMD Noise

Test Case 4: TCP/IP Throughput w. Impairment at 21 dB Below Carrier



Test Case 15: TCP/IP Throughput w. 2 Interferers
Impairment at 21 dB Below Carrier






Figure 10 - Packet Loss, CER and CCER for UDP packets for default and new modulation profile

While CCER is generally considered “a bad thing,” because it means there is noise present in the network and causing errors, it is also a good thing if it means fewer uncorrectable errors are occurring, as in this example. The CER has been converted to CCER and has reduced the packet loss, thus improving the customer experience.

1.4. Field Test and Deployment Results

These new configurations were tested in the field against a population of nodes that were exhibiting interference from CMD noise. Codeword errors were tracked across all the CMTS upstream channels for a population of CMTSs that had the configuration changed. For each five-minute time sample, the CER was classified as not degraded, degraded or severely degraded. These samples were then plotted and accumulated to understand the impact of the new modulation profile design. A similar population of CMTSs acted as experimental control and did not have the configuration changes applied over the same period of time. During the test the CNR of the channels for the configured and control population was also tracked to identify any changing network conditions that would differ between the groups. Before and after comparisons and comparisons between the experimental and control population were completed. The results were the following:

- CER Degradedness was reduced by 20% across all configured US interfaces vs. 12% experimental control
- CER degradedness for the most impacted CMD interfaces (23 MHz) was reduced 20% vs. 8.6% experimental control

Many examples were identified of improved CER, as shown in Figure 12. In these figures, after the new modulation profile configuration was changed, the CER was reduced to acceptable levels. The CCER is still showed that significant errors were occurring, and that the noise was still present in the network -- but because the errors had been corrected by the FEC, the customer experience was improved. This CCER dynamic is useful because it indicates there is still a noise source on the network that can be fixed to improve performance, but the customer experience has been much improved -- enabling network engineers to address the problems disaffecting the customer experience first.




Figure 11 - CER and CCER before and after modulation profile configuration change

A simplified example is shown in Figure 12 for one of the channels where the errors are seen to be happening regularly before the change, both corrected and uncorrectable. A time frame of continuing errors over several days was occurring during the configuration change window. Even though the noise remained in the network, causing errors, the errors were all being corrected after the new profile was deployed. When the noise spikes re-occurred in the future, the CER was maintained at a low level sufficient for improved customer experience.




Figure 12 - Single Interface example where configuration was modified in the middle of a CMD noise event over several days

2. Data Analytics

Over the last many years, MSOs have made a quantum leap in securing real-time data about the state of their networks. Some of the data is derived from the CMTS; other data is from the CPE and CMs; yet other data comes from specialized equipment that tracks noise and other effects throughout the system. These are unofficially known as the “Sources of Truth” (SOT) for MSO analytics teams. This data is then analyzed, curated and made available to multiple teams within the organization to seek incremental improvements to the customer experience.

Within Comcast, these sets of real time telemetry data are collected and presented via several SOTs. These data sets are analyzed using advanced analytic tools and dashboards set up to test various Proofs of Concepts (POC) before rolling them out system wide as analytics tools.

The goal of the data analytics effort has been to leverage available SOTs, drill down to the MAC address level and create associations that help network engineers to understand pervasive noise impairments, necessarily separating them from transient hits that mar service, but are hard to pin down. In addition to understanding network impairments, analytics can be used to optimize network configurations and HFC physical setup metrics.

Nodes could be classified as Green, Yellow or Red in proportion to customers affected over a period of time. Such classification, based on multiple internal constructs, enables a common focus for setting up priorities. Notice that the red nodes may be so designated because of a car colliding with the telephone pole, thus taking out service -- or because of excessive transient ingress coming in and marring the customer experience.

Figure 13 is a snap shot of a Green node with three upstream (US) channels. The graph on the left is the distribution of CM TX level, and the graph on the right shows SNR distribution across the CPEs for the same three RF US channels. While this is a popular way of constructing analysis, it does not show the full story.

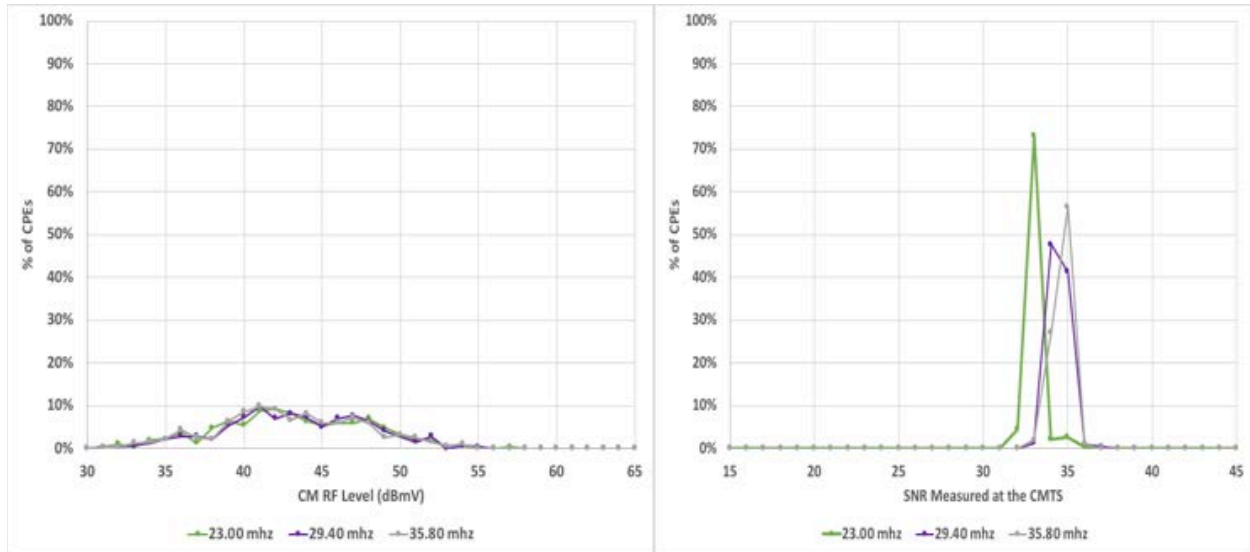


Figure 13 - Green node with 3 upstream channels

Figure 14 combines both of the above representations and shows the RF transmit level vs. SNR values, in one spot. This enables us to see that this Green node has a tight distribution of SNR across the RF levels (as expected) and points to a well-behaved node deserving of its Green designation.




Figure 14 - Green node 3 channel SNR and TX correlation

A Yellow node, on the other hand, has a slightly wider distribution on the SNR metric, which is also reflected in the scatter plot, as shown in Figure 15.

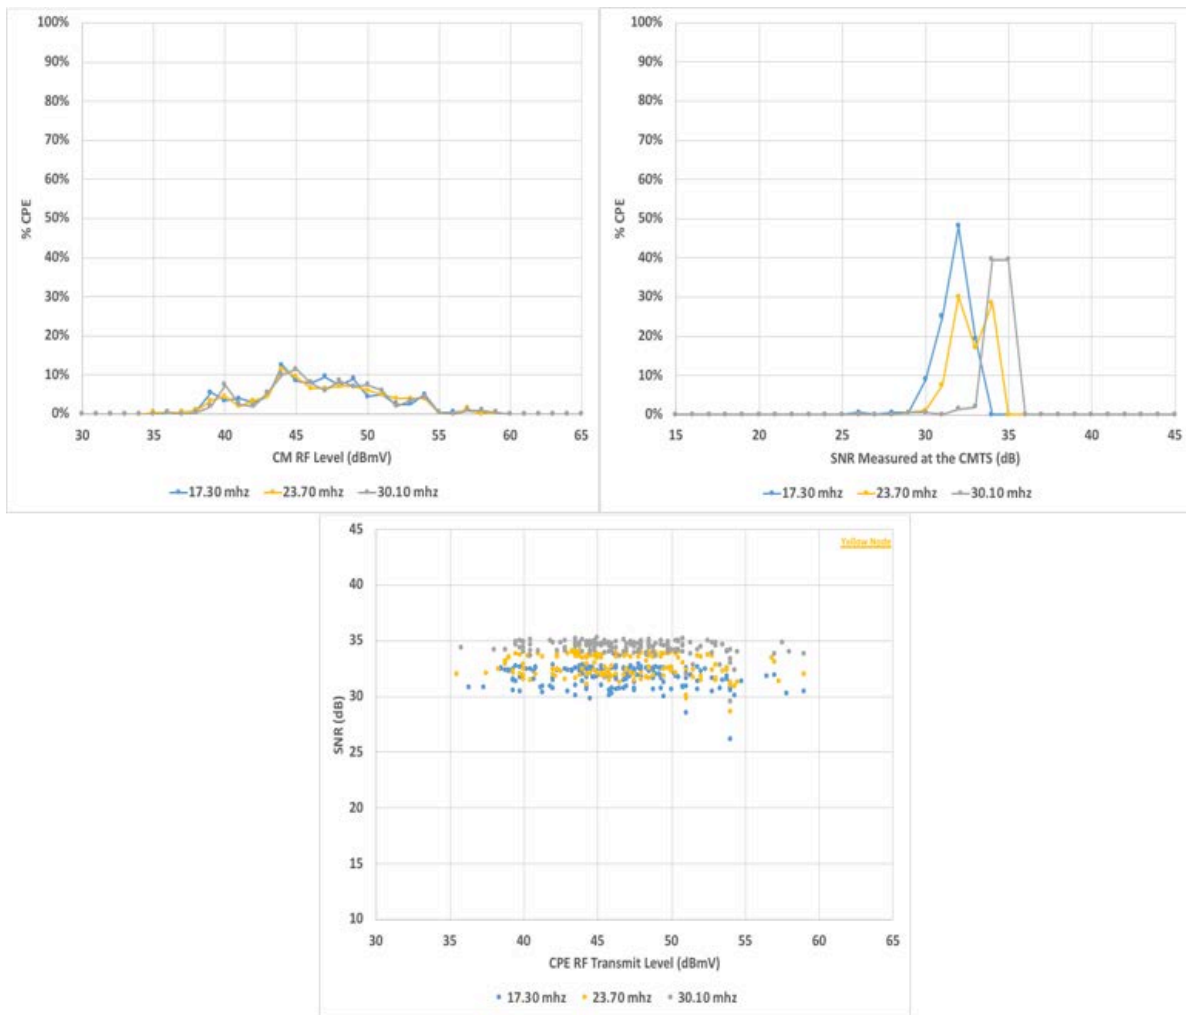


Figure 15 - Yellow Node with 3 US channels CM TX Power and SNR

A persistently wide distribution of SNR could be an indication of transient effects, which we will be able to more clearly in the next set of figures that represent a Red node. In Figure 16 the SNR is very widespread and one of the channels has a very poor SNR value, possibly impacting reliable transfer of information. A loss of code-word-errors is especially problematic for TCP/IP throughput, and is likely to impact customer experience negatively.

A look at the scatter plot indicates how profound and widespread this effect is, and gives us an indication of amounts of ingress in the system. The specific impairments on this node include the impact of the CMD discussed earlier.

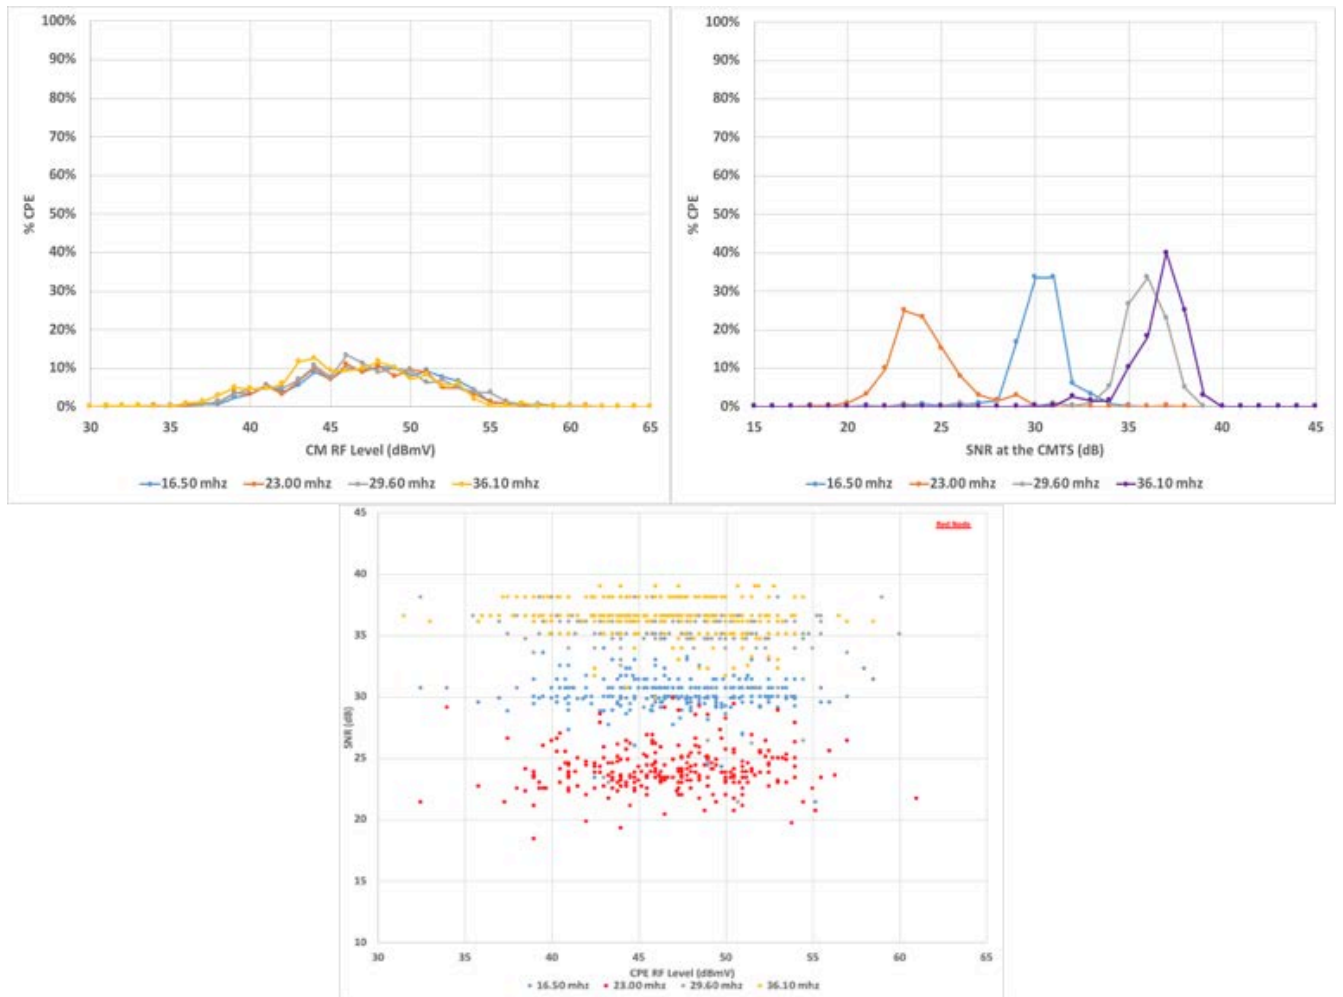


Figure 16 - Red node with 4 US channels CM TX Power and SNR

While the above analysis above is just a snapshot and has taken just three nodes into account, Comcast supports analytics for hundreds of thousands of nodes in continuous operation, spread over three divisions that pass 55M households (HHP). The challenge is to take this analysis, add additional metrics and scale it on dashboards.

2.1. Leading and Lagging Indicators

Customer experience, while vital, is a lagging indicator for our network. Rather than risk customer experience, it is important to gather sufficient metrics to understand some leading indicators, and use those to improve the customer experience before it degrades.

A preliminary form of data analysis could acquire the sets of data from multiple SOTs and arrange them so that the inter-relationships between multiple parameters is easier to project. A cross correlation matrix can be constructed as shown below in Figure 17.

US					
CM Tx	dBmV	Mean/Max/Min/SD	f1, f2, f3, f4	For each MAC address	
CM RX	dBmV	Mean/Max/Min/SD	f1, f2, f3, f4	For each MAC address	
CM MER	dB	Mean/Max/Min/SD	f1, f2, f3, f4	For each MAC address	
CM cCER	%	Mean/Max/Min/SD	f1, f2, f3, f4	For each MAC address	
CM ucCER	%	Mean/Max/Min/SD	f1, f2, f3, f4	For each MAC address	
DS					
CM MER	dB	Mean/Max/Min/SD	f1 fn	For each MAC address	
CM RF level	dBmV	Mean/Max/Min/SD	f1 fn	For each MAC address	
Mean/Max/Min/SD for each f1, f2, f3, f4 US and f1 ... fn DS across all the MACs in the Node					
Cross Correlation Matrix					
	CM Tx	CM RX	CM MER	CM cCER	CM ucCER
CM Tx		%	%	%	%
CM RX			%	%	%
CM MER				%	%
CM cCER					%
CM ucCER					

Figure 17 - Network performance cross-correlation matrix

For example, in the above set that can be measured continuously in real time or at least several times a day, a smaller standard deviation (SD) for all the parameters across the day, or of the SD across the multiple devices, would mean a tighter control of the network. However, a tighter control of one parameter, with a relatively higher variance in other parameters would be akin to transients that would then need to be brought under control, all of these being leading indicators of network health.

Of the metrics, the relationship between MER/SNR and the CCER or CER is the most an interesting relationship to help diagnose RF problems. For one thing, acquiring CER data at the resolution of a single MAC address can be difficult to collect, especially at scale. But its relationship to the MER, if tight and well behaved, would give us great confidence of the lack of non-linearities in any part of our system. This fact can be gleaned by a cross-correlation matrix of the kind described above.

Figure 18 illustrates a preliminary example of a dashboard that could track a large number of variables, on a node-by-node basis, drilled down to the MAC addresses along with a correlation matrix:

17Mhz						23Mhz					
17Mhz_rx	17Mhz_tx	17Mhz_mer	17Mhz_coor	17Mhz_uncor		23Mhz_rx	23Mhz_tx	23Mhz_mer	23Mhz_coor	23Mhz_uncor	
-0.06	40.12	34.36	0.12	0.03		-0.02	43.98	35.1	0.12	0.03	
null	null	null	null	null		0.06	50.58	34.8	0.12	0	
-0.04	42.28	34.42	0.3	0.01		-0.06	42.02	34.94	0.3	0.01	
-0.1	48.66	34.04	0.08	0		-0.08	49.24	34.64	0.08	0	
0.02	40.76	34.4	0.27	0.11		0	41.82	34.86	0.27	0.11	
0.02	47.22	34.34	0.1	0.1		-0.06	48.12	34.9	0.1	0.1	
-2.36	51	32.4	0.3	0.08		-0.04	42.12	34.86	0.31	0.11	
-0.06	39.42	34.32	0.31	0.11		-0.46	48.52	34.3	0.12	0.11	
-0.1	48.42	33.22	1.43	0.01		-1.12	51	33.56	0.3	0.08	
-0.28	47.66	33.96	0.12	0.11		-0.06	46.98	34.52	1.43	0.01	
-0.12	46.44	34.06	0.11	0		0	45.58	34.86	0.23	0.02	
0.04	44.88	34.32	0.23	0.02		-0.08	48.04	34.72	0.11	0	
0	42.96	34.28	0.33	0.06		-0.02	43.46	34.94	0.33	0.06	
0.08	35.88	34.36	0.23	0.03		0.02	36.18	34.9	0.25	0.03	
-0.02	47.12	34.28	0.06	0.02		0	47.88	34.82	0.06	0.02	
-5.24	51	29.44	0.18	2.51		-4.88	51	30.18	0.18	2.51	
-0.68	40.38	34.18	0.07	0.02		-0.56	40.94	34.78	0.07	0.02	
summary	17Mhz_rx	17Mhz_tx	17Mhz_mer	17Mhz_coor	17Mhz_uncor	summary	23Mhz_rx	23Mhz_tx	23Mhz_mer	23Mhz_coor	23Mhz_uncor
count	257	257	257	256	256	count	260	260	260	259	259
mean	-0.1	44.58	33.09	0.21	0.05	mean	-0.12	43.21	33.78	0.21	0.05
stddev	0.52	4.85	5.38	0.19	0.16	stddev	0.64	4.54	4.83	0.18	0.16
min	-5.24	32.24	0	0	0	min	-4.28	33.52	0	0	0
max	1.22	54	34.78	1.43	2.51	max	1.8	54	35.54	1.43	2.51
MATRIX	Rx	Tx	Mer	coor	uncor	MATRIX	Rx	Tx	Mer	coor	uncor
Rx	1	-16.87	1.27	-4.37	-48.69	Rx	1	-15.13	5.81	-1.29	-43.79
Tx	-16.87	1	62.84	19.06	9.53	Tx	-15.13	1	56.7	14.47	8.59
Mer	1.27	62.84	1	21.69	3.53	Mer	5.81	56.7	1	13.49	2.27
coor	-4.37	19.06	21.69	1	5.18	coor	-1.29	14.47	13.49	1	4.9
uncor	-48.69	9.53	3.53	5.18	1	uncor	-43.79	8.59	2.27	4.9	1
total_macs	rx_counts	tx_counts	mer_counts	coor_counts	uncor_counts	total_macs	rx_counts	tx_counts	mer_counts	coor_counts	uncor_counts
0.94	0.94	0.94	0.94	0.94	0.94	0.95	0.95	0.95	0.95	0.95	0.95
30Mhz						36Mhz					
30Mhz_rx	30Mhz_tx	30Mhz_mer	30Mhz_coor	30Mhz_uncor		36Mhz_rx	36Mhz_tx	36Mhz_mer	36Mhz_coor	36Mhz_uncor	
0.08	42.06	34.3	0.12	0.03		0.04	42.32	34.7	0.12	0.03	
null	null	null	null	null		null	null	null	null	null	
0.02	41.96	34.36	0.3	0.01		0.02	42.02	34.86	0.3	0.01	
-0.06	48.94	34.08	0.08	0		-0.1	48.78	34.72	0.08	0	
0.02	42.02	34.44	0.27	0.11		0.04	42.52	34.84	0.27	0.11	
0	48.38	34.36	0.1	0.1		0	48.92	34.96	0.1	0.1	
-0.04	40.34	34.4	0.31	0.11		-0.02	40.48	34.86	0.31	0.11	
-0.38	48.72	33.76	0.12	0.11		-0.54	49.52	34.34	0.12	0.11	
-0.86	50.96	33.12	0.3	0.08		-0.42	50.82	33.82	0.3	0.08	
-0.04	47.68	35	1.43	0.01		-0.08	47.18	35.04	1.43	0.01	
-0.06	45.52	34.22	0.23	0.02		0	45.98	34.74	0.23	0.02	
-0.04	48.06	34.06	0.11	0		-0.02	49.12	34.76	0.11	0	
-0.04	43.62	34.24	0.33	0.06		0.02	44.02	34.5	0.33	0.06	
0	36.26	34.36	0.25	0.03		0.08	36.82	34.82	0.25	0.03	
0	48.84	34.58	0.06	0.02		-0.02	49.38	34.36	0.06	0.02	
-4.12	51	30.72	0.18	2.51		-4.28	51	31.16	0.18	2.51	
-0.22	40.86	34.08	0.07	0.02		-0.42	41.48	34.56	0.07	0.02	
0.04	46.88	34.4	0.54	0.03		0	47.12	34.68	0.54	0.03	
-0.02	42.24	33.78	0.06	0		-0.08	42.42	34.56	0.06	0	
-0.86	46.9	34.24	0.51	0.14		-0.2	46.52	34.6	0.51	0.14	
summary	30Mhz_rx	30Mhz_tx	30Mhz_mer	30Mhz_coor	30Mhz_uncor	summary	36Mhz_rx	36Mhz_tx	36Mhz_mer	36Mhz_coor	36Mhz_uncor
count	260	260	259	259	259	count	262	262	262	261	261
mean	-0.12	45.28	33.91	0.21	0.05	mean	-0.13	45.61	34.4	0.21	0.05
stddev	0.5	4.53	1.96	0.19	0.16	stddev	0.64	4.52	1.49	0.19	0.16
min	-5.04	34.04	13.22	0	0	min	-5.3	34.12	13.5	0	0
max	1.94	57	35.04	1.43	2.51	max	1.62	57	36	1.43	2.51

Figure 18 - Cross-correlation matrix dashboard

At the current time, various metrics of the above dashboard are available from different SOTs. An effort that consolidates information from all the diverse SOTs would establish the internal consistency of our data acquisition and also enable visibility of key leading indicators.

Conclusion

The rise of CMD noise, coincident with the rise of switched power supplies in consumer CPE and other in-home electronics, catalyzed the identification and mitigation framework described in this paper. The framework consists of:

- 1) Identifying operational challenges in the field, finding the root cause and characterizing the noise
- 2) Applying detailed lab characterizations to effectively understand the options for mitigation, which enabled the design that assuages the impacts on customer experience
- 3) Documenting and modeling a solution to the CMD challenge
- 4) Field trialing and deploying the solution and verifying its efficacy

This full cycle example of the CMD mitigation framework demonstrates an opportunity to simultaneously mitigate a network issue and enable better operational efficiencies for plant maintenance. In the example of CMD noise and advanced analytics of CM network data the customer experience was improved with increased operational efficiency.

This paper describes a discovery and development process to identify opportunities to improve network performance, develop solutions and to use data analytics to validate the improvements. This paper demonstrates that ... *what gets measured gets done, and what gets analyzed gets transformed.*

Abbreviations

AC	Alternating Current
bps	bits per second
CCER	Correctable Codeword Error Ratio
CDN	Coupling Decoupling Network
CER	Uncorrectable Codeword Error Ratio
CM	Cable Modem
CMD	Common Mode Disturbance
CMTS	Cable Modem Termination System
CPE	Consumer Premise Equipment
dB	Decibel
DC	Direct Current
DOCSIS	Data over Cable Systems Interface Specification
DS	Downstream
FEC	Forward Error Correction
HFC	Hybrid Fiber-Coax
Hz	Hertz
ISBE	International Society of Broadband Experts
IUC	Interval Usage Code
LED	Light Emitting Diode
MAC	Media Access Control
MER	Modulation Error Ratio
MHz	1x10 ⁶ Hz
PoC	Proof of Concept
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RX	Receive
SCTE	Society of Cable Telecommunications Engineers
SIK	Self-Install-Kit
SNR	Signal to Noise Ratio
SOT	Source of Truth
STB	Set top Box
T&M	Test and Measurement
TCP	Transmission Control Protocol
TX	Transmit
UDP	User Datagram Protocol
UGS	Unsolicited Grant Service
US	Upstream

Bibliography & References

See footnotes

When Security and Privacy Collide

New Approaches are Needed

A Technical Paper prepared for SCTE•ISBE by

Sandy Wilbourn

VP Engineering

Akamai

Santa Clara, CA

+1 650 381 6129

rwilbourn@akamai.com

Craig Sprosts

Senior Director, Product Management

Akamai

Santa Clara, CA

+1 650 381 6043

csprosts@akamai.com

Table of Contents

Title	Page Number
Table of Contents	2
Introduction.....	3
Content.....	4
1. Security Research Today	4
2. Changing Privacy Landscape	4
3. Generating Security Insights	5
4. Gathering Network Data.....	6
5. Protecting Privacy	6
6. Building a Layered, Intelligent Processing System	6
7. Discovering New Core Domains	7
8. Adapting Natural Language Processing to Domain Names	9
9. Evaluating Quarantined Domains	10
10. Visualizing Security Data	10
Conclusion.....	11

List of Figures

Title	Page Number
Figure 1 - A read/write in-memory processing engine	8
Figure 2 - Output of a near real time processing engine evaluating live streamed DNS resolution traffic	9
Figure 3 - Two-dimensional visualization of results from a clustering engine	11

Introduction

Extensive publicity about gathering and use of personal data by popular online services has increased privacy concerns, especially in developed countries. This has led to consideration and passage of privacy regulations in many parts of the world. The most visible example is the European Union's General Data Protection Regulations (GDPR) which define a new regulatory framework for the management of personal data. These regulations are colliding in unexpected ways with Internet Service Providers desire to protect their subscribers from malicious activity.

As the May 2018 deadline for implementation of the GDPR regulations drew closer security researchers realized there was potential impact on the use of the *whois* database that stores data about domain name registrations. The *whois* database was widely used for security research because it contains useful information about domain name registrations like who is registering the name, their contact information (email), location and more. Since use of domain names is fundamental to activating and maintaining most security exploits, data about their heritage is useful.

The International Corporation for Assigned Names and Numbers (ICANN), the organization responsible for administering the *whois* database, has defined a temporarily specification that pares back data fields in *whois* significantly so it is compatible with GDPR. Information that's been useful for security research in the past isn't available. ICANN has convened a group to develop a long-term solution but it's not clear where it will lead.

There are other examples in the past where proposed privacy regulations had the potential to impair security research by limiting availability of data. In early 2016 the United States Federal Communications Commission began to formulate regulations that would have restricted gathering of various kinds of network data. In this case the industry and research community collaborated and advocated for revisions that would ensure privacy while allowing for capture and use of properly anonymized network data.

It's inevitable collisions between privacy and security will continue to occur. Solving the problem of diminishing data availability means security researchers have to maximize the utility of security data that remains. Security research will need to move from rigid, deterministic, and rule-based, where personal information was helpful; to behavioral, anomalies-based analysis across very large volumes of anonymized data. The future calls for overlaying multiple layers of data where no single layer produces a result.

This will require highly automated processing and machine learning. Advanced algorithms can expand coverage of activity related to known threats, and discover previously unknown attacks, without compromising precision (generating false positives). High-performance processing of real-time data can also improve agility, or how quickly threats are found. There's also the possibility of reducing research costs by extending the efforts of human experts with machines.

This paper will cover a recent example of privacy regulation impacting security research by outlining the issues that led to the *whois* problem and compliance with GDPR. It will then discuss a way forward: applying modern data processing techniques to large data sets to expand threat coverage, improve precision, and increase agility. A production machine learning system that analyzes live streamed, anonymized, DNS data gathered from DNS resolvers serving active Internet users all over the world will be described, along with the results it can generate.

Content

1. Security Research Today

Security researchers evaluate data to find anomalies, and then cross check or validate their findings against potentially many other data source to determine whether or not a threat exists. Then they characterize the threat and publish mechanisms to deter it. Researchers use a variety of data sources to conduct their work:

- Honeypots mimic systems like mail and web servers, databases, or other services that are commonly attacked, hoping they'll be perceived as legitimate targets so malware can be captured.
- Offline data sources track ownership details for Internet resources like web hosting and cloud services, or registration data for domain names.
- In some cases other network data collected from traces or scans might be used.

The ethics of security research is a sub-field in itself, and although privacy has always been a consideration, guidelines were often based on the policies and processes of the organization doing the research since formalized privacy regulations were not necessarily directly relevant. That's beginning to change with the advent of the European Union's (EU) General Data Protection Regulations (GDPR).

2. Changing Privacy Landscape

Data privacy became a visible issue many years ago in the EU, with citizens expressing serious concerns regarding the use of their personal information by online services. After several years of debate regulations were approved by the European Parliament in April 2016 and a two-year transition period for organizations to reach compliance was established, ending in May 2018. Significant fines will be levied for businesses that don't follow GDPR guidelines. Recognizing this organizations began to assess the data they were collecting through the lens of GDPR.

The International Corporation for Assigned Names and Numbers (ICANN) is the organization responsible for administering the whois database, which contains information about domain name registrations. Whois data has been an important tool for security researchers and law enforcement agencies investigating malicious activity like phishing, malware sites, botnets and many other kinds of online crime. Information about the heritage of domain names is useful to security research because domain names are so fundamental to most exploits. More on this topic in the next section of this paper.

In late 2017 ICANN publicized a memo that concluded *whois* needed to be restructured to be in compliance with GDPR.¹ In the short term they defined a temporary specification that removed data fields containing personal information² like contact details, leaving only basic information like organization name, state or province, and country. They also created a multi-stakeholder working group to define a long term solution to *whois* compliance with GDPR that would meet the needs of everyone who uses the data.³ It's at best unclear how this effort will turn out, given the complexity of the issues.⁴

¹ <https://www.icann.org/en/system/files/files/gdpr-memorandum-part2-18dec17-en.pdf>

² <https://www.icann.org/en/system/files/files/proposed-gtld-registration-data-temp-specs-11may18-en.pdf>

³ <https://www.icann.org/news/announcement-2018-07-02-en>

⁴ <https://www.internetgovernance.org/2018/07/03/stacking-the-deck-the-epdp-on-the-whois-temp-spec/>

A solution for the *whois* situation may be found, but heightened interest in privacy makes it likely more regulations will be implemented and it's probable other kinds of security data will also become inaccessible or obscured in ways that make it less useful. This, and widespread use of encryption, will make keeping pace with the volume and sophistication of today's security threats harder. New approaches are needed.

3. Generating Security Insights

Over the long-term balancing security research needs and privacy requirements will require making better use of available data. This means moving from rigid, deterministic, rule-based security, where personal information was helpful, to behavioral anomalies-based analysis across large volumes of data. The future calls for overlaying multiple layers of data where no single layer produces a result. Machine learning and other kinds of data processing are needed to identify increasingly sophisticated threats. Advanced algorithms can find more threats with less data, without compromising precision. New techniques can expand coverage of activity related to known threats and discover previously unknown attacks. Agility - how quickly threats are found - also improves, and research costs are reduced by extending the efforts of human experts.

A security data source that's starting to get a lot of attention is DNS resolution data sourced from resolvers or various kinds of network taps. DNS resolution data has a number of useful characteristics for security research. Domain names, and the Domain Name System (DNS) authorities and resolvers that support them, are fundamental to most security exploits. The DNS is widely used by malware developers because it connects everything on the internet, from anywhere. Virtually every network and device where an exploit might be activated will have access to the DNS. Conversely, any device that emits a DNS query known to be associated with associated malware.

The DNS has scaled remarkably as the Internet has grown and there's considerable infrastructure and tools for managing domain names. This enables highly dynamic connectivity, so exploits can move and change rapidly to avoid detection or takedowns. Malware developers can use a domain generation algorithm (DGA) to create an endless supply of random names to obfuscate their exploits. They only pay a modest fee to register the small percentage of domain names they actively use to enable their exploits. The use of DGAs is explored in depth in the paper: "A Comprehensive Study of Domain Generating Malware".⁵

From a practical standpoint DNS queries also tend to be one of the first steps in enabling malware on a host to function. A DNS query sent from a device to a known malicious destination indicates the device is associated with malicious activity, it's also usually the first "signal" that's visible on a network where it can be detected remotely. Identifying activity at this stage is extremely useful as an exploit can potentially be disrupted before it does any real damage.

This agility aspect of DNS data (and the value of DNS data more broadly for security research) was discussed in a widely publicized academic paper: "*A Lustrum of Malware Network Communication: Evolution and Insights*".⁶ The paper states: "We find that a significant percentage of malware domains can be seen in passive DNS several weeks, in many cases even months, before the actual malware sample was dynamically analyzed by the security community."

For completeness, malware developers have alternatives to the DNS. It's possible to code static IP addresses into exploits but once the address is discovered it's easy to block or takedown. Proprietary

⁵ <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/plohmann>

⁶ <https://www.computer.org/csdl/proceedings/sp/2017/5533/00/07958610.pdf>

protocols can also be created to facilitate communications and management, but it requires effort to implement and maintain them, and once they're discovered they can be blocked. DNS remains the only viable choice for simple, dynamic connectivity.

4. Gathering Network Data

A data science team at Akamai began processing DNS resolution data several years ago to detect and track malicious activity. The data is from diverse worldwide sources and live streamed 24x7. It's transported over a redundant network to multiple data centers that contain parallel, intelligent processing systems the team has developed so incoming data can be evaluated in near real time.

Obtaining data is always challenging because it involves extra effort on the part of the contributors. Fortunately most resolvers can be equipped with facilities to capture query data and ship it off to other systems. There's usually a cost in terms of query performance, but it can be modest with an efficient implementation for copying query data and sending it off the server. Service providers supplying data also need to provision links to transport the data to the Akamai data centers where it will be processed as well.

5. Protecting Privacy

User privacy has always been a consideration, even in the absence of regulations. Another advantage of using DNS queries gathered from resolvers as a security data source is it's minimally invasive of privacy. Unlike technologies that promiscuously gather and evaluate traffic in the data plane, DNS queries only contain source/destination IP addresses and domain name related data. Personally Identifiable Information (PII) like IP addresses can be anonymized so that it cannot be traced to an individual. This topic will be addressed in the next section.

Data used for research at Akamai is anonymized with the Lucent extension to Crypto-PAn, a well-known cryptography-based sanitization tool for anonymizing IP addresses. Service providers who own the resolvers control all aspects of the anonymization of their query data; they configure which potential PII is anonymized and create and manage the anonymization keys. A third party cannot reverse the anonymization, only the provider can, using keys they generate. They use one key for anonymizing all of the data in their network which is a bare passphrase consisting of any ASCII text, on any number of lines.

Data is also encrypted in transit. This requires provider systems to initiate secure connections to the destination servers. OpenSSL has proven to be a good solution. Connections from the provider network to Akamai servers are authenticated by looking up a host specified in an authtoken file supplied by Akamai, connecting to the host using TLS, and exchanging and verifying certificates.

6. Building a Layered, Intelligent Processing System

The team set goals of improving threat coverage and precision by applying intelligent processing to the DNS data. Another objective was to do all of the processing in near real time, so threats could be identified, validated, and published as quickly as possible. This has led to development of a number of systems for intelligent processing, summarized below and described in detail in the following sections.

- Preprocess the data to reduce noise so more processing power can be applied to data of interest
- Map relationships between domain names
- Assign domain reputation scores by joining with other data sources to evaluate "maliciousness"
- Expand coverage using techniques similar to natural language processing
- Correlate relationships between malicious domain names

- Visualize clusters using 2D and 3D graphs to better understand relationships

Each system operates as a separate “layer”, with each adding intelligence to the findings of others. In most cases no single layer offers conclusive evidence that a domain name is malicious, instead they all work together to formulate conclusions. In effect the network, or more accurately streamed network data, looks like a massive, extremely diverse, near real-time honeypot.

7. Discovering New Core Domains

One of the earliest revelations of the research was newly observed domain names tend to be more highly correlated with malicious activity. This makes sense intuitively because malware developers need to constantly change the face of their exploits to avoid detection and take down. One of the ways they do this is to constantly change the domain names associated with their exploits.

In creating a methodical approach for studying newly observed domain names we defined the concept of a “core” domain, which is also known as an “effective 2nd level domain” (e2LD). For instance:

[www.example1.com](#) and [www.example2.co.uk](#) are core domains or e2LDs. It can be seen that core domains usually capture domain ownership. For the past 5 years Akamai researchers have been tracking new core domains, essentially newly observed domain names, and in 2017 undertook a project to greatly improve the infrastructure in order to study them more intensively. Details of the new core domain work were presented at a DNS conference in 2017.⁷

The team developed a read/write in-memory processing engine that was capable of operating on a 1.5 million QPS data stream (scalable as the data stream grows). This engine was designed to enable real time processing to reduce noise in the data and evaluate the relevance of each query. Algorithms also flag other kinds of anomalous behavior, such as incoming queries for domains with query patterns that substantially differ from previous patterns. This engine effectively detects potential phishing, bot and other malware activity, DNS based DDoS attacks, and DNS tunnels.

⁷ <https://indico.dns-oarc.net/event/27/contributions/456/>

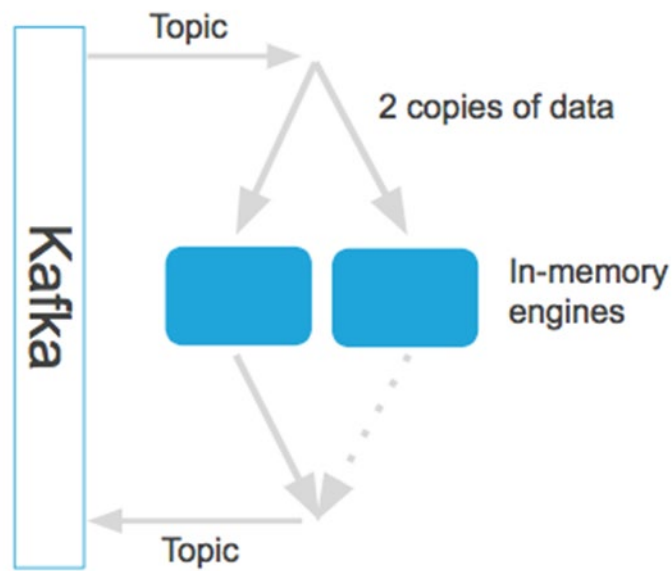


Figure 1 - A read/write in-memory processing engine

A read/write in-memory processing engine processes live streamed DNS queries at 1.5 million queries per second. This engine reduces noise in the data and evaluates the relevance of each query.

Output of the processing can be seen in Figure 2 below. A dashboard displays a number of statistics such as total queries processed, new core domains found, queries to new core domains that resolve and don't resolve (return an error code). The resolution status is represented in dark blue (resolved) vs light-blue (not resolved). Seeing the resolution status, including the answer itself, is useful because it turns out domains that are not registered are frequently used by botnets through DGA (domain generation algorithms). While blocking such domains may or may not help, there is still value in identifying an infected machine even if a query did not resolve.

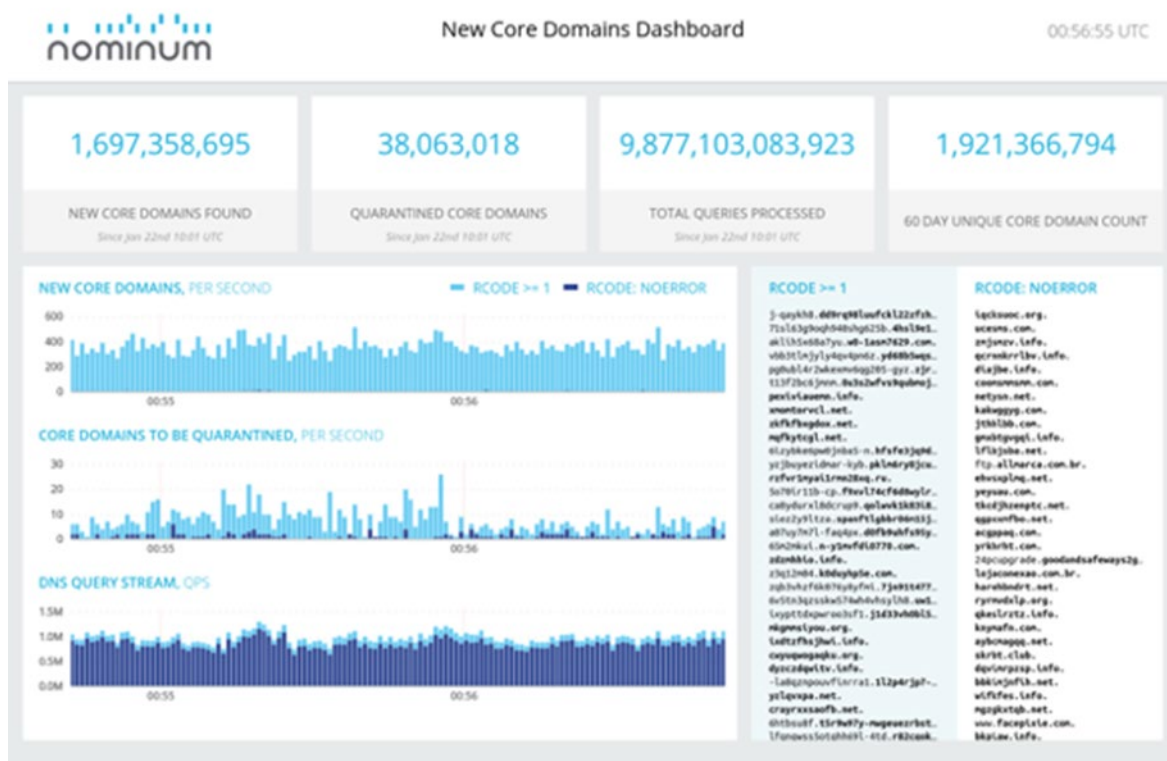


Figure 2 - Output of a near real time processing engine evaluating live streamed DNS resolution traffic

Output of a near real time processing engine evaluating live streamed DNS resolution traffic to find newly observed domain names. Domains discovered by this engine are subject to additional processing to validate their maliciousness and characterize their behavior.

Domains that will be quarantined are also displayed. These are the targets for more processing to determine whether or not they are actually malicious. Subsequent processing also reveals their intent so they can be categorized.

New core domains provide useful information about other vectors of attacks. For example, an uptick in the use of social media for distributing malware was uncovered. In this case waves of free airline ticket promotions by airlines. The domain names for the promotion used alternate character sets that looked like ordinary roman characters in order to trick users. The domain name in the url displayed a subtly different character but the actual domain name seen by the resolver was much different and easily detected by the new core domain logic.

8. Adapting Natural Language Processing to Domain Names

Security list providers catch some of the domain names an exploit uses from honeypots, but they typically don't capture all of the names in use. Malware can also include anti-honeypot techniques to fool the honeypot, for instance `pykspa` uses real and fake DGAs to confuse the honeypot output.

To expand coverage of malicious activity generated from the other layers in the system additional techniques borrowed from natural language processing are used to reveal relationships among seemingly random domain names and clients that query them. The model borrows concepts from the word2vec work

done at Google.⁸ Quarantined domain streams are fed into the model and it generates clusters which group the most correlated names together. The model applies an advanced neural network structure onto the original word2vec neural network by modeling the DNS query sequence and discovering the in-depth correlation among domain names in a massive DNS traffic stream.

9. Evaluating Quarantined Domains

Clusters that are discovered are validated using 3rd party security lists, typically generated by human researchers, which include malware C&C, malvertising, phishing, etc. Relationships between domains are mapped (analogous to a social graph) so likely neighboring domains that are malicious can be propagated. Algorithms overseeing these layers of guilt by association generate a Domain Reputation Score that categorizes domains to be designated as malicious or those worthy of even more analysis.

Measurements calculated by the research team showed propagating human security intelligence to clusters discovered with machine learning can expand coverage by 5x to 10x. To the point made earlier about the agility of DNS resolution data, malicious clusters are also regularly identified which didn't appear until hours or days later on 3rd party threat lists.

Malicious or suspicious domain names discovered in data Akamai collects are stored in a reputation knowledge-base. Continuous improvements to this database make associated machine-learning systems faster and more accurate so more malware can be effectively blocked before it causes damage.

Looking at the output of the algorithms for individual threats and then doing a deeper dive to see what other kinds of patterns emerge always offers interesting insight. For example, additional analysis of the machines that emitted the kill-switch domain for Wannacry showed there was a significant correlation with gaming use of those machines and TeamViewer, a tool for remote administration. This makes sense since leaving certain ports open increases the likelihood exploits will get into unpatched systems.

As another example, an evaluation of Petya's time sequence showed it took the dropper exactly 2 minutes from the time it was downloaded until it started querying the payload site. Only a couple of minutes! AV for these infected users, assuming there was one installed, didn't catch the dropper file. Instead it allowed it to install itself, and then make a query to the payload site.

10. Visualizing Security Data

Continuous improvements in graphing technology allow better visualization of threat activity. Results calculated using the correlation techniques above are fed into a model that groups the most correlated domain names together into clusters and places them on special 2D and 3D graphs so their relationships can be better understood. An example of the graphs that can be generated are shown in the figure below.

⁸ <https://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>




Figure 3 - Two-dimensional visualization of results from a clustering engine

Two-dimensional visualization of results from a clustering engine that uses unsupervised machine learning to correlate characteristics that reflect association with a common underlying threat.

Conclusion

It's inevitable collisions between privacy and security will continue to occur. This will make maintaining security in a privacy driven world harder. The telecommunications industry needs to monitor regulatory initiatives worldwide and advocate for policies that preserve privacy but allow for gathering and use of data used for security research. Even with advocacy, security data sources will continue to be obscured or blocked altogether, and with stiff penalties for privacy violations creators of data are likely to become more cautious so raw data will be less available, more opaque, and generally less useful. Yet malware developers won't relent, in some cases they've already implemented with multi-faceted exploits that evade defenses and propagate rapidly. Maintaining an edge will require making the most of available data sources.

Almost all threats have a footprint in the DNS and analysis of query traffic captured by DNS resolvers can provide early detection of malware since exploits have to resolve addresses of malicious resources under their control before they can use the functions they rely on to operate and propagate. Evaluating DNS query data also offers the possibility of improving coverage of diverse communications channels malware uses. Human driven security research will always be necessary but intelligent processing and machine learning will become essential tools that strongly complement agile, rich and diverse DNS data.

AUTHOR INDEX
2018 Technical Paper Proceedings

Aberastury, Marcos	384	Cooper, Kyle.....	1009
Ames, Rob	1009	Cooper, Michael	1051
Anand, Arvinder S.....	1361	Dalton, Curt	93
Andreoli-Fang, Jennifer.....	984	Day, Chris.....	1072
Ansari, Furquan	703	De Arca, Florencia.....	1221
Assylbekov, Shamil	1105	Delbar, Jos.....	539, 642
Baldry, Jon.....	1, 349	Dhawan, Sanjay	787
Bastian, Chris	821	Diaz, Gaston	384
Beesley, Bill	935	Dodd-Noble, Aeneas	984
Begen, Ali C.	442	Dorigo, Francesco.....	1275
Belt, Dave	1159	Fay, Michael	858
Bernstein, Alon.....	984, 1124	Feehan, Leonard	1455
Boone, Ted	1009	Flesch, J.R.	1411
Brooks, Roger.....	1148	Fiorenzo, Mariela.....	1221
Burg, Bernard	1480	Foroughi, Nader.....	682
Callesen, Uffe	886	Gantt, Kevin	1072
Campos, L. Alberto	359, 766, 1168	Gaydos, Robert.....	1037
Carro, Gabriel.....	1221	Ghai, Rajat.....	216, 465
Chalapati, Satish	295	Ghuman, Harj	272
Chapman, John T.....	620	Gibellini, Emilia	1221
Cheevers, Charles	560, 720, 867, 1411	Giladi, Alex	442
Chen, Tianwen.....	1480	Glapa, Martin J.	82
Cho, Junho	807	Goemaere, Patrick	465
Colby, Andrew	1148	Goeringer, Steve	1313
Conklin, Tom.....	1266	Goodwin, David	560
Coomans, Werner	807	Gopal, Vishnu.....	48

Grimaldi, Adrian.....	384	Krishnan, Yadhav	703
Gronvall, Erik	186	Kumar, Pankaj	1148
Gutknecht, Gary.....	956	Kuykendall, Peter	48
Harmath, Norberto.....	384	Laufer, Tal	886
Heaton, Eric	148, 321	Lavallée, Brian	1027
Helms, K. Scott.....	1328	Lefevre, Yannick	807
Hering, David	921	Levy, Devin	1105
Hernandez, Miguel	703	Liu, Fan.....	1480
Holobinko, John	117, 1072	Liu, Tong	252
Howald, Rob	1072	Loeffelholz, Todd	1072, 1112
Howard, Daniel.....	1072	Mahajan, Pravin.....	799
Howell, Daniel.....	1275	Malhotra, Anant.....	1148
Iannone, Patrick.....	807	March, Ryan	1480
Ih, Ron	839	Marut, Dan.....	1072
Jain, Mudit.....	1148	Masoud, Fady	1017
Jia, ZhenSheng (Steve).....	359, 766, 1168	Matatyaou, Asaf.....	1135
Jin, Hang.....	620	Mattingly, Martin.....	194
Job, David.....	272	May, Bradley	1292
Johnson, Tim	821	McGilvray, Glenn.....	117
Joseph, Jean-Philippe.....	82, 1385	McKibben, Bernard	984
Judge, David	921	McLeod, Bruce	1467
Justis, Colin	1495	Menon, Narayan	667
Kannan, Nav	720	Metsch, Thijs	1455
Katiyar, Sandeep.....	336, 1371	Miles, Kathleen.....	1072
Kipp, Neill A.	1250	Mobley, Michael.....	93
Kirsche, Dick	1072	Morley, Dave	13
Kloberdans, Michael.....	131	Mukhopadhyay, Amit.....	1385
Knittle, Curtis	359, 766	Mulqueen, Kieran	1455

Mutalik, Venk.....	1511	Ryan, Brendan	1455
Nair, Raj	1299	Singh, Amit.....	321
Navali, Prabhu	1299	Solomon, Joe	1037
Nguyen, Bao	1275	Spee, Rene	1072
Nicholson, Greg.....	168	Spoczynski, Marcin	1455
O'Dell, Michael.....	48	Sprosts, Craig	1532
O'Hanlon, Michael.....	148, 1455	Srinivasa, Sunil.....	1480
Ovadia, Shlomo	131	Stevens, J. Clarke.....	821
Panciera Molanes, Eduardo M.....	384	Stoneback, Dean	1072
Parayil, Shiby	528	Stratton, Mark.....	517
Patel, Jignesh	1009	Subramanya, Karthik	1511
Patel, Mehul.....	1037	Sundaresan, Karthik.....	587
Peck, Tobias	33	Sundelin, Andrew	506
Pickering, Ladan.....	935	Syed, Yasser	442
Pinckernell, Nick	49	Topazi, Christopher	1051
Putzeys, Jeroen	886	Ulm, John	1072
Quinn, Ruth	1455	Urrutia-Valdez, Carlos.....	1385
Raman, Narayan	703	Vale, R. J.	82
Ravisankar, Arun.....	307, 821	Van Caenegem, Tom	1385
Rehman, Abdul.....	1209	van Veen, Dora.....	807
Reyes, Elias Chavarria.....	984	Vasamsetty, Chaitanya	1009
Rice, Dan	1511	Vercammen, Bart.....	539, 642
Righetti, Claudio.....	1221	Villanueva, Earl.....	528
Ritchie, John	117	Villarruel, Fernando X.....	93, 194
Roque, Abel Vilca	1480	Walker, Richard J.	1135
Rudrapatna, Ashok	1385	Walsh, Jim	921
Ruff, Chris	454	Walsh, Joe.....	942
Rupe, Jason.....	1313, 1495	Wang, Jing.....	359

Wang, Jon-en	1511
Wang, Zhou	1209
Watson, Justin.....	1148
West, Lamar	93, 1072
Whitehouse, Dan	1072
Wilbourn, Sandy.....	1532
Wolcott, Larry	48, 1168
Wong, Curt	984
Woodrich, Jason	49
Wu, Jonathan	867
Xu, Mus	766
Zhang, Haipeng	766